Complete Software Guide for SRX Series Services Gateways, Release 15.1X49-D40 (Volume 2)

Modified: 2016-04-14

Copyright © 2016, Juniper Networks, Inc.

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Complete Software Guide for SRX Series Services Gateways, Release 15.1X49-D40 (Volume 2)
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation cxliii

Documentation and Release Notes cxliii
Using the Examples in This Manual cxliii
Merging a Full Example cxliv
Merging a Snippet cxliv
Documentation Conventions cxlv
Documentation Feedback cxlvii
Requesting Technical Support cxlvii
Self-Help Online Tools and Resources cxlvii
Opening a Case with JTAC cxlviii

Guide 1 Application Layer Gateways Feature Guide for Security Devices

Part 1 Overview

Chapter 1 Introduction to ALGs 3

ALG Overview 3
Understanding Custom ALG Services 4
Understanding IPv6 ALG Support for ICMP 5
ICMP Error Messages 6
ICMP ALG Functionality 6

Part 2 Configuring Data ALGs

Chapter 2 Data ALG Types 11

Understanding Data ALG Types 11

Chapter 3 Configuring the DNS ALG 13

DNS ALG Overview 13
Understanding the IPv6 DNS ALG for Routing, NAT, and NAT-PT 14
IPv6 DNS ALG Traffic in NAT mode 14
IPv6 DNS ALG Traffic in NAT-PT mode 15
Example: Configuring the DNS ALG 16
Understanding DNS and DDNS Doctoring 19
Disabling DNS and DDNS Doctoring 23

Chapter 4 Configuring the FTP ALG 25

FTP ALG Overview 25
Understanding FTP Commands 26
PORT Command 26
PASV Command 26

Copyright © 2016, Juniper Networks, Inc.

Extended FTP Commands . . . 26
EPRT Command . . . 26
EPSV mode . . . 27
Example: Configuring the FTP ALG . . . 28
Understanding the IPv6 FTP ALG for Routing . . . 31
FTP ALG Support for IPv6 . . . 31
EPRT mode . . . 31
EPSV mode . . . 32

Chapter 5 Configuring the IKE and ESP ALG . . . 33

Understanding the ALG for IKE and ESP . . . 33
Understanding IKE and ESP ALG Operation . . . 34
Example: Configuring the IKE and ESP ALG . . . 35
Example: Enabling the IKE and ESP ALG and Setting Timeouts . . . 41

Chapter 6 Configuring the PPTP ALG . . . 45

Understanding the PPTP ALG . . . 45
Understanding IPv6 Support for the PPTP ALG . . . 45
Example: Configuring the PPTP ALG . . . 46

Chapter 7 Configuring the RPC ALG . . . 61

Understanding RPC ALGs . . . 61
Understanding Sun RPC ALGs . . . 62
Enabling Sun RPC ALGs (CLI Procedure) . . . 63
Enabling Sun RPC ALGs (J-Web Procedure) . . . 63
Customizing Sun RPC Applications (CLI Procedure) . . . 64
Understanding Sun RPC Services . . . 64
Understanding Microsoft RPC ALGs . . . 67
Configuring the Microsoft RPC ALG . . . 68
Configuring the MS-RPC ALG with a Predefined Microsoft Application . . . 68
Configuring the MS-RPC ALG with a Wildcard UUID . . . 69
Configuring the MS-RPC ALG with a Specific UUID . . . 69
Enabling Microsoft RPC ALGs (CLI Procedure) . . . 70
Enabling Microsoft RPC ALGs (J-Web Procedure) . . . 70
Verifying the Microsoft RPC ALG Tables . . . 71
Understanding Microsoft RPC Services . . . 71
Customizing Microsoft RPC Applications (CLI Procedure) . . . 73

Chapter 8 Configuring the RSH ALG . . . 75

Understanding the RSH ALG . . . 75
Example: Configuring the RSH ALG . . . 75

Chapter 9 Configuring the RTSP ALG . . . 89

Understanding the RTSP ALG . . . 89
Overview . . . 89
RTSP Modes . . . 90
Standard Mode . . . 90
Interleave Mode . . . 90
Understanding IPv6 Support for RTSP ALG . . . 91

Copyright © 2016, Juniper Networks, Inc.

Understanding RTSP ALG Messages . . . 91
RTSP Messages Format . . . 91
RTSP Methods . . . 92
RTSP Status Code . . . 92
RTSP Header . . . 93
Understanding RTSP ALG Conversation and NAT . . . 93
Example: Configuring the RTSP ALG . . . 96

Chapter 10 **Configuring the SQLNET ALG . . . 103**

Understanding the SQLNET ALG . . . 103
Example: Configuring the SQLNET ALG . . . 104

Chapter 11 **Configuring the TALK ALG . . . 117**

Understanding the TALK ALG . . . 117
Example: Configuring the TALK ALG . . . 117

Chapter 12 **Configuring the TFTP ALG . . . 131**

Understanding the TFTP ALG . . . 131
Overview . . . 131
TFTP Packets . . . 131
TFTP Session . . . 132
Understanding TFTP ALG Conversation . . . 132
Understanding IPv6 Support for the TFTP ALG . . . 133
Example: Configuring the TFTP ALG . . . 134

Part 3 Configuring VoIP ALGs

Chapter 13 **VoIP ALG Types . . . 141**

Understanding VoIP ALG Types . . . 141

Chapter 14 **Configuring VoIP Rewrite Rules . . . 143**

Understanding VoIP DSCP Rewrite Rules . . . 143
Example: Configuring VoIP DSCP Rewrite Rules . . . 144

Chapter 15 **Configuring the H.323 ALG . . . 145**

Understanding H.323 ALG . . . 145
Understanding the Avaya H.323 ALG . . . 147
Avaya H.323 ALG-Specific Features . . . 147
Call Flow Details in the Avaya H.323 ALG . . . 148
H.323 ALG Configuration Overview . . . 149
Example: Passing H.323 ALG Traffic to a Gatekeeper in the Private Zone . . . 149
Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone . . . 154
Example: Using NAT with the H.323 ALG to Enable Incoming Calls . . . 160
Example: Using NAT with the H.323 ALG to Enable Outgoing Calls . . . 168
Understanding H.323 ALG Endpoint Registration Timeouts . . . 174
Example: Setting H.323 ALG Endpoint Registration Timeouts . . . 175
Understanding H.323 ALG Media Source Port Ranges . . . 176
Example: Setting H.323 ALG Media Source Port Ranges . . . 176
Understanding H.323 ALG DoS Attack Protection . . . 177
Example: Configuring H.323 ALG DoS Attack Protection . . . 178
Understanding H.323 ALG Unknown Message Types . . . 179

Copyright © 2016, Juniper Networks, Inc.

Example: Allowing Unknown H.323 ALG Message Types 179

Chapter 16 **Configuring the MGCP ALG 183**

Understanding the MGCP ALG 183
MGCP Security 184
Entities in MGCP 184
Endpoint 185
Connection 185
Call 185
Call Agent 185
Commands 186
Response Codes 188
MGCP ALG Configuration Overview 189
Example: Configuring Media Gateways in Subscriber Homes Using MGCP ALGs 189
Example: Configuring Three-Zone ISP-Hosted Service Using MGCP ALG and NAT 196
Understanding MGCP ALG Call Duration and Timeouts 208
Example: Setting MGCP ALG Call Duration 209
Example: Setting MGCP ALG Inactive Media Timeout 210
Example: Setting MGCP ALG Transaction Timeout 211
Understanding MGCP ALG DoS Attack Protection 212
Example: Configuring MGCP ALG DoS Attack Protection 213
Understanding MGCP ALG Unknown Message Types 214
Example: Allowing Unknown MGCP ALG Message Types 214

Chapter 17 **Configuring the SCCP ALG 217**

Understanding SCCP ALGs 217
SCCP Security 218
SCCP Components 219
SCCP Client 219
Call Manager 219
Cluster 219
SCCP Transactions 220
Client Initialization 220
Client Registration 220
Call Setup 220
Media Setup 220
SCCP Version 221
SCCP Control Messages and RTP Flow 221
SCCP Messages 221
SCCP ALG Configuration Overview 223
Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone 223
Understanding SCCP ALG Inactive Media Timeouts 233
Example: Setting SCCP ALG Inactive Media Timeouts 233
Understanding SCCP ALG Unknown Message Types 234
Example: Allowing Unknown SCCP ALG Message Types 235
Understanding SCCP ALG DoS Attack Protection 236

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring SCCP ALG DoS Attack Protection . . . 236
Verifying SCCP ALG Configurations . . . 237
Verifying SCCP ALG . . . 237
Verifying SCCP ALG Calls . . . 238
Verifying SCCP ALG Call Details . . . 238
Verifying SCCP ALG Counters . . . 239

Chapter 18 **Configuring the SIP ALG . . . 241**

Understanding the SIP ALG . . . 241
SIP ALG Operation . . . 242
SDP Session Descriptions . . . 244
Pinhole Creation . . . 244
Understanding IPv6 Support for SIP ALG . . . 246
Understanding SIP ALG Request Methods . . . 247
SIP ALG Configuration Overview . . . 248
Understanding SIP ALG Call Duration and Timeouts . . . 249
Example: Setting SIP ALG Call Duration and Timeouts . . . 250
Understanding SIP ALG DoS Attack Protection . . . 251
Example: Configuring SIP ALG DoS Attack Protection . . . 252
Understanding SIP ALG Unknown Message Types . . . 253
Example: Allowing Unknown SIP ALG Message Types . . . 253
Understanding SIP ALG Hold Resources . . . 254
Retaining SIP ALG Hold Resources (CLI Procedure) . . . 255
Retaining SIP ALG Hold Resources (J-Web Procedure) . . . 255
Understanding the SIP ALG and NAT . . . 256
Outgoing Calls . . . 257
Incoming Calls . . . 257
Forwarded Calls . . . 258
Call Termination . . . 258
Call Re-INVITE Messages . . . 258
Call Session Timers . . . 258
Call Cancellation . . . 258
Forking . . . 259
SIP Messages . . . 259
SIP Headers . . . 259
SIP Body . . . 261
SIP NAT Scenario . . . 261
Classes of SIP Responses . . . 263
NAT Mode in Pure IPv6 Mode (NAT66) for SIP IPv6 ALG . . . 264
NAT-PT . . . 264
NAT64 . . . 265
STUN and SIP ALG . . . 265
Understanding Incoming SIP ALG Call Support Using the SIP Registrar and NAT . . . 266
Example: Configuring Interface Source NAT for Incoming SIP Calls . . . 267
Decreasing Network Complexity by Configuring a Source NAT Pool for Incoming SIP Calls . . . 273
Example: Configuring Static NAT for Incoming SIP Calls . . . 281

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the SIP Proxy in the Private Zone and NAT in the Public Zone . . . 288
Example: Configuring a Three-Zone SIP ALG and NAT Scenario . . . 294
Verifying SIP ALG Configurations . . . 302
Verifying SIP ALG . . . 302
Verifying SIP ALG Calls . . . 302
Verifying SIP ALG Call Details . . . 302
Verifying SIP ALG Counters . . . 303
Verifying the Rate of SIP ALG Messages . . . 304

Part 4 Configuration Statements and Operational Commands

Chapter 19 Configuration Statements . . . 309

Applications Configuration Statement Hierarchy . . . 311
[edit security alg] Hierarchy Level . . . 312
[edit security nat] Hierarchy Level . . . 316
[edit security policies] Hierarchy Level . . . 320
[edit security zones] Hierarchy Level . . . 324
[edit security traceoptions] Hierarchy Level . . . 325
address (Security Destination NAT) . . . 326
alg . . . 327
alg (Applications) . . . 332
alg-manager . . . 333
allow-dns-reply . . . 333
application (Security Policies) . . . 334
application-protocol (Applications) . . . 335
application-screen (Security H323) . . . 336
application-screen (Security MGCP) . . . 337
application-screen (Security SCCP) . . . 338
application-screen (Security SIP) . . . 339
banner (Access FTP HTTP Telnet Authentication) . . . 340
call-flood . . . 340
c-timeout . . . 341
deny (Security SIP) . . . 342
destination-address (Security Destination NAT) . . . 343
destination-address (Security Policies) . . . 344
destination-address (Security Source NAT) . . . 345
destination-address (Security Static NAT) . . . 345
destination-nat . . . 346
destination-port (Applications) . . . 347
dns . . . 351
dns (System Services) . . . 352
dscp-rewrite . . . 353
endpoint-registration-timeout . . . 354
family inet (Interfaces) . . . 355
ftp (Access) . . . 357
ftp (Security ALG) . . . 358
gatekeeper . . . 359
h323 . . . 360

Copyright © 2016, Juniper Networks, Inc.

host-inbound-traffic . . . 361
ike-esp-nat . . . 362
ike (Security) . . . 363
inactive-media-timeout (Security MGCP) . . . 365
inactive-media-timeout (Security SCCP) . . . 365
inactive-media-timeout (Security SIP) . . . 366
map-entry-timeout . . . 366
maximum-call-duration (Security) . . . 367
maximum-message-length . . . 367
media-source-port-any . . . 368
message-flood (Security H323) . . . 368
mgcp . . . 369
msrpc . . . 370
nat . . . 371
nat-pat-address . . . 375
policy (Security Policies) . . . 376
pptp . . . 378
protect . . . 379
protocols (Security Zones Interfaces) . . . 380
retain-hold-resource . . . 381
rsh . . . 382
rtsp . . . 383
sccp . . . 384
security-zone . . . 385
sip (Security) . . . 387
source-address (Security Destination NAT) . . . 388
source-address (Security Policies) . . . 389
source-nat . . . 390
sql . . . 391
static-nat . . . 392
sunrpc . . . 393
support-lib . . . 394
system-services (Security Zones Host Inbound Traffic) . . . 395
t1-interval . . . 397
t4-interval . . . 397
talk . . . 398
term (Applications) . . . 399
tftp (Security ALG) . . . 400
traceoptions (Security ALG) . . . 401
traceoptions (Security H323 ALG) . . . 403
traceoptions (Security MGCP ALG) . . . 404
traceoptions (Security SCCP ALG) . . . 405
traceoptions (Security SIP ALG) . . . 406
traceoptions (System Services DNS) . . . 407
transaction-timeout . . . 409
unknown-message (Security H323 ALG) . . . 410
unknown-message (Security MGCP ALG) . . . 411
unknown-message (Security SCCP ALG) . . . 412
unknown-message (Security SIP ALG) . . . 413

Copyright © 2016, Juniper Networks, Inc.

Chapter 20 Operational Commands 415

clear security alg h323 counters 416
clear security alg ike-esp-nat 417
clear security alg sccp calls 418
clear security alg sccp counters 419
clear security alg sip calls 420
clear security alg sip counters 421
clear security flow session application 422
show chassis cluster data-plane statistics 424
show chassis cluster statistics 426
show security alg h323 counters 430
show security alg ike-esp-nat summary 432
show security alg msrpc 433
show security alg sccp calls 435
show security alg sccp counters 437
show security alg sip calls 439
show security alg sip counters 441
show security alg sip rate 445
show security alg status 447
show security flow gate 449
show security flow session 453
show security flow session application 459
show security flow session resource-manager 463
show security idp policy-templates 467
show security resource-manager group active 468
show security resource-manager resource active 471
show security resource-manager summary 474
show security zones 475
show security zones type 478

Guide 2 AppSecure Services Feature Guide for Security Devices

Chapter 21 Overview 483

Understanding AppSecure Services 483

Chapter 22 Understanding Application Identification 485

Understanding Application Identification Techniques 485
Junos OS Next-Generation Application Identification 485
Application Signature Mapping 486
Application Identification Match Sequence 486
Understanding the Junos OS Application Identification Database 488

Copyright © 2016, Juniper Networks, Inc.

Chapter 23 Installing Application Signature Package 489
Understanding the Junos OS Application Package Installation 489
Upgrading to Next-Generation Application Identification 490
Installing and Verifying Licenses for an Application Signature Package 491
Downloading and Installing the Junos OS Application Signature Package Manually 493
Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package 496
Example: Scheduling the Application Signature Package Updates 499
Scheduling the Application Signature Package Updates As Part of the IDP Security Package 501
Example: Downloading and Installing the Application Identification Package in Chassis Cluster Mode 503
Verifying the Junos OS Application Identification Extracted Application Package 506
Uninstalling the Junos OS Application Identification Application Package 507
Disabling and Reenabling Junos OS Application Identification 507

Chapter 24 Custom Application Signatures 509
Understanding Junos OS Application Identification Custom Application Signatures 509
ICMP-Based Mapping 510
Address-Based Mapping 510
IP Protocol-Based Mapping 511
Layer 7-Based Signatures 511
Example: Configuring Junos OS Application Identification Custom Application Signatures 512

Chapter 25 Configuring Application Groups 517
Customizing Application Groups for Junos OS Application Identification 517
Enabling Application Groups in Junos OS Application Identification 518
Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management 518

Chapter 26 Configuring Application System Cache 523
Understanding the Application System Cache 523
Deactivating Application System Cache Information for Application Identification (CLI Procedure) 524
Verifying Application System Cache Statistics 524

Chapter 27 Controlling Application Identification Performance 527
Onbox Application Identification Statistics 527
Understanding Jumbo Frames Support for Junos OS Application Identification Services 528
Improving the Application Traffic Throughput 528

Chapter 28 Configuring Encrypted Files Using SSL Proxy 531
SSL Proxy Overview 531
Supported Ciphers in Proxy Mode 533
Server Authentication 534
Trusted CA List 535

Copyright © 2016, Juniper Networks, Inc.

Root CA . . . 536
Client Authentication . . . 536
Whitelists . . . 536
Dynamic Resolution of Domain Names . . . 536
Session Resumption . . . 536
Session Renegotiation . . . 536
SSL Proxy Logs . . . 537
Leveraging Dynamic Application Identification . . . 538
Logical Systems Support . . . 539
Limitations . . . 539
Application Firewall, IDP, and Application Tracking with SSL Proxy Overview . . . 540
Configuring SSL Proxy . . . 540
SSL Proxy Configuration Overview . . . 541
Configuring a Root CA Certificate . . . 542
Configuring a CA Profile Group . . . 543
Configuring a Trusted CA Profile . . . 544
Importing a Root CA Certificate into a Browser . . . 545
Applying an SSL Proxy Profile to a Security Policy . . . 546
Creating a Whitelist of Exempted Destinations . . . 547
Configuring SSL Proxy Logging . . . 549
Configuring Ciphers . . . 549
Exporting Certificates to a Specified Location . . . 550
Ignoring Server Authentication . . . 550
Working with the Certificate Revocation Lists for SSL Proxy . . . 550
Disabling CRL Verification . . . 551
Allowing Sessions When CRL Information Is Not Available . . . 551
Allowing Sessions When CRL Status Is Unknown . . . 552
Enabling Debugging and Tracing for SSL Proxy . . . 552

Chapter 29 **Configuring Application Firewall . . . 555**

Application Firewall Overview . . . 555
Understanding Application Firewall Rule Sets . . . 556
Configuring an Application Firewall Within a Security Policy . . . 557
Application Group Support for Application Firewall . . . 557
Redirecting Users . . . 558
Session Logging for Application Firewalls . . . 559
Application Firewall Support in Chassis Cluster . . . 559
Example: Configuring Application Firewall Rule Sets Within a Security Policy . . . 560
Example: Configuring an Application Group for Application Firewall . . . 564
Example: Configuring Application Firewall When SSL Proxy Is Enabled . . . 568

Chapter 30 **Configuring Application Tracking . . . 573**

Understanding AppTrack . . . 573
Example: Configuring AppTrack . . . 574
Example: Configuring AppTrack When SSL Proxy Is Enabled . . . 580
Disabling AppTrack . . . 581

Copyright © 2016, Juniper Networks, Inc.

Chapter 31 Configuring Application QoS 583

Understanding Application QoS (AppQoS) 583
Unique Forwarding Classes and Queue Assignments 584
Application-Aware DSCP Code-Point and Loss Priority Settings 584
Rate Limiters and Profiles 586
Rate-Limiter Assignment 587
Rate-Limiter Action 589
AppQoS Security Policy Configuration 589
Example: Configuring AppQoS 589

Chapter 32 Configuration Statements 597

Class-of-Service Configuration Statement Hierarchy 599
Security Configuration Statement Hierarchy 603
Services Configuration Statement Hierarchy 605
System Configuration Statement Hierarchy 612
[edit security address-book] Hierarchy Level 643
[edit security application-firewall] Hierarchy Level 644
[edit security application-tracking] Hierarchy Level 644
[edit security log] Hierarchy Level 645
[edit security pki] Hierarchy Level 646
[edit security policies] Hierarchy Level 647
[edit security zones] Hierarchy Level 651
actions (Services SSL Proxy) 654
actions (Services SSL Initiation) 655
address-mapping (Application Identification) 656
appfw-profile (System) 657
appfw-rule 658
appfw-rule-set 659
application-firewall 660
application (Application Identification) 662
application-firewall (Application Services) 664
application-identification 665
application-group (Services) 667
application-services (Security Policies) 668
application-system-cache 669
application-system-cache-timeout (Services) 669
application-tracking 670
application-tracking (Security Zones) 670
application-traffic-control 671
application-traffic-control (Application Services) 672
block-message (Application Firewall) 673
context (Application Identification) 675
custom-ciphers 676
default-rule 677
direction (Application Identification) 678
disable (Application Tracking) 678
download (Services) 679
dynamic-application 680
dynamic-application-group 680

Copyright © 2016, Juniper Networks, Inc.

enable-flow-tracing (Services) . . . 681
enable-performance-mode . . . 681
enable-session-cache . . . 682
file (Services) . . . 682
files (Services) . . . 683
file (System Logging) . . . 684
first-update . . . 686
first-update-interval . . . 687
flag (Services) . . . 688
format (Security Log) . . . 688
forwarding-classes (CoS) . . . 689
global-config (Services) . . . 690
icmp-mapping (Application Identification) . . . 691
ip (Application Identification) . . . 691
ip-protocol-mapping (Application Identification) . . . 692
initiation (Services) . . . 693
level (Services) . . . 694
log (Security) . . . 695
log (Services) . . . 697
match (Services) . . . 698
no-application-identification (Services) . . . 698
no-application-system-cache (Services) . . . 699
no-remote-trace (Services) . . . 699
over (Application Identification) . . . 700
policies . . . 702
policy (Security Policies) . . . 707
port-range (Application Identification) . . . 709
preferred-ciphers . . . 710
profile (Application Firewall) . . . 710
profile (Rule Sets) . . . 711
profile (Services) . . . 712
profile (SSL Initiation) . . . 713
profile (SSL Termination) . . . 714
protocol-version . . . 715
proxy (Services) . . . 716
rate-limiters . . . 717
renegotiation (Services) . . . 718
root-ca (Services) . . . 718
rule-sets (CoS AppQoS) . . . 719
rule-sets (Security Application Firewall) . . . 721
security-zone . . . 722
server-certificate (Services) . . . 723
session-update-interval . . . 724
size (Services) . . . 724
ssl (Services) . . . 725
ssl-encryption . . . 727
ssl-proxy (Application Services) . . . 727
statistics (Services) . . . 728
stream (Security Log) . . . 729

Copyright © 2016, Juniper Networks, Inc.

termination (Services) . . . 730
then (Security Application Firewall) . . . 731
trusted-ca (Services) . . . 732
traceoptions (Security Application Firewall) . . . 733
traceoptions (Services SSL) . . . 735
traceoptions (Services Application Identification) . . . 736
transport (Security Log) . . . 738
whitelist (Services) . . . 739
zones . . . 740

Chapter 33 Operational Commands . . . 743

clear security application-firewall rule-set statistics . . . 745
clear security application-firewall rule-set statistics logical-system . . . 746
clear services application-identification application-statistics . . . 747
clear services application-identification application-statistics cumulative . . . 748
clear services application-identification application-statistics interval . . . 749
clear services application-identification application-system-cache (Junos OS) . . . 750
clear services application-identification counter (Values) . . . 751
clear services ssl proxy statistics . . . 752
request security pki ca-certificate ca-profile-group load . . . 753
request security pki local-certificate export . . . 755
request security pki local-certificate generate-self-signed . . . 756
request security pki local-certificate load . . . 757
request services application-identification application . . . 758
request services application-identification download . . . 760
request services application-identification download status . . . 761
request services application-identification group . . . 762
request services application-identification install . . . 764
request services application-identification install status . . . 765
request services application-identification proto-bundle-status . . . 766
request services application-identification uninstall . . . 767
request services application-identification uninstall status . . . 768
show class-of-service application-traffic-control counter . . . 769
show class-of-service application-traffic-control statistics rate-limiter . . . 771
show class-of-service application-traffic-control statistics rule . . . 773
show security application-firewall rule-set . . . 774
show security application-firewall rule-set logical-system . . . 777
show security application-tracking counters . . . 780
show security flow session . . . 781
show security flow session application-firewall . . . 787
show security pki ca-certificate . . . 793
show security pki local-certificate (View) . . . 797
show security policies . . . 802
show services application-identification application . . . 810
show services application-identification application-system-cache (View) . . . 813
show services application-identification commit-status . . . 815
show services application-identification counter (AppSecure) . . . 816
show services application-identification group . . . 819

Copyright © 2016, Juniper Networks, Inc.

show services application-identification statistics applications 821
show services application-identification statistics application-groups 823
show services application-identification status 825
show services application-identification version 828
show services ssl proxy statistics 829

Guide 3 Attack Detection and Prevention Feature Guide for Security Devices

Part 5 Overview

Chapter 34 Introduction to Attack Detection and Prevention 835

Attack Detection and Prevention Overview 835
Understanding IPv6 Support for Screens 836
IPv6 Extension Header Checking and Filtering 836
Maximum Number of Extension Headers 837
Bad Option Extension Headers 838
ICMPv6 Checking and Filtering 838
IPv6 Packet Header Checking and Filtering 839
Understanding Screens Options on SRX Series Devices 840
Understanding Screen Options on the SRX5000 Module Port Concentrator . . 843
Statistics-Based Screens 844
Differences Between IOC1 and IOC2 844
Signature-Based Screens 848
Example: Configuring Multiple Screening Options 849
Understanding Central Point Architecture Enhancements for Screens 854

Part 6 Configuring Screen Options to Protect Against Denial-of-Service Attacks

Chapter 35 Understanding DoS Attacks 859

DoS Attack Overview 859

Chapter 36 Protecting Against Firewall DoS Attacks 861

Firewall DoS Attacks Overview 861
Understanding Firewall Filters on the SRX5000 Module Port Concentrator 861
Understanding Session Table Flood Attacks 862
Understanding Source-Based Session Limits 863
Example: Setting Source-Based Session Limits 864
Understanding Destination-Based Session Limits 866
Example: Setting Destination-Based Session Limits 867
Understanding SYN-ACK-ACK Proxy Flood Attacks 869
Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack 870

Copyright © 2016, Juniper Networks, Inc.

Chapter 37 Protecting Against Network DoS Attacks 873

Network DoS Attacks Overview 873
Understanding SYN Flood Attacks 874
SYN Flood Protection 875
SYN Flood Options 876
Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection 877
Example: Enabling SYN Flood Protection for Webservers in the DMZ 879
Understanding Whitelists for SYN Flood Screens 885
Example: Configuring Whitelists for SYN Flood Screens 886
Understanding SYN Cookie Protection 887
SYN Cookie Options 889
Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection 890
Understanding ICMP Flood Attacks 893
Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection 894
Understanding UDP Flood Attacks 896
Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection 897
Understanding Land Attacks 899
Protecting Your Network Against Land Attacks by Enabling Land Attack 900
Understanding Screen IPv6 Tunneling Control 902
Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options 905

Chapter 38 Protecting Against OS-Specific DoS Attacks 913

OS-Specific DoS Attacks Overview 913
Understanding Ping of Death Attacks 913
Example: Protecting Against a Ping of Death Attack 914
Understanding Teardrop Attacks 915
Example: Protecting Against a Teardrop Attack 917
Understanding WinNuke Attacks 917
Example: Protecting Against a WinNuke Attack 919

Part 7 Configuring Reconnaissance Deterrence for Security Devices

Chapter 39 Protecting Against IP Sweep and Port Options 923

Reconnaissance Deterrence Overview 923
Understanding IP Address Sweeps 923
Example: Blocking IP Address Sweeps 924
Understanding TCP Port Scanning 926
Enhancing Traffic Management by Blocking Port Scans 927
Understanding UDP Port Scanning 930
Understanding Network Reconnaissance Using IP Options 930
Uses for IP Packet Header Options 931
Screen Options for Detecting IP Options Used for Reconnaissance 933
Example: Detecting Packets That Use IP Screen Options for Reconnaissance . . 933

Copyright © 2016, Juniper Networks, Inc.

Chapter 40 Protecting Against System Probes and Flag Set 937

Understanding Operating System Probes 937
Understanding Domain Name System Resolve 937
Understanding TCP Headers with SYN and FIN Flags Set 938
Example: Blocking Packets with SYN and FIN Flags Set 939
Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set 940
Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set 941
Understanding TCP Header with No Flags Set 943
Example: Blocking Packets with No Flags Set 944

Chapter 41 Protecting Against Attacker Evasion Techniques 947

Understanding Attacker Evasion Techniques 947
Understanding FIN Scans 948
Thwarting a FIN Scan 948
Understanding TCP SYN Checking 948
Setting TCP SYN Checking 950
Setting Strict SYN Checking 951
Understanding IP Spoofing 951
Example: Blocking IP Spoofing 951
Understanding IP Spoofing in Layer 2 Transparent Mode 953
Configuring IP Spoofing in Layer 2 Transparent Mode 954
Understanding IP Source Route Options 955
Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set 957
Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set 958

Part 8 Configuring Suspicious Packet Attributes for Security Devices

Chapter 42 Protecting Against ICMP and SYN Fragment Attacks 963

Suspicious Packet Attributes Overview 963
Understanding ICMP Fragment Protection 963
Example: Blocking Fragmented ICMP Packets 964
Understanding Large ICMP Packet Protection 965
Example: Blocking Large ICMP Packets 966
Understanding SYN Fragment Protection 966
Example: Dropping IP Packets Containing SYN Fragments 967

Chapter 43 Protecting Against IP Attacks 969

Understanding Bad IP Option Protection 969
Example: Blocking IP Packets with Incorrectly Formatted Options 970
Understanding Unknown Protocol Protection 971
Example: Dropping Packets Using an Unknown Protocol 972
Understanding IP Packet Fragment Protection 972
Example: Dropping Fragmented IP Packets 973

Copyright © 2016, Juniper Networks, Inc.

Part 9 Configuration Statements and Operational Commands

Chapter 44 Configuration Statements 977

Security Configuration Statement Hierarchy 978
[edit security screen] Hierarchy Level 979
attack-threshold 982
bad-inner-header 983
description (Security Screen) 983
destination-ip-based 984
destination-threshold 985
fin-no-ack 985
flood (Security ICMP) 986
flood (Security UDP) 987
gre 988
icmp (Security Screen) 989
ids-option 990
ipip 993
ip (Security Screen) 994
ip-sweep 997
ip-in-udp 998
land 998
large 999
limit-session 999
no-syn-check 1000
no-syn-check-in-tunnel 1000
ping-death 1001
port-scan 1002
screen (Security Zones) 1003
source-ip-based 1003
source-threshold 1004
strict-syn-check 1004
syn-ack-ack-proxy 1005
syn-check-required 1005
syn-fin 1006
syn-flood 1007
syn-flood-protection-mode 1008
syn-frag 1008
tcp (Security Screen) 1009
tcp-no-flag 1010
tcp-sweep 1011
timeout (Security Screen) 1012
traceoptions (Security Screen) 1013
trap 1014
tunnel (Security Screen) 1015
udp (Security Screen) 1016
udp-sweep 1017
white-list 1018
winnuke 1019

Copyright © 2016, Juniper Networks, Inc.

Chapter 45 Operational Commands 1021
clear security screen statistics 1022
clear security screen statistics interface 1023
clear security screen statistics zone 1025
show security screen ids-option 1027
show security screen statistics 1032
show security screen status 1040

Guide 4 Building Blocks Feature Guide for Security Devices

Part 10 Overview

Chapter 46 Introduction to Security Building Blocks 1045
Understanding Security Building Blocks for Security Devices 1045

Part 11 Configuring Security Zones and Interfaces

Chapter 47 Configuring Security Zones 1049
Security Zones and Interfaces Overview 1049
Understanding Security Zone Interfaces 1050
Understanding Functional Zones 1050
Understanding Security Zones 1050
Example: Creating Security Zones 1051

Chapter 48 Managing Inbound Traffic for Security Zones 1055
Understanding How to Control Inbound Traffic Based on Traffic Types 1055
Example: Controlling Inbound Traffic Based on Traffic Types 1056
Understanding How to Control Inbound Traffic Based on Protocols 1058
Example: Controlling Inbound Traffic Based on Protocols 1059
Supported System Services for Host Inbound Traffic 1061

Chapter 49 Identifying Duplicate Sessions by Configuring TCP-Reset Parameters . . 1065
Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter 1065
Example: Configuring the TCP-Reset Parameter 1065

Part 12 Configuring Address Books and Address Sets

Chapter 50 Configuring Address, Address Books, and Address Sets 1069
Understanding Address Books 1069
Predefined Addresses 1069
Network Prefixes in Address Books 1070
Wildcard Addresses in Address Books 1070
DNS Names in Address Books 1070
Understanding Address Sets 1071
Understanding Global Address Books 1071
Configuring Addresses and Address Sets 1072
Addresses and Address Sets 1072
Address Books and Security Zones 1073

Copyright © 2016, Juniper Networks, Inc.

Address Books and Security Policies 1073
Addresses Available for Security Policies 1074
Applying Policies to Address Sets 1074
Address Books and NAT 1075
Example: Configuring Address Books and Address Sets 1076
Limitations of Addresses and Address Sets in a Security Policy 1081

Part 13 Configuring Security Policies

Chapter 51 Enforcing Transit Traffic Rules by Configuring Security Policies 1085

Security Policies Overview 1085
Understanding Security Policy Rules 1087
Understanding Wildcard Addresses 1090
Understanding Security Policy Elements 1091
Understanding Security Policies for Self Traffic 1092
Security Policies Configuration Overview 1093
Configuring Policies Using the Firewall Wizard 1094
Example: Configuring a Security Policy to Permit or Deny All Traffic 1094
Example: Configuring a Security Policy to Permit or Deny Selected Traffic 1098
Example: Configuring a Security Policy to Permit or Deny Wildcard Address Traffic 1102
Example: Configuring a Security Policy to Redirect Traffic Logs to an External System Log Server 1105

Chapter 52 Configuring Negated Addresses 1109

Understanding Negated Address Support 1109
Example: Configuring Negated Addresses 1110

Chapter 53 Configuring Global Security Policy 1115

Global Policy Overview 1115
Example: Configuring a Global Policy with No Zone Restrictions 1117
Example: Configuring a Global Policy with Multiple Zones 1119

Chapter 54 Managing Security Policy Activation By Configuring Schedulers 1123

Security Policy Schedulers Overview 1123
Example: Configuring Schedulers for a Daily Schedule Excluding One Day 1124

Chapter 55 Configuring User Role Firewall Security Policies 1127

Understanding User Role Firewalls 1127
User Role Retrieval and the Policy Lookup Process 1128
Understanding the User Identification Table 1130
Local Authentication Table 1131
UAC Authentication Table 1133
Firewall Authentication Table 1134
Policy Provisioning With Users and Roles 1135
Obtaining Username and Role Information Through Firewall Authentication . . 1136
Configuring a User Role Firewall For Captive Portal Redirection 1137
Example: Configuring a User Role Firewall on an SRX Series Device 1138
Configuring Resource Policies Using UAC 1145

Copyright © 2016, Juniper Networks, Inc.

Chapter 56 Setting Security Policy Reorder 1149
Understanding Security Policy Ordering 1149
Example: Reordering the Policies 1151
Chapter 57 Monitoring and Troubleshooting Security Policies 1153
Matching Security Policies 1153
Tracking Policy Hit Counts 1155
Best Practices for Defining Policies on High-End SRX Series Devices 1155
Checking Memory Status 1157
Synchronizing a Security Policy on SRX Series Devices 1159
Verifying Scheduled Policies 1159
Verifying Shadow Policies 1160
Verifying All Shadow Policies 1160
Verifying a Policy Shadows One or More Policies 1161
Verifying a Policy Is Shadowed by One or More Policies 1161
Monitoring Policy Statistics 1162
Troubleshooting Security Policies 1162
Checking a Security Policy Commit Failure 1163
Verifying a Security Policy Commit 1163
Debugging Policy Lookup 1163
Chapter 58 Handling Security Policy Violations 1165
Understanding Searching and Sorting Audit Log 1165
Understanding Packet Flow Alarms and Auditing 1166
Example: Generating a Security Alarm in Response to Policy Violations 1167
Part 14 Configuring Security Policy Applications
Chapter 59 Configuring Applications and Application Sets 1171
Security Policy Applications Overview 1171
Policy Application Sets Overview 1172
Example: Configuring Applications and Application Sets 1172
Chapter 60 Configuring Custom Policy Applications 1175
Understanding Custom Policy Applications 1175
Custom Application Mappings 1175
Example: Adding and Modifying Custom Policy Applications 1176
Example: Configuring Custom Policy Application Term Options 1177
Chapter 61 Setting Policy Application Timeout 1181
Understanding Policy Application Timeout Configuration and Lookup 1181
Understanding Policy Application Timeouts Contingencies 1182
Example: Setting a Policy Application Timeout 1182
Chapter 62 Understanding Predefined Policy Applications 1185
Understanding Internet-Related Predefined Policy Applications 1185
Understanding Microsoft Predefined Policy Applications 1187
Understanding Dynamic Routing Protocols Predefined Policy Applications . . . 1188
Understanding Streaming Video Predefined Policy Applications 1188
Understanding Sun RPC Predefined Policy Applications 1189
Understanding Security and Tunnel Predefined Policy Applications 1190

Copyright © 2016, Juniper Networks, Inc.

Understanding IP-Related Predefined Policy Applications 1191
Understanding Instant Messaging Predefined Policy Applications 1191
Understanding Management Predefined Policy Applications 1192
Understanding Mail Predefined Policy Applications 1193
Understanding UNIX Predefined Policy Applications 1194
Understanding Miscellaneous Predefined Policy Applications 1194
Understanding the ICMP Predefined Policy Application 1195
Example: Defining a Custom ICMP Application 1199
Default Behavior of ICMP Unreachable Errors 1201

Part 15 Configuration Statements and Operational Commands

Chapter 63 Configuration Statements 1205

Applications Configuration Statement Hierarchy 1208
Security Configuration Statement Hierarchy 1209
[edit security address-book] Hierarchy Level 1210
[edit security policies] Hierarchy Level 1211
[edit security user-identification] Hierarchy Level 1215
[edit security zones] Hierarchy Level 1216
address (Security Address Book) 1218
address-book 1219
address-set 1220
alarms (Security) 1221
alarm-threshold 1222
alarm-without-drop 1223
application (Applications) 1224
application (Security Alarms) 1227
application (Security Policies) 1228
application-protocol (Applications) 1229
application-services (Security Policies) 1230
application-tracking (Security Zones) 1231
application-traffic-control (Application Services) 1231
attach 1232
audible (Security Alarms) 1232
authentication (Security Alarms) 1233
authentication-source 1234
captive-portal (Services UAC Policy) 1234
count (Security Policies) 1235
default-policy 1235
deny (Security Policies) 1236
description (Applications) 1236
description (Security Address Book) 1237
description (Security Policies) 1238
description (Security Zone) 1239
destination-address (Security Policies) 1240
destination-address (Security Policies Flag) 1241
destination-address-excluded 1242
destination-ip (Security Alarms) 1243
destination-port (Applications) 1244

Copyright © 2016, Juniper Networks, Inc.

dns-proxy . . . 1248
dynamic-dns . . . 1249
exclude (Schedulers) . . . 1250
firewall-authentication (Security Policies) . . . 1251
firewall-authentication (User Identification) . . . 1252
forward-only (DNS) . . . 1252
from-zone (Security Policies) . . . 1253
from-zone (Security Policies Global) . . . 1255
functional-zone . . . 1256
global (Security Policies) . . . 1257
host-inbound-traffic . . . 1259
icmp-code (Applications) . . . 1260
icmp-type (Applications) . . . 1260
inactivity-timeout (Applications) . . . 1261
interfaces (Security Zones) . . . 1262
initial-tcp-mss . . . 1263
ipsec-group-vpn (Security Policies) . . . 1264
ipsec-vpn (Security Policies) . . . 1264
local-authentication-table . . . 1265
log (Security Policies) . . . 1265
management (Security Zones) . . . 1266
match (Security Policies) . . . 1267
match (Security Policies Global) . . . 1268
no-policy-cold-synchronization . . . 1269
pair-policy . . . 1270
pass-through . . . 1271
permit (Security Policies) . . . 1272
policies . . . 1274
policy (Security Alarms) . . . 1279
policy (Security Policies) . . . 1280
policy-match . . . 1282
policy-rematch . . . 1283
policy-stats . . . 1284
potential-violation . . . 1285
protocol (Applications) . . . 1287
protocols (Security Zones Host Inbound Traffic) . . . 1288
protocols (Security Zones Interfaces) . . . 1290
range-address . . . 1291
redirect-wx (Application Services) . . . 1291
reject (Security) . . . 1292
reverse-tcp-mss . . . 1293
rpc-program-number (Applications) . . . 1294
scheduler (Security Policies) . . . 1295
scheduler-name . . . 1296
schedulers (Security Policies) . . . 1296
screen (Security Zones) . . . 1297
secure-domains . . . 1297
secure-neighbor-discovery . . . 1298
security-zone . . . 1299

Copyright © 2016, Juniper Networks, Inc.

sequence-check-required . . . 1300
services-offload (Security) . . . 1300
session-close . . . 1301
session-init . . . 1301
simple-mail-client-service . . . 1302
source-address (Security Policies) . . . 1303
source-address-excluded . . . 1303
source-identity . . . 1304
source-ip (Security Alarms) . . . 1305
source-port (Applications) . . . 1306
ssl-proxy (Application Services) . . . 1306
ssl-termination-profile . . . 1307
start-date . . . 1307
start-time (Schedulers) . . . 1308
stop-date . . . 1309
stop-time . . . 1310
syn-check-required . . . 1310
system-services (Security Zones Host Inbound Traffic) . . . 1311
system-services (Security Zones Interfaces) . . . 1313
tcp-options (Security Policies) . . . 1315
tcp-rst . . . 1316
term (Applications) . . . 1316
then (Security Policies) . . . 1317
to-zone (Security Policies) . . . 1319
to-zone (Security Policies Global) . . . 1321
traceoptions (Security Policies) . . . 1322
traceoptions (Security User Identification) . . . 1324
traceoptions (System Services DNS) . . . 1326
tunnel (Security Policies) . . . 1328
uac-policy (Application Services) . . . 1328
unified-access-control (Security) . . . 1329
user-firewall . . . 1330
user-identification . . . 1331
utm-policy . . . 1332
uuid (Applications) . . . 1333
vrrp . . . 1334
web-authentication . . . 1335
web-redirect . . . 1335
zones . . . 1336

Chapter 64 Operational Commands . . . 1339

clear security alarms . . . 1340
clear security policies hit-count . . . 1343
clear security policies statistics . . . 1344
clear system services dns dns-proxy . . . 1345
request security user-identification local-authorization-table add . . . 1346
request security user-identification local-authentication-table delete . . . 1348
show security alarms . . . 1349
show security firewall-authentication users address . . . 1353

Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users auth-type 1356
show security flow session application 1358
show security match-policies 1362
show security policies 1367
show security policies hit-count 1375
show security policies unknown-source-identity 1378
show security shadow-policies logical-system 1380
show security user-identification local-authentication-table 1381
show security user-identification role-provision all 1383
show security user-identification source-identity-provision all 1384
show security user-identification user-provision all 1385
show security zones 1386
show security zones type 1389
show system services dns dns-proxy 1392
show system services dynamic-dns 1395

Guide 5 Class of Service Feature Guide for Security Devices

Part 16 Overview

Chapter 65 Introduction to Class of Service 1401

Understanding Class of Service 1401
Benefits of CoS 1402
CoS Across the Network 1403
Junos OS CoS Components 1403
CoS Components Packet Flow 1405
CoS Process on Incoming Packets 1406
CoS Process on Outgoing Packets 1406
CoS Device Configuration Overview 1407
Understanding CoS Default Settings 1407

Part 17 Configuring Class of Service Components

Chapter 66 Assigning Service Levels with Classifiers 1411

Classification Overview 1411
Behavior Aggregate Classifiers 1412
Multifield Classifiers 1412
Default IP Precedence Classifier 1413
Understanding Packet Loss Priorities 1414
Default Behavior Aggregate Classification 1414
Sample Behavior Aggregate Classification 1416
Example: Configuring Behavior Aggregate Classifiers 1417

Chapter 67 Controlling Network Access with Traffic Policing 1427

Simple Filters and Policers Overview 1427
Two-Rate Three-Color Policer Overview 1428
Example: Configuring a Two-Rate Three-Color Policer 1429
Logical Interface (Aggregate) Policer Overview 1434
Two-Color Policer Configuration Overview 1434
Example: Configuring a Two-Color Logical Interface (Aggregate) Policer 1436

Copyright © 2016, Juniper Networks, Inc.

Guidelines for Configuring Simple Filters . . . 1442
Statement Hierarchy for Configuring Simple Filters . . . 1442
Simple Filter Protocol Families . . . 1443
Simple Filter Names . . . 1443
Simple Filter Terms . . . 1443
Simple Filter Match Conditions . . . 1443
Simple Filter Terminating Actions . . . 1445
Simple Filter Nonterminating Actions . . . 1445
Example: Configuring and Applying a Firewall Filter for a Multifield Classifier . . 1445

Chapter 68 Controlling Output Queues with Forwarding Classes . . . 1453

Forwarding Classes Overview . . . 1453
Forwarding Class Queue Assignments . . . 1454
Forwarding Policy Options . . . 1455
Example: Configuring Forwarding Classes . . . 1455
Example: Assigning Forwarding Classes to Output Queues . . . 1460
Example: Assigning a Forwarding Class to an Interface . . . 1462
Understanding the SPC High-Priority Queue . . . 1463
Example: Configuring the SPC High-Priority Queue . . . 1463

Chapter 69 Altering Outgoing Packets Headers with Rewrite Rules . . . 1467

Rewrite Rules Overview . . . 1467
Rewriting Frame Relay Headers . . . 1467
Assigning the Default Frame Relay Rewrite Rule to an Interface . . . 1468
Defining a Custom Frame Relay Rewrite Rule . . . 1468
Example: Configuring and Applying Rewrite Rules . . . 1469

Chapter 70 Defining Output Queue Properties with Schedulers . . . 1473

Schedulers Overview . . . 1473
Transmit Rate . . . 1474
Delay Buffer Size . . . 1475
Scheduling Priority . . . 1476
Shaping Rate . . . 1477
Default Scheduler Settings . . . 1478
Transmission Scheduling Overview . . . 1479
Excess Bandwidth Sharing and Minimum Logical Interface Shaping . . . 1480
Excess Bandwidth Sharing Proportional Rates . . . 1481
Calculated Weights Mapped to Hardware Weights . . . 1482
Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces . . 1483
Shared Bandwidth Among Logical Interfaces . . . 1484
Example: Configuring Class-of-Service Schedulers . . . 1485
Scheduler Buffer Size Overview . . . 1489
Maximum Delay Buffer Sizes Available to Channelized T1/E1 Interfaces . . 1490
Maximum Delay Buffer Size for vSRX Interfaces . . . 1491
Delay Buffer Size Allocation Methods . . . 1491
Delay Buffer Sizes for Queues . . . 1492
Example: Configuring a Large Delay Buffer on a Channelized T1 Interface 1493
Configuring Large Delay Buffers in CoS . . . 1496
Example: Configuring and Applying Scheduler Maps . . . 1500

Copyright © 2016, Juniper Networks, Inc.

Chapter 71 Removing Delays with Strict-Priority Queues 1505
Strict-Priority Queue Overview 1505
Understanding Strict-Priority Queues 1506
Example: Configuring Priority Scheduling 1506
Example: Configuring Strict-Priority Queuing 1508
Chapter 72 Controlling Congestion with Drop Profiles 1519
RED Drop Profiles Overview 1519
Default Drop Profiles 1520
RED Drop Profiles and Congestion Control 1520
Configuring RED Drop Profiles 1521
Example: Configuring RED Drop Profiles 1523
Chapter 73 Controlling Congestion with Adaptive Shapers 1527
Adaptive Shaping Overview 1527
Defining a Custom Frame Relay Loss Priority Map 1528
Example: Configuring and Applying an Adaptive Shaper 1528
Chapter 74 Limiting Traffic Using Virtual Channels 1531
Virtual Channels Overview 1531
Understanding Virtual Channels 1532
Example: Configuring Virtual Channels 1533
Chapter 75 Enabling Queuing for Tunnel Interfaces 1539
CoS Queuing for Tunnels Overview 1539
Benefits of CoS Queuing for Tunnel Interfaces 1539
How CoS Queuing Works 1540
Limitations on CoS Shapers for Tunnel Interfaces 1541
Understanding the ToS Value of a Tunnel Packet 1541
Example: Configuring CoS Queuing for GRE or IP-IP Tunnels 1542
Copying Outer IP Header DSCP and ECN to Inner IP Header 1546
Chapter 76 Naming Components with Code-Point Aliases 1549
Code-Point Aliases Overview 1549
Default CoS Values and Aliases 1550
Example: Defining Code-Point Aliases for Bits 1553
Part 18 Configuring Class of Service Scheduler Hierarchy
Chapter 77 Controlling Traffic by Configuring Scheduler Hierarchy 1557
Understanding Hierarchical Schedulers 1557
Understanding Internal Scheduler Nodes 1560
Example: Configuring a Four-Level Scheduler Hierarchy 1561
Example: Controlling Remaining Traffic 1573
Part 19 Configuring Class of Service for IPv6
Chapter 78 Configuring Class of Service for IPv6 Traffic 1581
CoS Functions for IPv6 Traffic Overview 1581
Understanding CoS with DSCP IPv6 BA Classifier 1583
Example: Configuring CoS with DSCP IPv6 BA Classifiers 1585

Copyright © 2016, Juniper Networks, Inc.

Understanding DSCP IPv6 Rewrite Rules 1588
Example: Configuring CoS with DSCP IPv6 Rewrite Rules 1588

Part 20 Configuring Class of Service for I/O Cards

Chapter 79 Configuring Class of Service for I/O Cards 1595

PIR-Only and CIR Mode Overview 1595
PIR-only Mode 1595
CIR Mode 1596
Understanding Priority Propagation 1597
Understanding IOC Hardware Properties 1598
Understanding IOC Map Queues 1600
WRED on the IOC Overview 1601
Shapers at the Logical Interface Level (Level 3) 1602
Shapers at the Interface Set Level (Level 2) 1603
Shapers at the Port Level (Level 1) 1604
MDRR on the IOC Overview 1605
CoS Support on the SRX5000 Module Port Concentrator Overview 1607
Example: Configuring CoS on High-End SRX Series Devices with an MPC 1608

Part 21 Configuration Statements and Operational Commands

Chapter 80 Configuration Statements 1621

Class-of-Service Configuration Statement Hierarchy 1622
adaptive-shaper 1626
adaptive-shapers 1627
application-traffic-control 1628
buffer-size (Schedulers) 1629
classifiers (CoS) 1630
code-points (CoS) 1631
copy-outer-dscp 1631
default (CoS) 1632
forwarding-classes (CoS) 1633
frame-relay-de (CoS Interfaces) 1634
frame-relay-de (CoS Loss Priority) 1635
frame-relay-de (CoS Rewrite Rule) 1635
ingress-policer-overhead 1636
interfaces (CoS) 1638
logical-interface-policer 1639
loss-priority (CoS Loss Priority) 1639
loss-priority (CoS Rewrite Rules) 1640
loss-priority-maps (CoS Interfaces) 1641
loss-priority-maps (CoS) 1641
rate-limiters 1642
rewrite-rules (CoS) 1643
rewrite-rules (CoS Interfaces) 1644
rule-sets (CoS AppQoS) 1645
scheduler-map (CoS Virtual Channels) 1647
shaping-rate (CoS Adaptive Shapers) 1647

Copyright © 2016, Juniper Networks, Inc.

shaping-rate (CoS Interfaces) . . . 1648
shaping-rate (CoS Virtual Channels) . . . 1649
trigger (CoS) . . . 1649
tunnel-queuing . . . 1650
virtual-channels . . . 1650
virtual-channel-group (CoS Interfaces) . . . 1651
virtual-channel-groups . . . 1652

Chapter 81 Operational Commands . . . 1653

show class-of-service application-traffic-control counter . . . 1654
show class-of-service application-traffic-control statistics rate-limiter . . . 1656
show class-of-service application-traffic-control statistics rule . . . 1658
show class-of-service forwarding-class . . . 1659

Guide 6 Flow-Based and Packet-Based Processing Feature Guide for Security Devices

Part 22 Overview

Chapter 82 Introduction to Processing on Security Devices . . . 1665

Juniper Networks Devices Processing Overview . . . 1665
Understanding Flow-Based Processing . . . 1666
Zones and Policies . . . 1667
Flows and Sessions . . . 1667
Understanding Packet-Based Processing . . . 1667
Configuring Packet-Based Processing . . . 1668
Stateless Firewall Filters . . . 1668
Class-of-Service Features . . . 1668
Screens . . . 1669
Understanding SRX Series Services Gateways Central Point Architecture . . . 1669
Load Distribution in Combo Mode . . . 1670
Sharing Processing Power and Memory in Combo Mode . . . 1670
Understanding Enhancements to Central Point Architecture for the SRX5000 Line . . . 1670
SRX5000 Line Devices Processing Overview . . . 1671
Understanding First-Packet Processing . . . 1672
Understanding Fast-Path Processing . . . 1673
Understanding the Data Path for Unicast Sessions . . . 1674
Session Lookup and Packet-Match Criteria . . . 1675
Understanding Session Creation: First-Packet Processing . . . 1675
Understanding Fast-Path Processing . . . 1678
Understanding Services Processing Units . . . 1681
Understanding Scheduler Characteristics . . . 1681
Understanding Network Processor Bundling . . . 1681
Network Processor Bundling Limitations . . . 1681
Understanding Central Point Session Limit Performance Enhancements . . . 1682
SRX320 Services Gateway Processing Overview . . . 1683
Understanding Flow Processing and Session Management . . . 1683
Understanding First-Packet Processing . . . 1683
Understanding Session Creation . . . 1684

Copyright © 2016, Juniper Networks, Inc.

Understanding Fast-Path Processing 1684
Understanding Central Point Architecture Flow Support for GTP and SCTP . . 1685

Part 23 Configuring Flow-Based Sessions

Chapter 83 Configuring Security Flow Sessions 1689

Understanding Session Characteristics for SRX Series Services Gateways . . . 1689
Understanding Aggressive Session Aging 1690
Example: Controlling Session Termination for SRX Series Services Gateways . . 1690
Understanding TCP Session Checks per Policy 1692
Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways 1693
Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways 1694
Example: Configuring TCP Packet Security Checks Per Policy 1696
Configuring the Timeout Value for Multicast Flow Sessions 1697
Clearing Sessions for SRX Series Services Gateways 1699
Terminating Sessions for SRX Series Services Gateways 1699
Terminating a Specific Session for SRX Series Services Gateways 1699
Using Filters to Specify the Sessions to Be Terminated for SRX Series Services Gateways 1699

Part 24 Improving Flow-Based Performance

Chapter 84 Expanding Session Capacity by Device 1703

Expanding Session Capacity by Device 1703
Reverting to Default Session Capacity on an SRX5800 Device 1704
Verifying the Current Session Capacity 1704

Chapter 85 Managing Sessions and Flow Distribution 1707

Understanding Load Distribution in High-End SRX Series Devices 1707
Hash-Based Forwarding on the SRX5K-MPC, SRX5K-MPC3-40G10G (IOC3), and the SRX5K-MPC3-100G10G (IOC3) 1707
Understanding Load Distribution on SRX5000 Line Devices Using the Packet-Ordering Function 1709
Disabling Packet-Ordering Mode on SRX5000 Line Devices 1710

Chapter 86 Reducing Long Packet-Processing Latency by Express Path 1713

Understanding Session Cache 1713
Overview 1713
Selective Session Cache Installation 1714
IPsec VPN Session Affinity Enhancement Using Session Cache 1716
Fragmentation Packet Ordering Using NP Session Cache 1716
Express Path Overview 1717
Understanding Express Path Functionality 1717
Understanding Express Path Support on SRX Series Devices 1718
Understanding Express Path Features 1719
Wing Statistics Counter 1719
Cross-Network Traffic 1720
LAG Support in Express Path Mode 1721

Copyright © 2016, Juniper Networks, Inc.

End-to-End Debugging . . . 1722
Per-Session Statistics CLI . . . 1722
Express Path Limitations . . . 1723
Express Path Support on NP-IOC Card . . . 1724
Express Path Support on SRX5K Modular Port Concentrator . . . 1725
Express Path Support on SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) . . . 1726
IPv6 Flow in Express Path Mode for IOC2 and IOC3 . . . 1728
Disabling TCP Packet Security Checks . . . 1729
IPv6 Flow in Express Path Mode . . . 1729
Understanding the Express Path Solution . . . 1731
Enabling and Disabling Express Path . . . 1732
Example: Enabling Express Path in Security Policies . . . 1733
Example: Configuring an IOC on SRX5000 Line Devices to Support Express Path . . . 1735
Example: Configuring an SRX5K-MPC on an SRX5000 Line Device to Support Express Path . . . 1737
Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path . . . 1739
Example: Configuring Express Path on an SRX5000 Line Device with IOC3 for IPv6 Traffic . . . 1742
Example: Configuring Low Latency . . . 1746

Part 25 Managing Flow-Based Processing for IPv6

Chapter 87 **Enabling IPv6 Flow-Based Processing . . . 1751**

IPv6 Advanced Flow . . . 1751
Understanding Sessions for IPv6 Flows . . . 1753
Understanding IPv6 Flow Processing on High-End SRX Series Devices . . . 1753
Enabling Flow-Based Processing for IPv6 Traffic . . . 1756
Using Filters to Display IPv6 Session and Flow Information for SRX Series Services Gateways . . . 1757

Chapter 88 **Managing IPv6 Packets . . . 1763**

The IPv6 Packet Header and SRX Series Overview . . . 1763
About the IPv6 Basic Packet Header . . . 1764
Understanding IPv6 Packet Header Extensions . . . 1766
About IPv6 Packet Header Verification Performed by the Flow Module for SRX Series Devices . . . 1767
Understanding Path MTU Messages for IPv6 Packets . . . 1767
Understanding How SRX Series Devices Handle Packet Fragmentation for IPv6 Flows . . . 1769
Understanding How SRX Series Devices Handle ICMPv6 Packets . . . 1769

Chapter 89 **Configuring IPv6 Dual-Stack . . . 1773**

Understanding IPv6 Dual-Stack Lite . . . 1773
Example: Configuring IPv6 Dual-Stack Lite . . . 1776

Copyright © 2016, Juniper Networks, Inc.

Part 26 Monitoring Flow-Based Sessions

Chapter 90 Monitoring Security Flow Sessions 1781

Monitoring Security Flow Sessions Overview 1781
Understanding How to Obtain Session Information for SRX Series Services Gateways 1782
Displaying Global Session Parameters for All SRX Series Services Gateways . . 1784
Displaying a Summary of Sessions for SRX Series Services Gateways 1785
Displaying Session and Flow Information About Sessions for SRX Series Services Gateways 1785
Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways 1786
Using Filters to Display Session and Flow Information for SRX Series Services Gateways 1787
Information Provided in Session Log Entries for SRX Series Services Gateways 1787

Chapter 91 Monitoring X2 Traffic By Configuring Mirror Filters 1793

Understanding X2 Traffic Monitoring 1793
X2 Traffic Monitoring Overview 1793
Limitations of X2 Traffic Monitoring 1795
X2 Traffic Terminology 1795
Example: Configuring Mirror Filters for X2 Traffic Monitoring 1796

Part 27 Configuring Packet-Based Forwarding

Chapter 92 Configuring Selective Stateless Packet-Based Services 1803

Understanding Packet-Based Processing 1803
Understanding Selective Stateless Packet-Based Services 1804
Selective Stateless Packet-Based Services Configuration Overview 1806
Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding 1808
Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding 1818

Part 28 Configuration Statements and Operational Commands

Chapter 93 Configuration Statements 1831

Chassis Configuration Statement Hierarchy 1833
Security Configuration Statement Hierarchy 1836
[edit security flow] Hierarchy Level 1837
[edit security forwarding-process] Hierarchy Level 1839
aging 1839
all-tcp 1840
allow-dns-reply 1840
allow-embedded-icmp 1841
application-services (Security Forwarding Process) 1842
apply-to-half-close-state 1843
ethernet-switching 1844
destination-header 1845

Copyright © 2016, Juniper Networks, Inc.

destination-port (Security Forwarding Options) . . . 1846
destination-prefix (Security Forwarding Options) . . . 1850
early-ageout . . . 1850
error . . . 1851
fin-invalidate-session . . . 1852
force-ip-reassembly . . . 1853
forwarding-process . . . 1853
fpc error . . . 1854
fru-poweron-sequence . . . 1856
gre-in . . . 1857
gre-out . . . 1858
high-watermark . . . 1858
hop-by-hop-header . . . 1859
icmpv6-malformed . . . 1860
inline-tap . . . 1860
interface-in (Security Forwarding Options) . . . 1861
interface-out (Security Forwarding Options) . . . 1861
ipv4-template (Services) . . . 1861
ipv6-extension-header . . . 1862
ipv6-extension-header-limit . . . 1863
ipv6-malformed-header . . . 1863
ipv6-template (Services) . . . 1864
low-latency . . . 1864
low-watermark . . . 1865
maximize-idp-sessions . . . 1865
mirror-filter (Security Forwarding Options) . . . 1866
mode (Security Forwarding Options) . . . 1867
no-sequence-check . . . 1868
no-tcp-reset . . . 1868
output (Security Forwarding Options) . . . 1869
packet-filter . . . 1870
packet-ordering-mode (Application Services) . . . 1871
pending-sess-queue-length . . . 1872
propagate-settings . . . 1872
protocol (Security Forwarding Options) . . . 1873
resource-manager . . . 1874
route-change-timeout . . . 1874
rst-invalidate-session . . . 1875
rst-sequence-check . . . 1875
sampling . . . 1876
services-offload . . . 1877
np-cache (Flexible PIC Concentrator) . . . 1878
source-port (Security Forwarding Options) . . . 1879
source-prefix (Security Forwarding Options) . . . 1883
syn-flood-protection-mode . . . 1883
tcp-initial-timeout . . . 1884
tcp-mss (Security Flow) . . . 1885
tcp-session . . . 1886
time-wait-state . . . 1887

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security) 1888
traceoptions (Security Flow) 1890
transport (Security Log) 1893
weight (Security) 1894

Chapter 94 Operational Commands 1895

clear firewall 1898
clear monitor security flow filter 1899
clear security flow ip-action 1900
clear security flow session all 1902
clear security flow session application 1903
clear security flow session destination-port 1905
clear security flow session family 1907
clear security flow session interface 1908
clear security flow session protocol 1910
clear security flow session resource-manager 1912
clear security flow session services-offload 1913
clear security flow session session-identifier 1916
clear security flow session source-port 1917
clear security flow session source-prefix 1919
clear security forward-options mirror filter 1921
monitor security flow file 1922
monitor security flow filter 1924
monitor security flow start 1926
monitor security flow stop 1927
show chassis environment (Security) 1928
show chassis fpc (View) 1932
show chassis fpc errors 1940
show chassis hardware (View) 1942
show chassis pic (Security) 1969
show chassis power 1971
show chassis power sequence 1973
show firewall (View) 1974
show interfaces (View Aggregated Ethernet) 1976
show interfaces (SRX Series) 1986
show interfaces diagnostics optics 2017
show interfaces flow-statistics 2022
show interfaces statistics (View) 2027
show interfaces swfabx 2028
show monitor security flow 2029
show security flow cp-session 2031
show security flow cp-session destination-port 2034
show security flow cp-session destination-prefix 2036
show security flow cp-session family 2038
show security flow cp-session protocol 2040
show security flow cp-session source-port 2043
show security flow cp-session source-prefix 2045
show security flow gate 2047
show security flow ip-action 2051

Copyright © 2016, Juniper Networks, Inc.

show security flow gate brief node . . . 2059
show security flow gate destination-port . . . 2065
show security flow gate destination-prefix . . . 2068
show security flow gate protocol . . . 2071
show security flow gate summary node . . . 2074
show security flow session . . . 2079
show security flow session brief node . . . 2085
show security flow session destination-port . . . 2089
show security flow session destination-prefix . . . 2093
show security flow session extensive node . . . 2097
show security flow session family . . . 2103
show security flow session interface . . . 2108
show security flow session nat . . . 2112
show security flow session policy-id . . . 2115
show security flow session protocol . . . 2118
show security flow session resource-manager . . . 2123
show security flow session services-offload . . . 2127
show security flow session session-identifier . . . 2132
show security flow session source-port . . . 2136
show security flow session source-prefix . . . 2140
show security flow session summary family . . . 2144
show security flow session summary node . . . 2146
show security flow session summary services-offload . . . 2152
show security flow session tunnel . . . 2156
show security flow statistics . . . 2162
show security flow status . . . 2164
show security forward-options mirror-filter . . . 2167
show security monitoring . . . 2169
show security policies . . . 2171
show security policies hit-count . . . 2179
show security resource-manager group active . . . 2182
show security resource-manager resource active . . . 2185
show security resource-manager settings . . . 2188
show security resource-manager summary . . . 2190
show security screen ids-option . . . 2191
show security screen statistics . . . 2196
show security softwires . . . 2204
show security zones . . . 2205
show security zones type . . . 2208

Guide 7 General Packet Radio Service Feature Guide for Security Devices

Part 29 Overview

Chapter 95 **Introduction to General Packet Radio Service . . . 2215**

GPRS Overview . . . 2215
Gp and Gn Interfaces . . . 2216
Gi Interface . . . 2216

Copyright © 2016, Juniper Networks, Inc.

Operational Modes 2217
GTP In-Service Software Upgrade 2218
Understanding Central Point Architecture Support for GTP 2218
Understanding GTP Handover Messages 2219

Part 30 Configuring GPRS Tunnel Protocol v1

Chapter 96 Configuring Policy-Based GTP 2223

Understanding Policy-Based GTP 2223
Example: Enabling GTP Inspection in Policies 2224

Chapter 97 Configuring GTP Inspection Objects 2229

Understanding GTP Inspection Objects 2229
Example: Creating a GTP Inspection Object 2229
Understanding GTP-U Inspection 2230

Chapter 98 Configuring GTP Message Filtering 2233

Understanding GTP Message Filtering 2233
Understanding GTP Message-Length Filtering 2233
Example: Setting the GTP Message Lengths 2234
Understanding GTP Message-Type Filtering 2235
Supported GTP Message Types 2235
Example: Permitting and Denying GTP Message Types 2238
Understanding GTP Message-Rate Limiting 2239
Understanding GTP Control Message Path Rate Limiting 2239
Example: Limiting the Message Rate and Path Rate for GTP Control Messages 2240
Example: Enabling GTP Sequence Number Validation 2244
Understanding GTP IP Fragmentation 2245

Chapter 99 Configuring GTP Information Elements 2247

Understanding GTP Information Elements 2247
Understanding GTP APN Filtering 2248
Example: Setting a GTP APN and a Selection Mode 2249
Understanding IMSI Prefix Filtering of GTP Packets 2250
Example: Setting a Combined IMSI Prefix and APN Filter 2251
Understanding R6, R7, R8, and R9 Information Elements Removal 2252
Supported R6 Information Elements 2252
Example: Removing R6 Information Elements from GTP Messages 2256
Supported R7 Information Elements 2257
Example: Removing R7 Information Elements from GTP Messages 2258
Supported R8 Information Elements 2259
Example: Removing R8 Information Elements from GTP Messages 2259
Supported R9 Information Elements 2260
Example: Removing R9 Information Elements from GTP Messages 2261
Understanding GTPv1 Information Element Removal 2262
Example: Removing GTPv1 Information Elements Using IE Number 2262

Copyright © 2016, Juniper Networks, Inc.

Chapter 100 Configuring NAT for GTP 2265
Understanding NAT for GTP 2265
Example: Configuring GTP Inspection in NAT 2266
Understanding Network Address Translation-Protocol Translation 2270
Example: Enhancing Traffic Engineering by Configuring NAT-PT Between an IPv4 and an IPv6 Endpoint with SCTP Multihoming 2271
Chapter 101 Configuring GGSN 2279
Understanding GGSN Redirection 2279
GGSN Pooling Scenarios Overview 2279
Understanding GGSN Pooling for Scenario 1 2279
Install Request Information to Remote SPU 2280
Workarounds for Scenario 1 2281
Understanding GGSN Pooling for Scenario 2 2281
Install Tunnel Information to Remote SPU 2282
Example: Configuring a GGSN Custom Policy 2283
Example: Configuring Custom Applications 2286
Part 31 Configuring GPRS Tunnel Protocol v2
Chapter 102 Configuring GTPv2 2291
Understanding GTPv2 2291
Understanding Policy-Based GTPv2 2293
Example: Enabling GTPv2 Inspection in Policies 2293
Understanding GTP Path Restart 2296
Example: Restarting a GTPv2 Path 2297
Understanding GTPv2 Tunnel Cleanup 2298
Example: Setting the Timeout Value for GTPv2 Tunnels 2299
Understanding GTPv2 Traffic Logging 2300
Example: Enabling GTPv2 Traffic Logging 2301
Chapter 103 Configuring GTPv2 Message Filtering 2303
Understanding GTPv2 Message Filtering 2303
Supported GTPv2 Message Types 2303
Example: Permitting and Denying GTPv2 Message Types 2307
Understanding GTPv2 Message-Length Filtering 2308
Example: Setting GTPv2 Message Lengths 2308
Understanding GTPv2 Message-Type Filtering 2309
Understanding GTPv2 Message-Rate Limiting 2310
Example: Limiting the GTPv2 Message Rate 2311
Chapter 104 GTPv2 Information Elements Overview 2313
Understanding GTPv2 Information Elements 2313
Understanding GTPv2 IMSI Prefix and APN Filtering 2313
Example: Setting a Combined GTPv2 IMSI Prefix and APN Filter 2315

Copyright © 2016, Juniper Networks, Inc.

Part 32 Configuring Stream Control Transmission Protocol

Chapter 105 Configuring SCTP 2319

Understanding Stream Control Transmission Protocol 2319
SCTP Features Overview 2323
Understanding Central Point Architecture Support for SCTP 2324
Understanding SCTP Behavior in Chassis Cluster 2324
Understanding SCTP Multihoming 2326
Understanding SCTP Multichunk Inspection 2327
SCTP Packet Structure Overview 2327
SCTP Configuration Overview 2329
Example: Configuring a GPRS SCTP Profile for Policy-Based Inspection to Reduce Security Risks 2329
Example: Configuring a Security Policy to Permit or Deny SCTP Traffic 2332

Part 33 Configuration Statements and Operational Commands

Chapter 106 Configuration Statements 2339

Security Configuration Statement Hierarchy 2340
[edit security gprs] Hierarchy Level 2341
[edit security forwarding-process] Hierarchy Level 2345
action (APN GTP) 2347
alarm-threshold (Security GPRS) 2348
apn 2349
application-services (Security Forwarding Process) 2350
association-timeout 2351
create-req 2351
delete-req 2352
drop (Security GTP) 2353
drop (Security SCTP) 2358
drop-threshold (Security GPRS) 2360
echo-req 2361
enable (GPRS GTP) 2361
enable-gtpu-distribution 2362
end-user-address-validated (GTP) 2362
forward 2363
gprs 2364
gprs-gtp-profile 2367
gprs-sctp-profile 2368
gtp 2369
gtp-in-gtp-denied 2371
handover-on-roaming-intf 2372
handshake-timeout 2372
imsi-prefix 2373
limit (Security SCTP) 2374
log (Security GTP) 2378
log (Security SCTP) 2379
max-message-length 2380
message-type 2381

Copyright © 2016, Juniper Networks, Inc.

min-message-length . . . 2382
multichunk-inspection . . . 2382
nullpdu . . . 2383
number . . . 2383
other . . . 2384
path-rate-limit . . . 2385
permit (Security SCTP) . . . 2386
profile (Security GTP) . . . 2387
profile (Security SCTP) . . . 2390
rate-limit (Security GTP) . . . 2394
remove-ie . . . 2395
req-timeout . . . 2395
restart-path . . . 2396
reverse . . . 2397
sctp . . . 2398
seq-number-validated (GTP) . . . 2403
timeout (Security GTP) . . . 2403
traceoptions (Security GTP) . . . 2404
traceoptions (Security SCTP) . . . 2406
u-tunnel-validated (GTP) . . . 2407
version (Security GTP) . . . 2408

Chapter 107 **Operational Commands . . . 2409**

clear gtp tunnels . . . 2410
clear security gprs gtp counters . . . 2411
clear security gprs sctp association . . . 2413
clear security gprs sctp counters . . . 2415
show gtp tunnels . . . 2416
show security gprs gtp counters . . . 2418
show security gprs gtp counters path-rate-limit . . . 2425
show security gprs gtp gsn statistics . . . 2426
show security gprs sctp association . . . 2427
show security gprs sctp counters . . . 2429

Guide 8 Interfaces Feature Guide for Security Devices

Part 34 Overview

Chapter 108 **Introduction to Interfaces . . . 2441**

Understanding Interfaces . . . 2441
Network Interfaces . . . 2442
Services Interfaces . . . 2443
Special Interfaces . . . 2445
Interface Naming Conventions . . . 2446
Understanding the Data Link Layer . . . 2447
Physical Addressing . . . 2448
Network Topology . . . 2448
Error Notification . . . 2448
Frame Sequencing . . . 2448
Flow Control . . . 2448

Copyright © 2016, Juniper Networks, Inc.

Data Link Sublayers . . . 2448
MAC Addressing . . . 2449
Monitoring Interfaces . . . 2449

Chapter 109 Configuring Interface Logical Properties . . . 2451

Understanding Interface Logical Properties . . . 2451
Understanding Protocol Families . . . 2452
Common Protocol Suites . . . 2452
Other Protocol Suites . . . 2452
Understanding IPv4 Addressing . . . 2453
IPv4 Classful Addressing . . . 2453
IPv4 Dotted Decimal Notation . . . 2454
IPv4 Subnetting . . . 2454
IPv4 Variable-Length Subnet Masks . . . 2455
Understanding IPv6 Address Space, Addressing, Address Format, and Address Types . . . 2455
Understanding IP Version 6 (IPv6) . . . 2455
Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them . . . 2456
IPv6 Address Scope . . . 2457
IPv6 Address Structure . . . 2457
Understanding IPv6 Address Space, Addressing, and Address Types . . . 2458
Understanding IPv6 Address Format . . . 2458
Limitations . . . 2459
Configuring the inet6 IPv6 Protocol Family . . . 2459
Enabling Flow-Based Processing for IPv6 Traffic . . . 2460
Configuring Flow Aggregation to Use Version 9 Flow Templates . . . 2461
Configuring the Traffic to Be Sampled . . . 2462
Configuring the Version 9 Template Properties . . . 2462
Restrictions . . . 2464
Fields Included in Each Template Type . . . 2465
inet Sampling Behavior . . . 2466
Verification . . . 2467
Examples: Configuring Version 9 Flow Templates . . . 2467
Understanding IPv6 Support VDSL2 Interfaces . . . 2470
Understanding MAC Limiting on Layer 3 Routing Interfaces . . . 2471
Overview . . . 2471
Limitations . . . 2473

Chapter 110 Understanding Interface Physical Properties . . . 2475

Understanding Interface Physical Properties . . . 2475
Understanding Bit Error Rate Testing . . . 2476
Understanding Interface Clocking . . . 2477
Data Stream Clocking . . . 2478
Explicit Clocking Signal Transmission . . . 2478
Understanding Frame Check Sequences . . . 2478
Cyclic Redundancy Checks and Checksums . . . 2478
Two-Dimensional Parity . . . 2479
MTU Default and Maximum Values . . . 2479
Understanding Jumbo Frames Support for Ethernet Interfaces . . . 2481

Copyright © 2016, Juniper Networks, Inc.

Chapter 111 Configuring VLAN Tagging 2483
Understanding Virtual LANs 2483
VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices 2485
Configuring VLAN Tagging 2485
Configuring Single-Tag Framing 2486
Configuring Dual Tagging 2486
Configuring Mixed Tagging 2486
Configuring Mixed Tagging Support for Untagged Packets 2487

Part 35 Configuring DS1 Interfaces

Chapter 112 Configuring DS1 Interfaces 2491
Understanding T1 and E1 Interfaces 2491
T1 Overview 2491
E1 Overview 2492
T1 and E1 Signals 2492
Encoding 2492
AMI Encoding 2492
B8ZS and HDB3 Encoding 2492
T1 and E1 Framing 2493
ESF Framing for T1 2493
T1 and E1 Loopback Signals 2493
Example: Configuring a T1 Interface 2494
Example: Deleting a T1 Interface 2496

Part 36 Configuring DSL Interfaces

Chapter 113 Configuring VDSL2 Interfaces 2501
VDSL2 Interface Technology Overview 2501
VDSL2 Network Deployment Topology 2502
VDSL2 Interface Support on SRX Series Devices 2503
VDSL2 Interface Compatibility with ADSL Interfaces 2504
VDSL2 Interfaces Supported Profiles 2505
VDSL2 Interfaces Supported Features 2506
Example: Configuring VDSL2 Interfaces (Basic) 2507
Example: Configuring VDSL2 Interfaces (Detail) 2513

Part 37 Configuring Ethernet Interfaces

Chapter 114 Performing Initial Configuration on Ethernet Interfaces 2541
Understanding Ethernet Interfaces 2541
Ethernet Access Control and Transmission 2542
Collisions and Detection 2542
Collision Detection 2542
Backoff Algorithm 2542
Collision Domains and LAN Segments 2543
Repeaters 2543
Bridges and Switches 2543

Copyright © 2016, Juniper Networks, Inc.

Broadcast Domains . . . 2544
Ethernet Frames . . . 2544
Understanding Static ARP Entries on Ethernet Interfaces . . . 2545
Understanding Promiscuous Mode on Ethernet Interface . . . 2545
Understanding Promiscuous Mode on the SRX5K-MPC . . . 2546
Example: Creating an Ethernet Interface . . . 2546
Example: Deleting an Ethernet Interface . . . 2547
Example: Configuring Static ARP Entries on Ethernet Interfaces . . . 2548
Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) . . . 2551
Example: Configuring Promiscuous Mode on the SRX5K-MPC . . . 2552

Chapter 115 Configuring Aggregated Ethernet Interfaces . . . 2557

Understanding Aggregated Ethernet Interfaces . . . 2557
LAGs . . . 2558
LACP . . . 2558
Aggregated Ethernet Interfaces Configuration Overview . . . 2560
Understanding the Aggregated Ethernet Interfaces Device Count . . . 2560
Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device . . . 2561
Understanding Physical Interfaces for Aggregated Ethernet Interfaces . . . 2562
Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces . . . 2562
Understanding Aggregated Ethernet Interface Link Speed . . . 2563
Example: Configuring Aggregated Ethernet Link Speed . . . 2564
Understanding Minimum Links for Aggregated Ethernet Interfaces . . . 2565
Example: Configuring Aggregated Ethernet Minimum Links . . . 2565
Understanding Aggregated Ethernet Interface Removal . . . 2566
Example: Deleting Aggregated Ethernet Interfaces . . . 2566
Example: Deleting Aggregated Ethernet Interface Contents . . . 2567
Verifying Aggregated Ethernet Interfaces . . . 2568
Verifying Aggregated Ethernet Interfaces (terse) . . . 2568
Verifying Aggregated Ethernet Interfaces (extensive) . . . 2569
Understanding VLAN Tagging for Aggregated Ethernet Interfaces . . . 2570
Understanding Promiscuous Mode for Aggregated Ethernet Interfaces . . . 2570

Chapter 116 Configuring Link Aggregation Control Protocol . . . 2571

Understanding LACP on Standalone Devices . . . 2571
Example: Configuring LACP on Standalone Devices . . . 2572
Verifying LACP on Standalone Devices . . . 2573
Verifying LACP Statistics . . . 2573
Verifying LACP Aggregated Ethernet Interfaces . . . 2574
Understanding LACP on Chassis Clusters . . . 2575
Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups . . . 2575
Minimum Links . . . 2576
Sub-LAGs . . . 2576
Supporting Hitless Failover . . . 2576
Managing Link Aggregation Control PDUs . . . 2577
Example: Configuring LACP on Chassis Clusters . . . 2577
Verifying LACP on Redundant Ethernet Interfaces . . . 2579

Copyright © 2016, Juniper Networks, Inc.

LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs) . . . 2580
LAG and LACP Support on the SRX5000 Module Port Concentrator 2580
LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode 2581
Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3 2582

Chapter 117 **Configuring Gigabit Ethernet Physical Interface Modules 2587**

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM 2587
Supported Features 2587
Interface Names and Settings 2588
Available Link Speeds and Modes 2588
Link Settings 2589
Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface 2589
Understanding the 2-Port 10-Gigabit Ethernet XPIM 2595
Supported Features 2596
Interface Names and Settings 2596
Copper and Fiber Operating Modes 2597
Link Speeds 2597
Link Settings 2597
Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface 2598

Chapter 118 **Configuring Ethernet OAM Link Fault Management 2603**

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways 2603
Example: Configuring Ethernet OAM Link Fault Management 2605

Chapter 119 **Configuring Power over Ethernet 2611**

Understanding Power over Ethernet 2611
SRX Series Services Gateway PoE Specifications 2611
PoE Classes and Power Ratings 2612
PoE Options 2613
Example: Configuring PoE on All Interfaces 2613
Example: Configuring PoE on an Individual Interface 2615
Example: Disabling a PoE Interface 2618

Part 38 **Configuring Interface Encapsulation**

Chapter 120 **Interface Encapsulation Overview 2623**

Understanding Physical Encapsulation on an Interface 2623
Understanding Frame Relay Encapsulation on an Interface 2624
Virtual Circuits 2624
Switched and Permanent Virtual Circuits 2624
Data-Link Connection Identifiers 2625
Congestion Control and Discard Eligibility 2625
Understanding Point-to-Point Protocol 2625
Link Control Protocol 2626
PPP Authentication 2626
Network Control Protocols 2627
Magic Numbers 2627

Copyright © 2016, Juniper Networks, Inc.

CSU/DSU Devices . . . 2628
Understanding High-Level Data Link Control . . . 2628
HDLC Stations . . . 2628
HDLC Operational Modes . . . 2629

Chapter 121 Configuring Point-to-Point Protocol over Ethernet . . . 2631

Understanding Point-to-Point Protocol over Ethernet . . . 2631
PPPoE Discovery Stage . . . 2632
PPPoE Session Stage . . . 2633
Understanding PPPoE Interfaces . . . 2634
Example: Configuring PPPoE Interfaces . . . 2635
Understanding PPPoE Ethernet Interfaces . . . 2641
Example: Configuring PPPoE Encapsulation on an Ethernet Interface . . . 2641
Understanding CHAP Authentication on a PPPoE Interface . . . 2642
Example: Configuring CHAP Authentication on a PPPoE Interface . . . 2643
Verifying Credit-Flow Control . . . 2644
Verifying PPPoE Interfaces . . . 2645
Verifying R2CP Interfaces . . . 2646
Displaying Statistics for PPPoE . . . 2647
Setting Tracing Options for PPPoE . . . 2647

Chapter 122 Configuring PPPoE-Based Radio-to-Router Protocol . . . 2649

PPPoE-Based Radio-to-Router Protocols Overview . . . 2649
Understanding the PPPoE-Based Radio-to-Router Protocol . . . 2650
Configuring PPPoE-Based Radio-to-Router Protocols . . . 2652
Example: Configuring the PPPoE-Based Radio-to-Router Protocol . . . 2652
Credit Flow Control for PPPoE . . . 2655
PPPoE Credit-Based Flow Control Configuration . . . 2655

Chapter 123 Configuring R2CP Radio-to-Router Protocol . . . 2657

R2CP Radio-to-Router Protocol Overview . . . 2657
Configuring the R2CP Radio-to-Router Protocol . . . 2658

Part 39 Configuration Statements and Operational Commands

Chapter 124 Configuration Statements . . . 2665

Chassis Configuration Statement Hierarchy . . . 2668
Class-of-Service Configuration Statement Hierarchy . . . 2671
Interfaces Configuration Statement Hierarchy . . . 2675
PoE Configuration Statement Hierarchy . . . 2690
accept-source-mac . . . 2691
access-point-name . . . 2692
apply-groups . . . 2692
arp-resp . . . 2693
authentication-method (Interfaces) . . . 2693
bandwidth (Interfaces) . . . 2694
bundle (Interfaces) . . . 2694
cbr rate . . . 2695
cellular-options . . . 2695
classifiers (CoS) . . . 2696

Copyright © 2016, Juniper Networks, Inc.

client-identifier (Interfaces) . . . 2697
code-points (CoS) . . . 2697
compression-device (Interfaces) . . . 2698
credit (Interfaces) . . . 2698
data-rate . . . 2699
disable (PoE) . . . 2699
dhcp (Interfaces) . . . 2700
duration (PoE) . . . 2700
encapsulation (Interfaces) . . . 2701
family inet (Interfaces) . . . 2702
family inet6 . . . 2705
flag (Interfaces) . . . 2707
flexible-vlan-tagging (Interfaces) . . . 2708
flow-control (Interfaces) . . . 2708
flow-monitoring (Services) . . . 2709
forwarding-classes (CoS) . . . 2710
fpc (Interfaces) . . . 2711
gratuitous-arp-reply . . . 2711
gsm-options . . . 2712
guard-band (PoE) . . . 2712
hub-assist . . . 2713
inline-jflow (Forwarding Options) . . . 2713
interface (PIC Bundle) . . . 2714
interface (PoE) . . . 2715
interfaces (CoS) . . . 2716
interval (Interfaces) . . . 2717
interval (PoE) . . . 2717
ipv4-template (Services) . . . 2718
ipv6-template (Services) . . . 2718
lacp (Interfaces) . . . 2719
latency (Interfaces) . . . 2719
lease-time . . . 2720
line-rate (Interfaces) . . . 2720
link-speed (Interfaces) . . . 2721
loopback (Interfaces) . . . 2721
loss-priority (CoS Loss Priority) . . . 2722
loss-priority (CoS Rewrite Rules) . . . 2723
loss-priority-maps (CoS Interfaces) . . . 2724
loss-priority-maps (CoS) . . . 2724
management (PoE) . . . 2725
maximum-power (PoE) . . . 2725
media-type (Interfaces) . . . 2726
minimum-links (Interfaces) . . . 2726
native-vlan-id (Interfaces) . . . 2727
next-hop-tunnel . . . 2727
no-dns-propagation . . . 2728
option-refresh-rate (Services) . . . 2728
pic-mode (Chassis T1 Mode) . . . 2729
periodic (Interfaces) . . . 2730

Copyright © 2016, Juniper Networks, Inc.

ppp-over-ether . . . 2730
pppoe . . . 2731
pppoe-options . . . 2732
priority (PoE) . . . 2733
profile (Access) . . . 2734
profiles . . . 2737
promiscuous-mode (Interfaces) . . . 2738
quality (Interfaces) . . . 2738
r2cp . . . 2739
radio-router (Interfaces) . . . 2740
redundancy-group (Interfaces) . . . 2741
redundant-ether-options . . . 2741
redundant-parent (Interfaces Fast Ethernet) . . . 2742
redundant-parent (Interfaces Gigabit Ethernet) . . . 2742
resource (Interfaces) . . . 2743
retransmission-attempt . . . 2743
retransmission-interval (Interfaces) . . . 2744
roaming-mode . . . 2744
scheduler-map (CoS Virtual Channels) . . . 2745
select-profile . . . 2745
server-address . . . 2746
shaping-rate (CoS Interfaces) . . . 2747
simple-filter (Interfaces) . . . 2748
sip-password . . . 2748
sip-user-id . . . 2749
source-address-filter (Interfaces) . . . 2750
source-filtering (Interfaces) . . . 2751
speed (Interfaces) . . . 2751
telemetries (PoE) . . . 2752
template-refresh-rate (Services) . . . 2752
threshold (Interfaces) . . . 2753
traceoptions (Interfaces) . . . 2753
update-server . . . 2754
vbr rate . . . 2754
vdsl-profile . . . 2755
vendor-id (Interfaces) . . . 2755
vlan-tagging (Interfaces) . . . 2756
web-authentication (Interfaces) . . . 2757

Chapter 125 Operational Commands . . . 2759

clear dhcpv6 server binding (Local Server) . . . 2761
clear ethernet-switching statistics mac-learning . . . 2762
clear interfaces statistics swfabx . . . 2763
clear ipv6 neighbors . . . 2764
clear lacp statistics interfaces . . . 2765
restart (Reset) . . . 2766
show chassis fpc (View) . . . 2771
show chassis hardware (View) . . . 2779
show ethernet-switching mac-learning-log (View) . . . 2806

Copyright © 2016, Juniper Networks, Inc.

show ethernet-switching table (View) . . . 2808
show igmp-snooping route (View) . . . 2813
show interfaces (SRX Series) . . . 2815
show interfaces diagnostics optics . . . 2846
show interfaces flow-statistics . . . 2851
show interfaces queue . . . 2856
show interfaces statistics (View) . . . 2859
show interfaces terse zone . . . 2860
show ipv6 neighbors . . . 2861
show lacp interfaces (View) . . . 2863
show lacp statistics interfaces (View) . . . 2867
show oam ethernet link-fault-management . . . 2868
show poe controller (View) . . . 2873
show pppoe interfaces . . . 2874
show pppoe statistics . . . 2878
show poe telemetries . . . 2880
show services accounting . . . 2882
show services accounting aggregation (View) . . . 2885
show services accounting aggregation template (View) . . . 2886
show services accounting flow-detail (View) . . . 2887

Guide 9 Layer 2 Bridging and Transparent Mode for Security Devices

Part 40 Overview

Chapter 126 **Introduction to Layer 2 Bridging and Switching . . . 2893**

Layer 2 Transparent and Switching Overview . . . 2893

Part 41 Configuring Layer 2 Bridging and Transparent Mode

Chapter 127 **Configuring Bridging and Transparent Mode . . . 2897**

Layer 2 Bridging and Transparent Mode Overview . . . 2897
Layer 2 Bridging Exceptions on SRX Series Devices . . . 2898
Layer 2 Transparent Mode on the SRX5000 Line Module Port Concentrator . . . 2899
Understanding VLANs . . . 2900
Example: Configuring VLANs . . . 2901
Enhanced Layer 2 CLI Configuration Statement and Command Changes . . . 2902

Chapter 128 **Configuring Interfaces . . . 2905**

Understanding Transparent Mode Conditions . . . 2905
Understanding Layer 2 Interfaces . . . 2906
Example: Configuring Layer 2 Logical Interfaces . . . 2906
Understanding Integrated Routing and Bridging Interfaces . . . 2907
Example: Configuring an IRB Interface . . . 2909
Understanding Mixed Mode (Layer 2 and Layer 3) . . . 2910
Example: Improving Security Services by Configuring an SRX Series Device Using Mixed Mode (Layer 2 and Layer 3) . . . 2913

Copyright © 2016, Juniper Networks, Inc.

Chapter 129 Configuring Security Zones and Security Policies 2921
Understanding Layer 2 Security Zones 2921
Example: Configuring Layer 2 Security Zones 2922
Understanding Security Policies in Transparent Mode 2923
Example: Configuring Security Policies in Transparent Mode 2924
Understanding Firewall User Authentication in Transparent Mode 2926
Chapter 130 Configuring Layer 2 Forwarding Tables 2927
Understanding Layer 2 Forwarding Tables 2927
Example: Configuring the Default Learning for Unknown MAC Addresses 2929
Chapter 131 Configuring Layer 2 Transparent Mode Chassis Clusters 2931
Understanding Layer 2 Transparent Mode Chassis Clusters 2931
Example: Configuring Redundant Ethernet Interfaces for Layer 2 Transparent Mode Chassis Clusters 2933
Chapter 132 Configuring IP Spoofing in Layer 2 Transparent Mode 2935
Understanding IP Spoofing in Layer 2 Transparent Mode 2935
Configuring IP Spoofing in Layer 2 Transparent Mode 2936
Chapter 133 Configuring Class of Service in Transparent Mode 2939
Class of Service Functions in Transparent Mode Overview 2939
Understanding BA Traffic Classification on Transparent Mode Devices 2940
Example: Configuring BA Classifiers on Transparent Mode Devices 2941
Understanding Rewrite of Packet Headers on Transparent Mode Devices 2943
Example: Configuring Rewrite Rules on Transparent Mode Devices 2944
Chapter 134 Configuring IPv6 Flows 2947
Understanding IPv6 Flows in Transparent Mode 2947
Enabling Flow-Based Processing for IPv6 Traffic 2948
Example: Configuring Transparent Mode for IPv6 Flows 2950
Chapter 135 Configuring Secure Wire 2955
Understanding Secure Wire 2955
Example: Simplifying SRX Series Device Deployment with Secure Wire over Access Mode Interfaces 2957
Example: Simplifying SRX Series Device Deployment with Secure Wire over Trunk Mode Interfaces 2960
Example: Simplifying SRX Series Device Deployment with Secure Wire over Aggregated Interface Member Links 2964
Example: Simplifying Chassis Cluster Deployment with Secure Wire over Redundant Ethernet Interfaces 2968
Example: Simplifying Chassis Cluster Deployment with Secure Wire over Aggregated Redundant Ethernet Interfaces 2972

Copyright © 2016, Juniper Networks, Inc.

Part 42 Configuring Layer 2 Next generation Switching Mode

Chapter 136 Configuring Switching Modes 2983

Understanding Switching Modes 2983
Ethernet Ports Switching Overview 2984
Supported Devices and Ports 2984
Integrated Bridging and Routing 2984
Types of Switch Ports 2985
uPIM in a Daisy Chain 2985
Example: Configuring Switching Modes 2985
Verifying Switching Mode Configuration 2986

Chapter 137 Configuring VLANs 2989

Understanding VLANs 2989
Example: Configuring VLANs 2991

Chapter 138 Configuring Link Aggregation Control Protocol 2993

Understanding Link Aggregation Control Protocol 2993
Link Aggregation Benefits 2994
Link Aggregation Configuration Guidelines 2994
Example: Configuring Link Aggregation Control Protocol 2997

Chapter 139 Configuring Port Security 2999

Port Security Overview 2999
Understanding MAC Limiting 2999
Example: Configuring MAC Limiting 3000

Part 43 Configuration Statements and Operational Commands

Chapter 140 Configuration Statements 3005

[edit security forwarding-options] Hierarchy Level 3006
Access Configuration Statement Hierarchy 3007
Class-of-Service Configuration Statement Hierarchy 3015
Interfaces Configuration Statement Hierarchy 3019
VLANS Configuration Statement Hierarchy 3035
authentication-order (Access Profile) 3040
code-points (CoS) 3041
destination-address (Security Policies) 3042
domain-type (VLANs) 3042
encapsulation (Interfaces) 3043
ethernet-switching 3044
family inet (Interfaces) 3045
family inet6 3048
flow (Security Flow) 3051
forwarding-classes (CoS) 3053
global-mac-table-aging-time (Protocols) 3054
global-mac-limit (Protocols) 3055
global-mode (Protocols) 3056
global-no-mac-learning (Protocols) 3056
host-inbound-traffic 3057
inet6 (Security Forwarding Options) 3058

Copyright © 2016, Juniper Networks, Inc.

interfaces (CoS) . . . 3059
interfaces (Security Zones) . . . 3060
interface (Switching Options) . . . 3061
interface (VLANs) . . . 3062
l2-learning (Protocols) . . . 3063
loss-priority (CoS Loss Priority) . . . 3064
match (Security Policies) . . . 3065
native-vlan-id (Interfaces) . . . 3066
peer-selection-service . . . 3067
pgcp-service . . . 3068
policy (Security Policies) . . . 3069
port (Access RADIUS) . . . 3071
profile (Access) . . . 3072
radius-server (Access) . . . 3075
redundancy-group (Interfaces) . . . 3076
secure-wire . . . 3076
security-zone . . . 3077
shaping-rate (CoS Interfaces) . . . 3079
source-address (Access RADIUS) . . . 3080
source-address (Security Policies) . . . 3081
static-mac (VLANs) . . . 3082
switch-options (VLANs) . . . 3083
system-services (Security Zones Interfaces) . . . 3084
unframed | no-unframed (Interfaces) . . . 3085
vlans . . . 3086
vlan-id (VLAN) . . . 3092
vlan members (VLANs) . . . 3093
vlan-tagging (Interfaces) . . . 3094

Chapter 141 **Operational Commands . . . 3095**

clear oam ethernet connectivity-fault-management path-database . . . 3096
clear oam ethernet connectivity-fault-management statistics . . . 3097
clear security flow ip-action . . . 3098
clear security flow session family . . . 3100
show ethernet-switching mac-learning-log (View) . . . 3101
show ethernet-switching table (View) . . . 3103
show igmp-snooping route (View) . . . 3108
show igmp-snooping vlans (View) . . . 3110
show interfaces (SRX Series) . . . 3112
show oam ethernet connectivity-fault-management adjacencies . . . 3143
show oam ethernet connectivity-fault-management forwarding-state . . . 3144
show oam ethernet connectivity-fault-management interfaces . . . 3146
show oam ethernet connectivity-fault-management mep-database . . . 3148
show oam ethernet connectivity-fault-management mep-statistics . . . 3152
show oam ethernet connectivity-fault-management mip . . . 3155
show oam ethernet connectivity-fault-management path-database . . . 3156
show oam ethernet connectivity-fault-management routes . . . 3158
show oam ethernet link-fault-management . . . 3160
show security flow gate family . . . 3165

Copyright © 2016, Juniper Networks, Inc.

show security flow ip-action . . . 3167
show security flow session family . . . 3175
show security flow statistics . . . 3180
show security flow status . . . 3182
show security forward-options secure-wire . . . 3185
show security policies . . . 3187
show security zones . . . 3195
show vlans . . . 3198

Guide 10 Logical Systems Feature Guide for Security Devices

Part 44 Overview

Chapter 142 Introduction to Logical Systems . . . 3205

Understanding Logical Systems for SRX Series Services Gateways . . . 3205
Understanding the Fundamentals and Constraints of Logical Systems . . . 3208
Understanding Licenses for Logical Systems on SRX Series Devices . . . 3209
Understanding the Interconnect Logical System and Logical Tunnel Interfaces . . . 3210
Understanding Flow in Logical Systems for SRX Series Devices . . . 3211
Understanding Junos OS SRX Series Services Gateways Architecture . . . 3213
Session Creation for Devices Running Logical Systems . . . 3214
Understanding Flow on Logical Systems . . . 3214
Understanding Packet Classification . . . 3214
Handling Pass-Through Traffic for Logical Systems . . . 3215
Pass-Through Traffic Within a Logical System . . . 3215
Pass-Through Traffic Between Logical Systems . . . 3215
Handling Self-Traffic . . . 3216
Self-Initiated Traffic . . . 3216
Traffic Terminated on a Logical System . . . 3217
Understanding Session and Gate Limitation Control . . . 3218
Understanding Sessions . . . 3218
About Configuring Sessions . . . 3218

Chapter 143 Understanding Master Logical Systems . . . 3221

Understanding the Master Logical System and the Master Administrator Role . . . 3221
SRX Series Logical System Master Administrator Configuration Tasks Overview . . . 3222

Chapter 144 Understanding User Logical Systems . . . 3225

User Logical System Configuration Overview . . . 3225
Understanding User Logical Systems and the User Logical System Administrator Role . . . 3227
Example: Configuring User Logical Systems . . . 3228

Copyright © 2016, Juniper Networks, Inc.

Part 45 Getting Started for Master Administrators

Chapter 145 Configuring Device for Master Logical Systems 3241

Example: Configuring a Root Password for the Device (Master Administrators Only) 3241
Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only) 3242

Part 46 Configuring Security Features

Chapter 146 Configuring Master Logical System Security Profiles 3253

Understanding Logical System Security Profiles (Master Administrators Only) 3253
Logical Systems Security Profiles 3254
How the System Assesses Resources Assignment and Use Across Logical Systems 3254
Cases: Assessments of Reserved Resources Assigned Through Security Profiles 3256
Example: Configuring Logical Systems Security Profiles (Master Administrators Only) 3258

Chapter 147 Configuring Master Logical System Security Features 3267

Understanding Logical System Firewall Authentication 3267
Example: Configuring Access Profiles (Master Administrators Only) 3269
Example: Configuring Security Features for the Master Logical System 3271
IDP in Logical Systems Overview 3276
IDP Policies 3277
IDP Installation and Licensing for Logical Systems 3277
Understanding IDP Features in Logical Systems 3278
Rulebases 3278
Protocol Decoders 3278
SSL Inspection 3279
Inline Tap Mode 3279
Multi-Detectors 3279
Logging and Monitoring 3279
Example: Configuring an IDP Policy for the Master Logical System 3281
Understanding Logical System Application Identification Services 3286
Understanding Logical System Application Firewall Services 3287
Example: Configuring Application Firewall Services for a Master Logical System 3288
Understanding Logical System Application Tracking Services 3292
Understanding Route-Based VPN Tunnels in Logical Systems 3293
Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) 3294

Chapter 148 Configuring User Logical System Security Features 3299

Understanding Logical System Zones 3299
Example: Configuring Zones for a User Logical System 3301
Understanding Logical System Screen Options 3304
Example: Configuring Screen Options for a User Logical System 3304

Copyright © 2016, Juniper Networks, Inc.

Understanding Logical System Security Policies 3306
Security Policies in Logical Systems 3306
Application Timeouts 3307
Security Policy Allocation 3307
Example: Configuring Security Policies in a User Logical System 3308
Understanding Logical System Firewall Authentication 3311
Example: Configuring Firewall Authentication for a User Logical System 3313
IDP in Logical Systems Overview 3317
IDP Policies 3317
IDP Installation and Licensing for Logical Systems 3318
Understanding IDP Features in Logical Systems 3318
Rulebases 3318
Protocol Decoders 3319
SSL Inspection 3319
Inline Tap Mode 3319
Multi-Detectors 3319
Logging and Monitoring 3320
Example: Configuring an IDP Policy for a User Logical System 3321
Example: Enabling IDP in a User Logical System Security Policy 3323
Understanding Logical System Application Identification Services 3326
Example: Configuring Application Firewall Services for a User Logical System 3327
Understanding Logical System Application Tracking Services 3330
Example: Configuring AppTrack for a User Logical System 3331
Understanding Route-Based VPN Tunnels in Logical Systems 3334
Example: Configuring a Route-Based VPN Tunnel in a User Logical System . . 3335

Part 47 Configuring Routing and Interfaces Features

Chapter 149 **Configuring Master Logical System Routing and Interfaces 3341**

Understanding Logical System Interfaces and Routing Instances 3341
Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) 3342
Example: Configuring OSPF Routing Protocol for the Master Logical System . . 3350

Chapter 150 **Configuring User Logical System Routing, Interfaces, and NAT Features 3355**

Understanding Logical System Network Address Translation 3355
Example: Configuring Network Address Translation for a User Logical System 3356
Understanding Logical System Interfaces and Routing Instances 3359
Example: Configuring Interfaces and Routing Instances for a User Logical System 3360
Example: Configuring OSPF Routing Protocol for a User Logical System 3362

Copyright © 2016, Juniper Networks, Inc.

Part 48 Configuring Logical Systems in Chassis Cluster

Chapter 151 Configuring Logical Systems When Device is in Chassis Cluster Mode . . 3369

Understanding Logical Systems in the Context of Chassis Cluster 3369
Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) 3370
Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) 3403

Part 49 Configuring IPv6 for Logical Systems

Chapter 152 Configuring IPv6 Addresses for Logical Systems 3439

IPv6 Addresses in Logical Systems Overview 3439
Understanding IPv6 Dual-Stack Lite in Logical Systems 3440
Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) 3441
Example: Configuring IPv6 Zones for a User Logical System 3449
Example: Configuring IPv6 Security Policies for a User Logical System 3452
Example: Configuring IPv6 Dual-Stack Lite for a User Logical System 3456

Part 50 Configuring System Resources Allocation

Chapter 153 System Resources Allocation (Master Administrators Only) 3461

Understanding CPU Allocation and Control 3461
CPU Control 3462
Reserved CPU Utilization Quota for Logical Systems 3462
CPU Control Target 3463
Shared CPU Resources and CPU Quotas 3463
CPU Utilization Scenario 1 3464
CPU Utilization Scenario 2 3464
CPU Utilization Scenario 3 3464
Monitoring CPU Utilization 3465
Example: Configuring CPU Utilization (Master Administrators Only) 3465
Example: Deleting an SRX Series Services Gateway Logical System (Master Administrators Only) 3468

Part 51 Troubleshooting

Chapter 154 Troubleshooting Logical Systems (Master Administrators Only) 3475

Understanding Security Logs and Logical Systems 3475
Understanding Data Path Debugging for Logical Systems 3476
Performing Tracing for Logical Systems (Master Administrators Only) 3477
Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only) 3481

Part 52 Configuration Statements and Operational Commands

Chapter 155 Configuration Statements 3485

Chassis Configuration Statement Hierarchy 3487
Logical-Systems Configuration Statement Hierarchy 3490

Copyright © 2016, Juniper Networks, Inc.

Security Configuration Statement Hierarchy 3492
System Configuration Statement Hierarchy 3493
[edit security address-book] Hierarchy Level 3524
[edit security application-firewall] Hierarchy Level 3524
[edit security application-tracking] Hierarchy Level 3525
[edit security datapath-debug] Hierarchy Level 3526
[edit security firewall-authentication] Hierarchy Level 3527
[edit security flow] Hierarchy Level 3527
[edit security idp] Hierarchy Level 3528
[edit security ike] Hierarchy Level 3537
[edit security ipsec] Hierarchy Level 3539
[edit security log] Hierarchy Level 3541
[edit security nat] Hierarchy Level 3542
[edit security policies] Hierarchy Level 3545
[edit security screen] Hierarchy Level 3550
[edit security softwires] Hierarchy Level 3552
[edit security zones] Hierarchy Level 3553
address-book 3555
address-book (System) 3556
appfw-profile (System) 3557
appfw-rule 3558
appfw-rule-set 3559
application-firewall 3560
application-tracking 3561
auth-entry 3562
cluster (Chassis) 3563
cpu 3565
datapath-debug 3566
dslite-softwire-initiator 3568
file (System Logging) 3569
firewall-authentication (Security) 3571
flow (Security Flow) 3572
flow-gate 3574
flow-session 3575
idp (Security) 3577
idp-policy 3585
ike (Security) 3586
ipsec (Security) 3588
log (Security) 3591
logical-system (System Security Profile) 3593
logical-systems (All) 3594
nat 3595
nat-cone-binding 3599
nat-destination-pool 3600
nat-destination-rule 3601
nat-interface-port-ol (System) 3602
nat-nopat-address 3603
nat-pat-address 3604
nat-pat-portnum 3605

Copyright © 2016, Juniper Networks, Inc.

nat-port-ol-ipnumber . . . 3606
nat-rule-referenced-prefix (System) . . . 3607
nat-source-pool . . . 3608
nat-source-rule . . . 3609
nat-static-rule . . . 3610
policies . . . 3611
policy (System Security Profile) . . . 3616
policy-with-count . . . 3617
profile (Access) . . . 3618
purging . . . 3621
root-authentication . . . 3622
root-logical-system . . . 3623
scheduler (System Security Profile) . . . 3624
screen (Security) . . . 3625
security-profile . . . 3628
security-profile-resources . . . 3631
softwires . . . 3632
zone (System Security Profile) . . . 3633
zones . . . 3634

Chapter 156 Operational Commands . . . 3637

clear security application-firewall rule-set statistics logical-system . . . 3639
clear security dns-cache . . . 3640
request security datapath-debug capture start . . . 3641
request security datapath-debug capture stop . . . 3642
set chassis cluster cluster-id node reboot . . . 3643
show chassis cluster status . . . 3644
show log . . . 3647
show security application-firewall rule-set . . . 3651
show security application-firewall rule-set logical-system . . . 3654
show security application-tracking counters . . . 3657
show security datapath-debug capture . . . 3658
show security datapath-debug counter . . . 3659
show security dns-cache . . . 3660
show security firewall-authentication history . . . 3662
show security firewall-authentication users . . . 3664
show security flow session . . . 3666
show security idp logical-system policy-association . . . 3672
show security ike security-associations . . . 3673
show security ipsec security-associations . . . 3681
show security match-policies . . . 3693
show security nat destination rule . . . 3698
show security nat destination summary . . . 3701
show security nat source rule . . . 3703
show security nat source summary . . . 3707
show security nat static rule . . . 3709
show security policies . . . 3713
show security screen statistics . . . 3721
show system security-profile . . . 3729

Copyright © 2016, Juniper Networks, Inc.

show security softwires 3734
show security zones 3735

Guide 11 MPLS Feature Guide for Security Devices

Part 53 Overview

Chapter 157 Introduction to MPLS 3743

MPLS Overview 3743
Label Switching 3744
Label-Switched Paths 3744
Label-Switching Routers 3745
Labels 3746
Label Operations 3746
Penultimate Hop Popping 3747
LSP Establishment 3747
Static LSPs 3747
Dynamic LSPs 3747
MPLS Configuration Overview 3748
Example: Deleting Security Services 3748
Example: Enabling MPLS 3750

Part 54 Configuring Traffic Engineering

Chapter 158 Configuring MPLS Traffic Engineering and Signaling Protocols 3755

MPLS Traffic Engineering and Signaling Protocols Overview 3755
Understanding the LDP Signaling Protocol 3756
Example: Configuring LDP-Signaled LSPs 3757
Understanding the RSVP Signaling Protocol 3761
RSVP Fundamentals 3762
Bandwidth Reservation Requirement 3762
Explicit Route Objects 3762
Constrained Shortest Path First 3763
Link Coloring 3764
Example: Configuring RSVP-Signaled LSPs 3765
Understanding Point-to-Multipoint LSPs 3768
Point-to-Multipoint LSP Configuration Overview 3770
Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP 3770

Part 55 Configuring MPLS VPNs

Chapter 159 Introduction to MPLS VPNs 3791

MPLS VPN Overview 3791
MPLS VPN Topology 3791
MPLS VPN Routing 3793
VRF Instances 3793
Route Distinguishers 3793
Configuring a BGP Session for MPLS VPNs (CLI Procedure) 3794
Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure) 3795
Configuring Routing Options for MPLS VPNs (CLI Procedure) 3795

Copyright © 2016, Juniper Networks, Inc.

Configuring a Routing Instance for MPLS VPNs (CLI Procedure) 3796

Chapter 160 **Configuring MPLS Layer 2 VPNs 3799**

Understanding MPLS Layer 2 VPNs 3799
MPLS Layer 2 VPN Configuration Overview 3799
Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure) 3801
Configuring Interfaces for Layer 2 VPNs (CLI Procedure) 3802
Verifying an MPLS Layer 2 VPN Configuration 3803

Chapter 161 **Configuring MPLS Layer 2 Circuit VPNs 3805**

Understanding MPLS Layer 2 Circuits 3805
MPLS Layer 2 Circuit Configuration Overview 3806
Configuring an MPLS Layer 2 Circuit (CLI Procedure) 3807
Verifying an MPLS Layer 2 Circuit Configuration 3807
Configuring an IGP and the LDP Signaling Protocol (CLI Procedure) 3808

Chapter 162 **Configuring MPLS Layer 3 VPNs 3809**

Understanding MPLS Layer 3 VPNs 3809
MPLS Layer 3 VPN Configuration Overview 3810
Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure) 3811
Verifying an MPLS Layer 3 VPN Configuration 3812

Part 56 Configuring CLNS VPNs

Chapter 163 **Introduction to CLNS 3815**

CLNS Overview 3815
CLNS Configuration Overview 3815

Chapter 164 **Configuring ES-IS for CLNS 3819**

Understanding ES-IS for CLNS 3819
Example: Configuring ES-IS for CLNS 3820

Chapter 165 **Configuring IS-IS for CLNS 3823**

Understanding IS-IS for CLNS 3823
Example: Configuring IS-IS for CLNS 3823

Chapter 166 **Configuring Static Routes for CLNS 3827**

Understanding Static Routes for CLNS 3827
Example: Configuring Static Routes for CLNS 3827

Chapter 167 **Configuring BGP for CLNS 3831**

Understanding BGP for CLNS VPNs 3831
Example: Configuring BGP for CLNS VPNs 3832
Example: Configuring a VPN Routing Instance for CLNS 3834
Verifying a CLNS VPN Configuration 3835

Part 57 Configuring VPLS

Chapter 168 **Introduction to VPLS 3841**

VPLS Overview 3841
Sample VPLS Topology 3842
VPLS on PE Routers 3843

Copyright © 2016, Juniper Networks, Inc.

Using an Ethernet Switch as the VPLS CE Device . . . 3845
VPLS Exceptions on SRX Series Devices . . . 3845
VPLS Configuration Overview . . . 3846

Chapter 169 Configuring Interfaces . . . 3849

Understanding VPLS Interfaces . . . 3849
Interface Name . . . 3849
Encapsulation Type . . . 3849
Flexible VLAN Tagging . . . 3850
VLAN Rewrite . . . 3850
Example: Configuring Routing Interfaces on the VPLS PE Router . . . 3851
Example: Configuring the Interface to the VPLS CE Device . . . 3852
VPLS Filters and Policers Overview . . . 3853
Example: Configuring VPLS Filters . . . 3853
Example: Configuring VPLS Policers . . . 3855

Chapter 170 Configuring Routing Instances . . . 3859

Understanding VPLS Routing Instances . . . 3859
BGP Signaling . . . 3860
VPLS Routing Table . . . 3860
Trace Options . . . 3861
Example: Configuring the VPLS Routing Instance . . . 3862
Example: Configuring Automatic Site Identifiers for VPLS . . . 3864

Chapter 171 Configuring Routing and Signaling Protocols . . . 3867

Example: Configuring OSPF on the VPLS PE Router . . . 3867
Example: Configuring RSVP on the VPLS PE Router . . . 3868
Example: Configuring MPLS on the VPLS PE Router . . . 3869
Example: Configuring LDP on the VPLS PE Router . . . 3871
Example: Configuring VPLS over GRE with IPsec VPNs . . . 3872
Example: Configuring VPLS with BGP Signaling . . . 3889
Example: Configuring BGP on the VPLS PE Router . . . 3902
Example: Configuring Routing Options on the VPLS PE Router . . . 3903

Chapter 172 Configuring Encapsulation . . . 3905

Understanding VPLS VLAN Encapsulation . . . 3905
Understanding VPLS VLAN Encapsulation on a Logical Interface . . . 3906
Example: Configuring VPLS VLAN Encapsulation . . . 3906
Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces . . . 3909
Example: Configuring Extended VLAN VPLS Encapsulation . . . 3910

Part 58 Configuration Statements and Operational Commands

Chapter 173 Configuration Statements . . . 3917

Accounting-Options Configuration Statement Hierarchy . . . 3917
Firewall Configuration Statement Hierarchy . . . 3919
Forwarding-Options Configuration Statement Hierarchy . . . 3929
Interfaces Configuration Statement Hierarchy . . . 3941
Policy-Options Configuration Statement Hierarchy . . . 3957
[edit security address-book] Hierarchy Level . . . 3958

Copyright © 2016, Juniper Networks, Inc.

[edit security forwarding-options] Hierarchy Level . . . 3959
[edit security ike] Hierarchy Level . . . 3959
[edit security ipsec] Hierarchy Level . . . 3961
[edit security policies] Hierarchy Level . . . 3963
condition (Policy Options) . . . 3967
family (Security Forwarding Options) . . . 3968
flow-server (Forwarding Options) . . . 3969
forwarding-options (Security) . . . 3971
fragment . . . 3972
hash-key (Forwarding Options) . . . 3973
iso (Security Forwarding Options) . . . 3974
mpls (Security Forwarding Options) . . . 3975
multicast-scope . . . 3976
policer (Firewall) . . . 3977
simple-filter (Firewall) . . . 3978
template (Flow Monitoring) . . . 3980
traceoptions (Security Flow) . . . 3981
version9 (Flow Server) . . . 3984

Chapter 174 **Operational Commands . . . 3985**

show bgp neighbor (View) . . . 3986
show interfaces flow-statistics . . . 3990
show interfaces statistics (View) . . . 3995
show security flow status . . . 3996
show security ipsec security-associations . . . 3999
show security ipsec statistics . . . 4011

Guide 12 Multicast Feature Guide for Security Devices

Part 59 Overview

Chapter 175 **Introduction to Multicast . . . 4019**

Multicast Overview . . . 4019
Comparing Multicast to Unicast . . . 4020
IP Multicast Uses . . . 4021
IP Multicast Terminology . . . 4022
Reverse-Path Forwarding for Loop Prevention . . . 4023
Shortest-Path Tree for Loop Prevention . . . 4023
Administrative Scoping for Loop Prevention . . . 4024
Multicast Leaf and Branch Terminology . . . 4024
IP Multicast Addressing . . . 4024
Multicast Addresses . . . 4025
Layer 2 Frames and IPv4 Multicast Addresses . . . 4025
Multicast Interface Lists . . . 4027
Multicast Routing Protocols . . . 4028
T Series Router Multicast Performance . . . 4031
Understanding Layer 3 Multicast Functionality on the SRX5K-MPC . . . 4031
Supported IP Multicast Protocol Standards . . . 4032

Copyright © 2016, Juniper Networks, Inc.

Part 60 Managing Group Membership

Chapter 176 Configuring IGMP 4037

Understanding Group Membership Protocols 4037
Understanding IGMP 4038
Configuring IGMP 4040
Enabling IGMP 4041
Modifying the IGMP Host-Query Message Interval 4042
Modifying the IGMP Query Response Interval 4043
Specifying Immediate-Leave Host Removal for IGMP 4044
Filtering Unwanted IGMP Reports at the IGMP Interface Level 4045
Accepting IGMP Messages from Remote Subnetworks 4046
Modifying the IGMP Last-Member Query Interval 4047
Modifying the IGMP Robustness Variable 4047
Limiting the Maximum IGMP Message Rate 4049
Changing the IGMP Version 4049
Enabling IGMP Static Group Membership 4050
Recording IGMP Join and Leave Events 4056
Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces 4057
Tracing IGMP Protocol Traffic 4059
Disabling IGMP 4060
IGMP and Nonstop Active Routing 4060

Chapter 177 Examples: Configuring MLD 4063

Understanding MLD 4063
Configuring MLD 4066
Enabling MLD 4067
Modifying the MLD Version 4068
Modifying the MLD Host-Query Message Interval 4068
Modifying the MLD Query Response Interval 4069
Modifying the MLD Last-Member Query Interval 4070
Specifying Immediate-Leave Host Removal for MLD 4071
Filtering Unwanted MLD Reports at the MLD Interface Level 4072
Example: Modifying the MLD Robustness Variable 4073
Limiting the Maximum MLD Message Rate 4075
Enabling MLD Static Group Membership 4075
Example: Recording MLD Join and Leave Events 4082
Configuring the Number of MLD Multicast Group Joins on Logical Interfaces . . 4084
Disabling MLD 4085

Part 61 Configuring Protocol Independent Multicast

Chapter 178 Understanding PIM 4089

PIM Overview 4089
Basic PIM Network Components 4091

Copyright © 2016, Juniper Networks, Inc.

Chapter 179 Configuring PIM Basics 4093

Configuring Basic PIM Settings 4093
PIM Configuration Statements 4093
Changing the PIM Version 4096
Modifying the PIM Hello Interval 4096
Preserving Multicast Performance by Disabling Response to the ping Utility 4097
PIM on Aggregated Interfaces 4098
Configuring PIM Trace Options 4098
Disabling PIM 4100
Disabling the PIM Protocol 4101
Disabling PIM on an Interface 4101
Disabling PIM for a Family 4102
Disabling PIM for a Rendezvous Point 4102
Configuring a Designated Router for PIM 4103
Configuring Interface Priority for PIM Designated Router Selection 4103
Configuring PIM Designated Router Election on Point-to-Point Links 4104

Chapter 180 Routing Content to Densely Clustered Receivers with PIM Dense Mode 4107

Configuring PIM Dense Mode 4107
Understanding PIM Dense Mode 4107
Configuring PIM Dense Mode Properties 4109
Configuring PIM Sparse-Dense Mode 4110
Understanding PIM Sparse-Dense Mode 4110
Mixing PIM Sparse and Dense Modes 4110
Configuring PIM Sparse-Dense Mode Properties 4111

Chapter 181 Routing Content to Larger, Sparser Groups with PIM Sparse Mode 4113

Examples: Configuring PIM Sparse Mode 4113
Understanding PIM Sparse Mode 4113
Rendezvous Point 4115
RP Mapping Options 4116
Designated Router 4116
Tunnel Services PICs and Multicast 4116
Enabling PIM Sparse Mode 4118
Configuring PIM Join Load Balancing 4119
Modifying the Join State Timeout 4122
Example: Enabling Join Suppression 4122
Example: Configuring PIM Sparse Mode over an IPsec VPN 4127
Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces . . . 4131
Configuring Static RP 4135
Understanding Static RP 4135
Configuring Local PIM RPs 4136
Configuring the Static PIM RP Address on the Non-RP Routing Device . . . 4137

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Anycast RP 4139
Understanding RP Mapping with Anycast RP 4139
Example: Configuring Multiple RPs in a Domain with Anycast RP 4140
Example: Configuring PIM Anycast With or Without MSDP 4142
Configuring a PIM Anycast RP Router Using Only PIM 4146
Configuring PIM Bootstrap Router 4147
Understanding the PIM Bootstrap Router 4147
Configuring PIM Bootstrap Properties for IPv4 4147
Configuring PIM Bootstrap Properties for IPv4 or IPv6 4149
Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain 4150
Example: Configuring PIM BSR Filters 4151
Configuring PIM Auto-RP 4151
Understanding PIM Auto-RP 4151
Configuring PIM Auto-RP 4152
Configuring Embedded RP 4156
Understanding Embedded RP for IPv6 Multicast 4156
Configuring PIM Embedded RP for IPv6 4158
Configuring PIM Filtering 4159
Understanding Multicast Message Filters 4159
Filtering MAC Addresses 4160
Filtering RP and DR Register Messages 4160
Filtering MSDP SA Messages 4161
Configuring Interface-Level PIM Neighbor Policies 4161
Filtering Outgoing PIM Join Messages 4162
Filtering Incoming PIM Join Messages 4163
Configuring Register Message Filters on a PIM RP and DR 4164
Examples: Configuring PIM RPT and SPT Cutover 4166
Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees 4167
Building an RPT Between the RP and Receivers 4168
PIM Sparse Mode Source Registration 4169
Multicast Shortest-Path Tree 4171
SPT Cutover 4172
SPT Cutover Control 4175
Example: Configuring the PIM Assert Timeout 4175
Example: Configuring the PIM SPT Threshold Policy 4177

Chapter 182 **Receiving Content Directly from the Source with SSM 4183**

Example: Configuring Source-Specific Multicast 4183
Understanding PIM Source-Specific Mode 4183
PIM SSM 4184
Source-Specific Multicast Groups Overview 4186
Example: Configuring Source-Specific Multicast Groups with Any-Source Override 4187
Example: Configuring an SSM-Only Domain 4190
Example: Configuring PIM SSM on a Network 4191

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring SSM Mapping 4192
Example: Configuring SSM Maps for Different Groups to Different Sources . . . 4195
Multiple SSM Maps and Groups for Interfaces 4195
Example: Configuring Multiple SSM Maps Per Interface 4195

Chapter 183 Minimizing Routing State Information with Bidirectional PIM 4199

Example: Configuring Bidirectional PIM 4199
Understanding Bidirectional PIM 4199
Designated Forwarder Election 4201
Bidirectional PIM Modes 4202
Bidirectional Rendezvous Points 4202
PIM Bootstrap and Auto-RP Support 4203
IGMP and MLD Support 4203
Bidirectional PIM and Graceful Restart 4203
Junos OS Enhancements to Bidirectional PIM 4204
Limitations of Bidirectional PIM 4204
Example: Configuring Bidirectional PIM 4205

Chapter 184 Rapidly Detecting Communication Failures with PIM and the BFD Protocol 4219

Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol . . 4219
Understanding Bidirectional Forwarding Detection Authentication for PIM 4219
BFD Authentication Algorithms 4220
Security Authentication Keychains 4220
Strict Versus Loose Authentication 4221
Configuring BFD for PIM 4221
Configuring BFD Authentication for PIM 4222
Configuring BFD Authentication Parameters 4223
Viewing Authentication Information for BFD Sessions 4224

Chapter 185 Configuring PIM Options 4227

Example: Configuring Nonstop Active Routing for PIM 4227
Understanding Nonstop Active Routing for PIM 4227
Example: Configuring Nonstop Active Routing with PIM 4228
Configuring PIM Sparse Mode Graceful Restart 4238
Configuring PIM-to-IGMP and PIM-to-MLD Message Translation 4240
Understanding PIM-to-IGMP and PIM-to-MLD Message Translation 4240
Configuring PIM-to-IGMP Message Translation 4241
Configuring PIM-to-MLD Message Translation 4242

Part 62 Configuring Multicast Routing Protocols

Chapter 186 Improving Multicast Reliability with PGM 4247

Configuring PGM 4247
Understanding Pragmatic General Multicast 4247
PGM Architecture and PGM Routers 4248
PGM-Enabled Source 4249
PGM-Enabled Receivers 4249
PGM-Enabled Routers 4250

Copyright © 2016, Juniper Networks, Inc.

PGM Configuration Guidelines 4251

Chapter 187 **Connecting Routing Domains Using MSDP 4253**

Examples: Configuring MSDP 4253
Configuring MSDP 4253
Example: Configuring MSDP in a Routing Instance 4254
Configuring the Interface to Accept Traffic from a Remote Source 4262
Example: Configuring MSDP with Active Source Limits and Mesh Groups . . 4262
Tracing MSDP Protocol Traffic 4268
Disabling MSDP 4269

Chapter 188 **Handling Session Announcements with SAP 4271**

Configuring the Session Announcement Protocol 4271

Chapter 189 **Facilitating Multicast Delivery Across Unicast-Only Networks with AMT 4273**

Example: Configuring Automatic IP Multicast Without Explicit Tunnels 4273
Understanding AMT 4273
AMT Applications 4274
AMT Operation 4276
Configuring the AMT Protocol 4277
Configuring Default IGMP Parameters for AMT Interfaces 4279
Example: Configuring the AMT Protocol 4281

Chapter 190 **Routing Content to Densely Clustered Receivers with DVMRP 4287**

Examples: Configuring DVMRP 4287
Understanding DVMRP 4287
Configuring DVMRP 4288
Example: Configuring DVMRP 4288
Example: Configuring DVMRP to Announce Unicast Routes 4292
Tracing DVMRP Protocol Traffic 4295

Part 63 Configuring Multicast VPNs

Chapter 191 **Configuring PIM Join Load Balancing 4299**

PIM Join Load Balancing on Multipath MVPN Routes Overview 4299

Chapter 192 **Configuring Next-Generation Multicast VPNs 4303**

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN 4303

Part 64 Configuring General Multicast Routing Options

Chapter 193 **Preventing Routing Loops with Reverse Path Forwarding 4315**

Examples: Configuring Reverse Path Forwarding 4315
Understanding Multicast Reverse Path Forwarding 4315
RPF Table 4316
Multicast RPF Configuration Guidelines 4317
Example: Configuring a Dedicated PIM RPF Routing Table 4317
Example: Configuring RPF Policies 4321
Example: Configuring PIM RPF Selection 4323

Copyright © 2016, Juniper Networks, Inc.

Chapter 194 Enabling Multicast Between Layer 2 and Layer 3 Devices Using Snooping 4329

Example: Configuring IGMP Snooping 4329
Understanding Multicast Snooping 4329
Understanding IGMP Snooping 4330
IGMP Snooping Interfaces and Forwarding 4331
IGMP Snooping and Proxies 4331
Multicast-Router Interfaces and IGMP Snooping Proxy Mode 4332
Host-Side Interfaces and IGMP Snooping Proxy Mode 4333
IGMP Snooping and Bridge Domains 4333
Configuring IGMP Snooping 4333
Configuring VLAN-Specific IGMP Snooping Parameters 4334
Example: Configuring IGMP Snooping 4335
Configuring IGMP Snooping Trace Operations 4341
Example: Configuring Multicast Snooping 4343
Understanding Multicast Snooping 4343
Understanding Multicast Snooping and VPLS Root Protection 4344
Configuring Multicast Snooping 4344
Example: Configuring Multicast Snooping 4345
Enabling Bulk Updates for Multicast Snooping 4350
Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces 4351

Chapter 195 Configuring Multicast Routing Options 4353

Examples: Configuring Bandwidth Management 4353
Understanding Bandwidth Management for Multicast 4353
Bandwidth Management and PIM Graceful Restart 4354
Bandwidth Management and Source Redundancy 4354
Logical Systems and Bandwidth Oversubscription 4354
Example: Defining Interface Bandwidth Maximums 4355
Example: Configuring Multicast with Subscriber VLANs 4358
Configuring Multicast Routing over IP Demux Interfaces 4371
Classifying Packets by Egress Interface 4372
Examples: Configuring the Multicast Forwarding Cache 4374
Understanding the Multicast Forwarding Cache 4374
Example: Configuring the Multicast Forwarding Cache 4374
Example: Configuring a Multicast Flow Map 4377
Example: Configuring Ingress PE Redundancy 4381
Understanding Ingress PE Redundancy 4381
Example: Configuring Ingress PE Redundancy 4381

Part 65 Configuration Statements and Operational Commands

Chapter 196 Configuration Statements 4389

accept-remote-source 4396
accounting (Protocols IGMP) 4396
accounting (Protocols IGMP AMT Interface) 4397
accounting (Protocols IGMP Interface) 4397
accounting (Protocols MLD) 4398

Copyright © 2016, Juniper Networks, Inc.

accounting (Protocols MLD Interface) . . . 4398
active-source-limit . . . 4399
address (Anycast RPs) . . . 4400
address (Bidirectional Rendezvous Points) . . . 4401
address (Local RPs) . . . 4402
address (Static RPs) . . . 4403
algorithm . . . 4404
amt (IGMP) . . . 4405
amt (Protocols) . . . 4406
anycast-pim . . . 4407
anycast-prefix . . . 4408
asm-override-ssm . . . 4409
assert-timeout . . . 4410
authentication . . . 4411
authentication-key . . . 4412
auto-rp . . . 4413
backoff-period . . . 4414
backup-pe-group . . . 4415
backups . . . 4416
bandwidth . . . 4417
bfd-liveness-detection . . . 4418
bidirectional (Interface) . . . 4419
bidirectional (RP) . . . 4420
bootstrap . . . 4421
bootstrap-export . . . 4422
bootstrap-import . . . 4423
bootstrap-priority . . . 4424
data-encapsulation . . . 4425
default-peer . . . 4426
defaults . . . 4427
dense-groups . . . 4428
detection-time (BFD for PIM) . . . 4429
df-election . . . 4430
disable (PIM) . . . 4431
disable (PIM Graceful Restart) . . . 4432
disable (Protocols DVMRP) . . . 4432
disable (Protocols IGMP) . . . 4433
disable (Protocols MLD) . . . 4433
disable (Protocols MSDP) . . . 4434
disable (Protocols SAP) . . . 4435
dr-election-on-p2p . . . 4435
dr-register-policy . . . 4436
dvmrp . . . 4437
embedded-rp . . . 4438
exclude (Protocols IGMP) . . . 4438
exclude (Protocols MLD) . . . 4439
export (Protocols DVMRP) . . . 4439
export (Protocols MSDP) . . . 4440
export (Protocols PIM) . . . 4441

Copyright © 2016, Juniper Networks, Inc.

export (Bootstrap) . . . 4442
family (Bootstrap) . . . 4443
family (Local RP) . . . 4444
family (Protocols AMT Relay) . . . 4445
family (Protocols PIM) . . . 4446
flood-groups . . . 4447
flow-map . . . 4448
forwarding-cache (Bridge Domains) . . . 4449
forwarding-cache (Flow Maps) . . . 4449
forwarding-cache (Multicast) . . . 4450
graceful-restart (Multicast Snooping) . . . 4451
graceful-restart (Protocols PIM) . . . 4452
group (Bridge Domains) . . . 4453
group (Protocols IGMP) . . . 4454
group (Protocols MLD) . . . 4455
group (Protocols MSDP) . . . 4456
group (RPF Selection) . . . 4457
group-count . . . 4458
group-count (Protocols MLD) . . . 4458
group-increment (Protocols IGMP) . . . 4459
group-increment (Protocols MLD) . . . 4459
group-limit (Protocols IGMP) . . . 4460
group-limit (IGMP and MLD Snooping) . . . 4461
group-limit (Protocols MLD) . . . 4462
group-policy (Protocols IGMP) . . . 4462
group-policy (Protocols IGMP AMT Interface) . . . 4463
group-policy (Protocols MLD) . . . 4463
group-ranges . . . 4464
hello-interval . . . 4465
hold-time (Protocols DVMRP) . . . 4465
hold-time (Protocols PIM) . . . 4466
host-only-interface . . . 4467
igmp . . . 4468
igmp-snooping . . . 4470
ignore-stp-topology-change . . . 4471
immediate-leave (Bridge Domains) . . . 4472
immediate-leave (Protocols IGMP) . . . 4474
immediate-leave (Protocols MLD) . . . 4475
import (Protocols DVMRP) . . . 4476
import (Protocols MSDP) . . . 4477
import (Protocols PIM) . . . 4478
import (Protocols PIM Bootstrap) . . . 4479
infinity . . . 4480
inet . . . 4481
interface (Bridge Domains) . . . 4482
interface (Protocols DVMRP) . . . 4483
interface (Protocols IGMP) . . . 4484
interface (Protocols MLD) . . . 4485
interface (Protocols PIM) . . . 4486

Copyright © 2016, Juniper Networks, Inc.

interface (Routing Options) 4487
interface (Scoping) 4488
join-load-balance 4489
join-prune-timeout 4490
key-chain 4490
listen 4491
local 4492
local-address (Protocols AMT) 4493
local-address (Protocols MSDP) 4494
local-address (Protocols PIM) 4495
local-address (Routing Options) 4496
loose-check 4497
mapping-agent-election 4498
maximum (MSDP Active Source Messages) 4499
maximum-bandwidth 4500
maximum-rps 4501
maximum-transmit-rate (Protocols IGMP) 4501
metric (Protocols DVMRP) 4502
minimum-interval (PIM BFD Liveness Detection) 4502
minimum-interval (PIM BFD Transmit Interval) 4503
minimum-receive-interval 4504
mld 4505
mode (Protocols DVMRP) 4506
mode (Protocols MSDP) 4507
mode (Protocols PIM) 4508
msdp 4509
multicast (Dynamic Profiles Routing Options) 4511
multicast-router-interface 4513
multicast-snooping-options 4514
multichassis-lag-replicate-state 4515
multiplier 4516
neighbor-policy 4516
nexthop-hold-time 4517
next-hop (PIM RPF Selection) 4517
no-adaptation (PIM BFD Liveness Detection) 4518
no-bidirectional-mode 4519
no-multicast-echo 4520
no-qos-adjust 4521
offer-period 4522
oif-map (IGMP Interface) 4523
oif-map (MLD Interface) 4523
override (PIM Static RP) 4524
override-interval 4525
passive (IGMP) 4526
passive (MLD) 4527
peer 4528
pgm 4529
pim 4530
pim-to-igmp-proxy 4534

Copyright © 2016, Juniper Networks, Inc.

pim-to-mld-proxy 4535
policy (Flow Maps) 4536
policy (SSM Maps) 4536
prefix-list (PIM RPF Selection) 4537
priority (Bootstrap) 4538
priority (PIM Interfaces) 4539
priority (PIM RPs) 4540
propagation-delay 4541
proxy 4542
query-interval (Bridge Domains) 4543
query-interval (Protocols IGMP) 4544
query-interval (Protocols IGMP AMT) 4545
query-interval (Protocols MLD) 4545
query-last-member-interval (Bridge Domains) 4546
query-last-member-interval (Protocols IGMP) 4547
query-last-member-interval (Protocols MLD) 4547
query-response-interval (Bridge Domains) 4548
query-response-interval (Protocols IGMP) 4549
query-response-interval (Protocols IGMP AMT) 4550
query-response-interval (Protocols MLD) 4551
redundant-sources 4552
relay (AMT Protocol) 4553
relay (IGMP) 4554
reset-tracking-bit 4555
restart-duration 4556
reverse-oif-mapping 4557
rib-group (Protocol DVMRP) 4557
rib-group (Protocols MSDP) 4558
rib-group (Protocols PIM) 4559
robust-count (Bridge Domains) 4560
robust-count (Protocols IGMP) 4561
robust-count (Protocols IGMP AMT) 4561
robust-count (Protocols MLD) 4562
robustness-count 4563
rp 4564
rp-register-policy 4566
rp-set 4567
rpf-check-policy 4568
rpf-selection 4569
sap 4570
scope 4571
scope-policy 4572
secret-key-timeout 4573
source (Bridge Domains) 4573
source (PIM RPF Selection) 4574
source (Protocols IGMP) 4575
source (Protocols MLD) 4575
source (Protocols MSDP) 4576
source (Routing Instances) 4577

Copyright © 2016, Juniper Networks, Inc.

source-address . . . 4577
source-count (Protocols IGMP) . . . 4578
source-count (Protocols MLD) . . . 4578
source-increment (Protocols IGMP) . . . 4579
source-increment (Protocols MLD) . . . 4579
spt-threshold . . . 4580
ssm-groups . . . 4581
ssm-map (Protocols IGMP AMT) . . . 4582
ssm-map (Protocols MLD) . . . 4582
ssm-map (Multicast Routing Options) . . . 4583
ssm-map (Protocols IGMP) . . . 4583
ssm-map-policy (IGMP) . . . 4584
ssm-map-policy (MLD) . . . 4584
static (Bridge Domains) . . . 4585
static (Protocols MLD) . . . 4586
static (Protocols IGMP) . . . 4587
static (Protocols PIM) . . . 4588
subscriber-leave-timer . . . 4589
threshold (Bridge Domains) . . . 4590
threshold (Protocols MSDP) . . . 4591
threshold (PIM BFD Detection Time) . . . 4592
threshold (PIM BFD Transmit Interval) . . . 4593
timeout (Flow Maps) . . . 4594
timeout (Multicast) . . . 4595
traceoptions (Multicast Snooping Options) . . . 4596
traceoptions (Protocols AMT) . . . 4598
traceoptions (Protocols DVMRP) . . . 4601
traceoptions (Protocols IGMP) . . . 4604
traceoptions (Protocols IGMP Snooping) . . . 4607
traceoptions (Protocols MLD) . . . 4609
traceoptions (Protocols MSDP) . . . 4612
traceoptions (Protocols PGM) . . . 4615
traceoptions (Protocols PIM) . . . 4617
transmit-interval (PIM BFD Liveness Detection) . . . 4620
tunnel-devices (Routing Instances) . . . 4621
tunnel-limit (Protocols AMT) . . . 4622
upstream-interface . . . 4623
version (BFD) . . . 4624
version (PIM) . . . 4625
version (Protocols IGMP) . . . 4626
version (Protocols IGMP AMT) . . . 4626
version (Protocols MLD) . . . 4627
vlan (Bridge Domains) . . . 4628
vpn-group-address . . . 4629
wildcard-source (PIM RPF Selection) . . . 4629

Chapter 197 Operational Commands . . . 4631

clear amt statistics . . . 4633
clear amt tunnel . . . 4634

Copyright © 2016, Juniper Networks, Inc.

clear igmp membership 4635
clear igmp snooping membership 4638
clear igmp snooping statistics 4639
clear igmp statistics 4640
clear mld membership 4642
clear mld statistics 4643
clear multicast snooping statistics 4644
clear pgm negative-acknowledgments 4645
clear pgm source-path-messages 4646
clear pgm statistics 4647
clear pim join 4648
clear pim join-distribution 4650
clear pim register 4652
clear pim statistics 4654
request pim multicast-tunnel rebalance 4657
show amt statistics 4658
show amt summary 4661
show amt tunnel 4663
show dvmrp interfaces 4666
show dvmrp neighbors 4668
show dvmrp prefix 4670
show dvmrp prunes 4672
show igmp group 4674
show igmp interface 4678
show igmp snooping interface 4682
show igmp snooping membership 4685
show igmp snooping statistics 4689
show igmp statistics 4692
show mld group 4695
show mld interface 4699
show mld statistics 4702
show msdp 4705
show msdp source 4707
show msdp source-active 4709
show msdp statistics 4711
show multicast backup-pe-groups 4715
show multicast flow-map 4717
show multicast interface 4719
show multicast pim-to-igmp-proxy 4721
show multicast pim-to-mld-proxy 4723
show multicast route 4725
show multicast rpf 4731
show multicast scope 4735
show multicast sessions 4737
show multicast snooping route 4740
show multicast snooping statistics 4743
show multicast usage 4746
show pgm negative-acknowledgments 4749
show pgm source-path-messages 4751

Copyright © 2016, Juniper Networks, Inc.

show pgm statistics 4752
show pim bidirectional df-election 4755
show pim bidirectional df-election interface 4758
show pim bootstrap 4761
show pim interfaces 4763
show pim join 4766
show pim neighbors 4787
show pim rps 4791
show pim source 4798
show pim statistics 4801
show policy 4811
show route 4813
show route table 4819
show sap listen 4834

Guide 13 Network Address Translation Feature Guide for Security Devices

Chapter 198 Overview 4839

Introduction to NAT 4839
Understanding NAT Rule Sets and Rules 4840
NAT Rule Sets 4840
NAT Rules 4841
Rule Processing 4841
NAT Rule Capacity 4842

Chapter 199 Configuring General NAT Options 4845

Configuring NAT Using the NAT Wizard 4845
Example: Configuring NAT for Multiple ISPs 4845
Configuring Proxy ARP (CLI Procedure) 4857
Verifying NAT Configuration 4858
Monitoring Incoming Table Information 4859
Monitoring Interface NAT Port Information 4860

Chapter 200 Configuring Source NAT 4863

Understanding Source NAT 4863
Source NAT Configuration Overview 4864
Example: Configuring Source NAT for Egress Interface Translation 4865
Example: Configuring Source NAT for Single Address Translation 4869
Example: Configuring Source and Destination NAT Translations 4873
Understanding Central Point Architecture Enhancements for NAT 4880
Understanding Reverse NAT Enhancements for Central Point Architecture . . 4880
Understanding Source NAT Rules 4881
Example: Configuring Source NAT with Multiple Rules 4882
Disabling Port Randomization for Source NAT (CLI Procedure) 4888
Monitoring Source NAT Information 4889

Chapter 201 Configuring Source NAT Pools 4895

Understanding Source NAT Pools 4895
Understanding Source NAT Pool Capacities 4897

Copyright © 2016, Juniper Networks, Inc.

Understanding Persistent Addresses 4898
Example: Configuring Capacity for Source NAT Pools with PAT 4898
Understanding Source NAT Pools with Address Pooling 4900
Understanding Source NAT Pools with Address Shifting 4901
Example: Configuring Source NAT with Address Shifting 4901
Understanding Source NAT Pools with PAT 4906
Example: Configuring Source NAT for Multiple Addresses with PAT 4907
Understanding Source NAT Pools Without PAT 4912
Example: Configuring a Single IP Address in a Source NAT Pool Without PAT . . 4913
Example: Configuring Source NAT for Multiple Addresses Without PAT 4917
Understanding Source NAT Pools with Shared Address 4921

Chapter 202 Configuring Destination NAT 4923

Understanding Destination NAT 4923
Understanding Destination NAT Address Pools 4924
Understanding Destination NAT Rules 4925
Destination NAT Configuration Overview 4926
Example: Improving Security By Configuring Destination NAT for Single Address Translation 4926
Example: Configuring Destination NAT for IP Address and Port Translation . . 4934
Example: Configuring Destination NAT for Subnet Translation 4939
Monitoring Destination NAT Information 4944

Chapter 203 Configuring Static NAT 4947

Understanding Static NAT 4947
Understanding Static NAT Rules 4948
Static NAT Configuration Overview 4949
Example: Configuring Static NAT for Single Address Translation 4949
Example: Configuring Static NAT for Subnet Translation 4953
Example: Configuring Static NAT for Port Mapping 4958
Monitoring Static NAT Information 4965

Chapter 204 Configuring Persistent NAT and NAT64 4969

Understanding Persistent NAT and NAT64 4969
Understanding Session Traversal Utilities for NAT (STUN) Protocol 4971
Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation . . . 4972
Persistent NAT and NAT64 Configuration Overview 4974
Example: Configuring Address Persistent NAT64 Pools 4975
Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT 4977
Example: Configuring Persistent NAT with Source NAT Address Pool (CLI) . . . 4983
Example: Configuring Address-Dependent Filtering for IPv6 Clients 4984
Example: Configuring Endpoint-Independent Filtering for IPv6 Clients 4987
Example: Setting Maximum Persistent NAT Bindings 4990

Chapter 205 Configuring NAT Hairpinning 4993

Persistent NAT Hairpinning Overview 4993
Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting 4994

Copyright © 2016, Juniper Networks, Inc.

Chapter 206 Configuring NAT for Multicast Flows 5001
Understanding NAT for Multicast Flows 5001
Example: Configuring NAT for Multicast Flows 5002
Chapter 207 Configuring IPv6 NAT 5011
IPv6 NAT Overview 5011
Source NAT Translations Supported by IPv6 NAT 5011
Destination NAT Mappings Supported by IPv6 NAT 5012
Static NAT Mappings Supported by IPv6 NAT 5012
IPv6 NAT PT Overview 5013
IPv6 NAT-PT Communication Overview 5014
Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping 5015
Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping 5018
Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping 5022
Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping 5026
Chapter 208 Configuring IPv6 Dual-Stack 5031
Understanding IPv6 Dual-Stack Lite 5031
Example: Configuring IPv6 Dual-Stack Lite 5034
Chapter 209 Configuration Statements 5037
Security Configuration Statement Hierarchy 5039
[edit security nat] Hierarchy Level 5041
address (Security ARP Proxy) 5045
address (Security Destination NAT) 5045
address (Security NDP Proxy) 5046
address-mapping 5046
address-persistent (Security Source NAT) 5047
address-persistent (Security Source NAT Pool) 5047
address-pooling (Security Source NAT) 5048
address-shared (Security Source NAT) 5049
application (Security Destination NAT) 5049
application (Security Source NAT) 5050
application-services (Security Forwarding Process) 5051
clear-threshold 5052
description (Security NAT Pool) 5053
description (Security NAT Rule) 5054
description (Security NAT Rule Set) 5055
destination (Security Destination NAT) 5056
destination-address (Security Destination NAT) 5057
destination-address (Security Source NAT) 5058
destination-address (Security Static NAT) 5058
destination-address-name (Security Destination NAT) 5059
destination-address-name (Security Source NAT) 5059
destination-address-name (Security Static NAT) 5060
destination-nat 5060

Copyright © 2016, Juniper Networks, Inc.

destination-port (Security Destination NAT) 5061
destination-port (Security Source NAT) 5061
destination-port (Security Static NAT) 5062
from (Security NAT) 5062
host-address-base 5063
inactivity-timeout (Security Persistent NAT) 5063
inet (Security Static NAT) 5064
interface (Security NAT ARP Proxy) 5065
interface (Security NAT NDP Proxy) 5065
interface (Security Source NAT) 5066
interface (Security Source NAT Rule Set) 5066
mapped-port (Security Static NAT) 5067
match (Security Destination NAT) 5068
match (Security Source NAT) 5069
match (Security Static NAT) 5070
max-session-number 5070
overflow-pool 5071
nptv6-prefix 5072
nptv6-prefix-name 5072
permit (Security Persistent NAT) 5073
persistent-nat 5074
pool (Security Destination NAT) 5075
pool (Security Source NAT) 5076
pool (Security Source NAT Rule Set) 5077
pool-default-port-range 5078
pool-default-twin-port-range 5079
pool-utilization-alarm 5080
pool-utilization-alarm (Security Source NAT Pool) 5081
port (Security Source NAT) 5082
port-overloading (Security Source NAT Interface) 5083
port-overloading-factor (Security Source NAT Interface) 5084
port-overloading-factor (Security Source NAT Pool) 5085
port-randomization 5086
port-round-robin 5086
prefix (Security Static NAT) 5087
prefix-name (Security Static NAT) 5088
protocol (Security Destination NAT) 5088
protocol (Security Source NAT) 5089
proxy-arp (Security NAT) 5089
proxy-ndp (Security NAT) 5090
raise-threshold 5090
routing-instance (Security Destination NAT) 5091
routing-instance (Security Source NAT) 5091
rule (Security Destination NAT) 5092
rule (Security Source NAT) 5093
rule (Security Static NAT) 5095
rule-session-count-alarm (Security Destination NAT Rule Set) 5096
rule-session-count-alarm (Security Source NAT Rule Set) 5097
rule-session-count-alarm (Security Static NAT Rule Set) 5098

Copyright © 2016, Juniper Networks, Inc.

rule-set (Security Destination NAT) 5099
rule-set (Security Source NAT) 5100
rule-set (Security Static NAT) 5102
source (Security Source NAT) 5104
source-address (Security Destination NAT) 5106
source-address (Security Source NAT) 5107
source-address (Security Static NAT Rule Set) 5107
source-address-name (Security Destination NAT) 5108
source-address-name (Security Source NAT) 5108
source-address-name (Security Static NAT Rule Set) 5109
source-nat 5110
source-port (Security Source NAT Rule Set) 5111
source-port (Security Static NAT Rule Set) 5111
static (Security NAT) 5112
static-nat 5113
to (Security Source NAT) 5114
then (Security Destination NAT) 5114
then (Security Source NAT) 5115
then (Security Static NAT) 5116
traceoptions (Security NAT) 5117

Chapter 210 Operational Commands 5119

clear security nat incoming-table 5120
clear security nat source persistent-nat-table 5121
clear security nat statistics destination pool 5122
clear security nat statistics destination rule 5123
clear security nat statistics source pool 5124
clear security nat statistics source rule 5125
clear security nat statistics static rule 5126
show security nat destination pool 5127
show security nat destination rule 5129
show security nat destination rule-application 5132
show security nat destination summary 5134
show security nat incoming-table 5136
show security nat interface-nat-ports 5138
show security nat resource-usage source-pool 5141
show security nat source deterministic 5144
show security nat source paired-address 5146
show security nat source persistent-nat-table 5148
show security nat source pool 5151
show security nat source port-block 5155
show security nat source rule 5157
show security nat source rule-application 5161
show security nat source summary 5163
show security nat static rule 5165

Copyright © 2016, Juniper Networks, Inc.

Guide 14 Authentication and Integrated User Firewalls Feature Guide for Security Devices

Part 66 Overview

Chapter 211 Introduction to User Authentication 5173

Understanding User Authentication for Security Devices 5173
Understanding the Three-Tiered User Firewall Features 5173

Part 67 Configuring Firewall User Authentication

Chapter 212 Understanding Firewall Authentication 5179

Firewall User Authentication Overview 5179
Obtaining Username and Role Information Through Firewall Authentication . . 5180

Chapter 213 Configuring Encrypted Files Using SSL Proxy 5183

SSL Proxy Overview 5183
- Supported Ciphers in Proxy Mode 5185
- Server Authentication 5186
- Trusted CA List 5187
- Root CA 5187
- Client Authentication 5188
- Whitelists 5188
- Dynamic Resolution of Domain Names 5188
- Session Resumption 5188
- Session Renegotiation 5188
- SSL Proxy Logs 5189
- Leveraging Dynamic Application Identification 5190
- Logical Systems Support 5190
- Limitations 5191

Configuring SSL Proxy 5191
- SSL Proxy Configuration Overview 5192
- Configuring a Root CA Certificate 5193
- Configuring a CA Profile Group 5195
- Configuring a Trusted CA Profile 5196
- Importing a Root CA Certificate into a Browser 5197
- Applying an SSL Proxy Profile to a Security Policy 5198
- Creating a Whitelist of Exempted Destinations 5199
- Configuring SSL Proxy Logging 5200
- Configuring Ciphers 5200
- Exporting Certificates to a Specified Location 5201
- Ignoring Server Authentication 5201

Enabling Debugging and Tracing for SSL Proxy 5202

Chapter 214 Configuring Pass-Through Authentication 5205

Understanding Pass-Through Authentication 5205
Example: Configuring Pass-Through Authentication 5207
Example: Configuring HTTPS Traffic to Trigger Pass-Through Authentication . . 5213

Copyright © 2016, Juniper Networks, Inc.

Chapter 215 Configuring Web Authentication 5221
Understanding Web Authentication 5221
Example: Configuring Web Authentication 5223
Example: Configuring HTTPS Traffic to Trigger Web Authentication 5230
Chapter 216 Configuring External Authentication Servers 5235
Understanding External Authentication Servers 5235
Understanding SecurID User Authentication 5236
Example: Configuring RADIUS and LDAP User Authentication 5236
Example: Configuring SecurID User Authentication 5241
Example: Deleting the SecurID Node Secret File 5244
Chapter 217 Configuring Client Groups 5247
Understanding Client Groups for Firewall Authentication 5247
Example: Configuring the Access Profile 5247
Example: Configuring Local Users for Client Groups 5248
Chapter 218 Customizing the Firewall Authentication Banner 5251
Understanding Firewall Authentication Banner Customization 5251
Example: Customizing a Firewall Authentication Banner 5251
Part 68 Configuring Infranet Authentication
Chapter 219 Configuring UAC in a Junos OS Environment 5257
Understanding UAC in a Junos OS Environment 5257
Enabling UAC in a Junos OS Environment (CLI Procedure) 5259
Chapter 220 Establishing Communications Between Devices 5261
Understanding Communications Between the Junos OS Enforcer and the IC Series UAC Appliance 5261
Understanding Communications Between Junos OS Enforcer and a Cluster of IC Series UAC Appliances 5262
Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure) 5262
Chapter 221 Configuring Policy Enforcement 5265
Understanding Junos OS Enforcer Policy Enforcement 5265
Configuring Junos OS Enforcer Failover Options (CLI Procedure) 5266
Testing Junos OS Enforcer Policy Access Decisions Using Test-Only Mode (CLI Procedure) 5267
Verifying Junos OS Enforcer Policy Enforcement 5268
Displaying IC Series UAC Appliance Authentication Table Entries from the Junos OS Enforcer 5268
Displaying IC Series UAC Appliance Resource Access Policies from the Junos OS Enforcer 5268
Chapter 222 Classifying Traffic with User Roles 5269
Understanding Unified Access Control 5269
Acquiring User Role Information from an Active Directory Authentication Server 5269

Copyright © 2016, Juniper Networks, Inc.

Chapter 223 Configuring Endpoint Security 5287
Understanding Endpoint Security Using the Infranet Agent with the Junos OS Enforcer 5287
Configuring Endpoint Security Using the Infranet Agent with the Junos OS Enforcer 5288
Chapter 224 Configuring IPsec 5289
Understanding Junos OS Enforcer Implementations Using IPsec 5289
Example: Configuring the Device as a Junos OS Enforcer Using IPsec (CLI) . . 5290
Chapter 225 Configuring Captive Portal 5299
Understanding the Captive Portal on the Junos OS Enforcer 5299
Understanding Captive Portal Configuration on the Junos OS Enforcer 5301
Understanding the Captive Portal Redirect URL Options 5301
Example: Creating a Captive Portal Policy on the Junos OS Enforcer 5302
Example: Configuring a Redirect URL for Captive Portal 5305
Part 69 Configuring Integrated User Firewall
Chapter 226 Understanding Integrated User Firewalls 5309
Overview of Integrated User Firewall 5309
Integrated User Firewall and Authentication Sources 5309
Benefits of Integrated User Firewall 5310
How the Integrated User Firewall Works 5310
Deployment Scenario for User Firewall Integration with Windows Active Directory 5311
Limitations 5311
Understanding Active Directory Authentication Tables 5312
Active Directory Authentication as an Authentication Source 5312
Active Directory Authentication Tables 5313
State Information for Active Directory Authentication Table Entries 5314
Active Directory Authentication Table Management 5315
Timeout Interval for Table Entries 5316
LDAP Functionality in Integrated User Firewall 5317
Role of LDAP in Integrated User Firewall 5318
LDAP Server Configuration and Base Distinguished Name 5318
LDAP's Authentication Method 5318
LDAP Server's Username, Password, and Server Address 5318
Caching and Calculation of User-to-Group Mappings 5319
Updating Group Information in the Authentication Entry Table 5319
LDAP Server Status and Statistics 5319
Active Directory Autodiscovery 5319
Example: Configuring Integrated User Firewall 5320
Chapter 227 Managing Event Logs 5329
Understanding How the WMIC Reads the Event Log on the Domain Controller 5329
Windows Management Instrumentation Client 5329
WMIC Reads the Event Log on the Domain Controller 5330
Specifying IP Filters to Limit IP-to-User Mapping 5330

Copyright © 2016, Juniper Networks, Inc.

Event Log Verification and Statistics 5330
Using Firewall Authentication as an Alternative to WMIC 5331
WMIC Limitations 5331
Firewall Authentication as a Backup Method for IP Address-to-User Mappings 5331
Understanding Integrated User Firewall Domain PC Probing 5332
Overview of Domain PC Probing 5332
Probing Domain PCs for User Information 5332
Probe Response 5333
Probe Configuration 5334
Probe Rate and Statistics 5334

Part 70 Configuration Statements and Operational Commands

Chapter 228 Configuration Statements 5339

Access Configuration Statement Hierarchy 5342
Security Configuration Statement Hierarchy 5351
Services Configuration Statement Hierarchy 5352
System Configuration Statement Hierarchy 5360
[edit security firewall-authentication] Hierarchy Level 5390
[edit security policies] Hierarchy Level 5391
[edit security user-identification] Hierarchy Level 5395
[edit security zones] Hierarchy Level 5395
actions (Services SSL Proxy) 5398
active-directory-access 5400
active-directory-authentication-table 5402
address (Services) 5403
admin-search 5403
application (Security Policies) 5404
application-services (Security Policies) 5405
assemble 5406
authentication-source 5406
banner (Access FTP HTTP Telnet Authentication) 5407
banner (Access Web Authentication) 5407
base-distinguished-name 5408
ca-profile (Services) 5408
captive-portal (Services UAC) 5409
captive-portal (Services UAC Policy) 5409
certificate-verification 5410
client-group 5411
client-idle-timeout (Access Profile) 5411
client-name-filter 5412
client-session-timeout (Access Profile) 5412
configuration-file 5413
count 5413
custom-ciphers 5414
default-profile 5414
distinguished-name (Access) 5415
domain-name (Access Profile) 5415

Copyright © 2016, Juniper Networks, Inc.

enable-flow-tracing (Services) . . . 5416
enable-session-cache . . . 5416
fail . . . 5417
file (Services) . . . 5417
files (Services) . . . 5418
file (System Logging) . . . 5419
firewall-authentication . . . 5422
firewall-authentication (Security) . . . 5423
firewall-authentication (Security Policies) . . . 5424
firewall-authentication (User Identification) . . . 5425
firewall-authentication-service . . . 5425
firewall-user . . . 5426
flag (Services) . . . 5426
from-zone (Security Policies) . . . 5427
ftp (Access) . . . 5429
group-profile (Access) . . . 5430
http (Access) . . . 5431
http (Services) . . . 5432
https (Services) . . . 5433
infranet-controller . . . 5434
interface (Services) . . . 5435
interval (Services) . . . 5435
ip-address (Access Profile) . . . 5436
ip-user-mapping . . . 5437
ldap-options . . . 5438
ldap-server . . . 5439
level (Services) . . . 5439
lifetime-seconds (Security IKE) . . . 5440
link (Access) . . . 5440
local-authentication-table . . . 5441
log (Services) . . . 5442
login (Access) . . . 5443
match (Services) . . . 5443
network (Access) . . . 5444
no-remote-trace (Services) . . . 5444
pass-through . . . 5445
password (Access) . . . 5446
password (Services) . . . 5446
permit (Security Policies) . . . 5447
policies . . . 5449
pool (Access) . . . 5454
port (Access LDAP) . . . 5455
port (Services) . . . 5456
preferred-ciphers . . . 5456
prefix (Access IPv6) . . . 5457
protocol-version . . . 5457
radius-options (Access) . . . 5458
radius-server (Access) . . . 5459
range (Access) . . . 5460

Copyright © 2016, Juniper Networks, Inc.

redirect-traffic . . . 5461
redirect-url . . . 5462
retry (Access LDAP) . . . 5463
retry (Access RADIUS) . . . 5463
revert-interval (Access LDAP) . . . 5464
revert-interval (Access RADIUS) . . . 5464
root-ca (Services) . . . 5465
routing-instance (Access LDAP) . . . 5465
routing-instance (Access RADIUS) . . . 5466
search . . . 5466
search-filter . . . 5467
secret (Access Profile) . . . 5467
securid-server . . . 5468
separator . . . 5469
server-certificate (Services) . . . 5469
server-certificate-subject . . . 5470
session-options (Access Profile) . . . 5470
size (Services) . . . 5471
source-address (Access LDAP) . . . 5471
source-address (Access RADIUS) . . . 5472
ssl (Services) . . . 5473
ssl-termination-profile . . . 5475
success . . . 5475
telnet (Access) . . . 5476
termination (Services) . . . 5477
test-only-mode . . . 5477
then (Security Policies) . . . 5478
timeout (Access LDAP) . . . 5480
timeout (Access RADIUS) . . . 5480
timeout (Services) . . . 5481
timeout-action . . . 5482
to-zone (Security Policies) . . . 5483
traceoptions (Access) . . . 5486
traceoptions (Active Directory Access) . . . 5488
traceoptions (Security Firewall Authentication) . . . 5490
traceoptions (Services SSL) . . . 5491
traceoptions (Services UAC) . . . 5492
trusted-ca (Services) . . . 5493
uac-policy (Application Services) . . . 5493
uac-service . . . 5494
unified-access-control (Security) . . . 5495
unified-access-control (Services) . . . 5496
user-group-mapping . . . 5497
user-identification (Services) . . . 5499
web-authentication . . . 5501
web-authentication (Access) . . . 5502
web-authentication (Interfaces) . . . 5503
web-management (System Services) . . . 5504
web-redirect . . . 5507

Copyright © 2016, Juniper Networks, Inc.

web-redirect-to-https 5508
whitelist (Services) 5508
wins-server (Access) 5509

Chapter 229 Operational Commands 5511

clear network-access requests pending 5513
clear network-access requests statistics 5514
clear network-access securid-node-secret-file 5515
clear security firewall-authentication history 5516
clear security firewall-authentication history address 5517
clear security firewall-authentication history identifier 5518
clear security firewall-authentication users 5520
clear security firewall-authentication users address 5521
clear security firewall-authentication users identifier 5522
clear security user-identification local-authentication-table 5523
clear services user-identification active-directory-access 5524
request services user-identification active-directory-access active-directory-authentication-table delete 5525
request services user-identification active-directory-access domain-controller 5527
request services user-identification active-directory-access ip-user-probe . . . 5528
show network-access requests pending 5529
show network-access requests statistics 5531
show network-access securid-node-secret-file 5532
show security firewall-authentication history 5533
show security firewall-authentication history address 5535
show security firewall-authentication history identifier 5538
show security firewall-authentication users 5541
show security firewall-authentication users address 5543
show security firewall-authentication users identifier 5546
show security policies 5549
show services unified-access-control authentication-table 5557
show services unified-access-control counters 5559
show services unified-access-control policies 5561
show services unified-access-control roles 5563
show services unified-access-control status 5564
show services user-identification active-directory-access active-directory-authentication-table 5565
show services user-identification active-directory-access domain-controller status 5569
show services user-identification active-directory-access statistics 5572
show services user-identification active-directory-access user-group-mapping 5575

Copyright © 2016, Juniper Networks, Inc.

Guide 15 UTM Feature Guide for Security Devices

Part 71 Overview

Chapter 230 **Understanding Unified Threat Management 5583**

Unified Threat Management Overview 5583
Understanding UTM Custom Objects 5585

Chapter 231 **Managing UTM Licensing 5587**

Understanding UTM Licensing 5587
Updating UTM Licenses (CLI Procedure) 5588

Chapter 232 **Configuring WELF Logging 5589**

Understanding WELF Logging for UTM Features 5589
Example: Configuring WELF Logging for UTM Features 5590

Chapter 233 **Configuring UTM for Chassis Cluster 5593**

Understanding UTM Support for Active/Active Chassis Cluster 5593
Understanding Chassis Cluster Support for UTM Modules 5595

Part 72 Configuring Antispam Filtering

Chapter 234 **Understanding Antispam Filtering 5599**

Antispam Filtering Overview 5599
Handling Spam Messages 5599
Blocking Detected Spam 5599
Tagging Detected Spam 5600

Chapter 235 **Configuring Server-Based Antispam Filtering 5601**

Understanding Server-Based Antispam Filtering 5601
Server-Based Antispam Filtering Configuration Overview 5602
Example: Configuring Server-Based Antispam Filtering 5603

Chapter 236 **Configuring Local List Antispam Filtering 5609**

Understanding Local List Antispam Filtering 5609
Local List Antispam Filtering Configuration Overview 5610
Example: Configuring Local List Antispam Filtering 5610

Part 73 Configuring Express Antivirus Protection and Pattern Updates

Chapter 237 **Configuring Express Antivirus Protection 5621**

Express Antivirus Protection Overview 5621
Express Antivirus Packet-Based Scanning Versus File-Based Scanning . . 5621
Express Antivirus Expanded MIME Decoding Support 5622
Express Antivirus Scan Result Handling 5622
Express Antivirus Intelligent Prescreening 5622
Express Antivirus Limitations 5622
Express Antivirus Configuration Overview 5623
Example: Configuring Express Antivirus Custom Objects 5624
Configuring Express Antivirus Custom Objects (J-Web Procedure) 5626
Example: Configuring Express Antivirus Feature Profiles 5629
Configuring Express Antivirus Feature Profiles (J-Web Procedure) 5634

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Express Antivirus UTM Policies 5636
Configuring Express Antivirus UTM Policies (J-Web Procedure) 5637
Example: Attaching Express Antivirus UTM Policies to Security Policies 5638
Attaching Express Antivirus UTM Policies to Security Policies (J-Web Procedure) 5639

Chapter 238 **Configuring Express Antivirus Pattern Updates 5641**

Understanding Express Antivirus Scanner Pattern Updates 5641
Example: Automatically Updating Express Antivirus Patterns 5642
Example: Automatically Updating Express Antivirus Patterns (J-Web) 5643
Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure) 5643

Part 74 Configuring Full Antivirus Protection and Pattern Updates

Chapter 239 **Configuring Full Antivirus Protection 5647**

Full Antivirus Protection Overview 5647
Full Antivirus Configuration Overview 5648
Full Antivirus Pattern Update Configuration Overview 5649
Example: Configuring Full Antivirus Custom Objects 5650
Configuring Full Antivirus Custom Objects (J-Web Procedure) 5653
Example: Configuring Full Antivirus Feature Profiles 5655
Configuring Full Antivirus Feature Profiles (J-Web Procedure) 5661
Example: Configuring Full Antivirus UTM Policies 5664
Configuring Full Antivirus UTM Policies (J-Web Procedure) 5665
Example: Attaching Full Antivirus UTM Policies to Security Policies 5665
Attaching Full Antivirus UTM Policies to Security Policies (J-Web Procedure) . . 5666

Chapter 240 **Configuring Full Antivirus Pattern Updates 5669**

Understanding Full Antivirus Pattern Updates 5669
Example: Configuring the Full Antivirus Pattern Update Server 5670
Example: Automatically Updating Full Antivirus Patterns (J-Web) 5671
Example: Automatically Updating Full Antivirus Patterns 5672
Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) 5673

Chapter 241 **Configuring File Scanning 5675**

Understanding the Full Antivirus Internal Scan Engine 5675
Understanding Full Antivirus Scan Mode Support 5676
Configuring Full Antivirus File Extension Scanning (CLI Procedure) 5677
Example: Configuring Full Antivirus File Extension Scanning 5677
Understanding Full Antivirus Scan Level Settings 5678
Example: Configuring Full Antivirus Scan Settings at Different Levels 5679
Understanding Full Antivirus Intelligent Prescreening 5681
Example: Configuring Full Antivirus Intelligent Prescreening 5681
Understanding Full Antivirus Content Size Limits 5682
Configuring Full Antivirus Content Size Limits (CLI Procedure) 5683
Understanding Full Antivirus Decompression Layer Limits 5683
Configuring Full Antivirus Decompression Layer Limits (CLI Procedure) 5684
Understanding Full Antivirus Scanning Timeouts 5684

Copyright © 2016, Juniper Networks, Inc.

Configuring Full Antivirus Scanning Timeouts (CLI Procedure) 5684
Understanding Full Antivirus Scan Session Throttling 5685
Configuring Full Antivirus Scan Session Throttling (CLI Procedure) 5685

Chapter 242 Configuring Scan Results and Fallback Options 5687

Understanding Full Antivirus Scan Result Handling 5687
Monitoring Antivirus Scan Engine Status 5687
Monitoring Antivirus Session Status 5688
Monitoring Antivirus Scan Results 5689
Understanding Antivirus Scanning Fallback Options 5691
Example: Configuring Antivirus Scanning Fallback Options 5692

Chapter 243 Configuring Application Protocol Scanning 5695

Understanding Full Antivirus Application Protocol Scanning 5695
Understanding HTTP Scanning 5696
Enabling HTTP Scanning (CLI Procedure) 5697
Understanding FTP Antivirus Scanning 5697
Enabling FTP Antivirus Scanning (CLI Procedure) 5698
Understanding SMTP Antivirus Scanning 5698
Understanding SMTP Antivirus Mail Message Replacement 5699
Understanding SMTP Antivirus Sender Notification 5699
Understanding SMTP Antivirus Subject Tagging 5700
Enabling SMTP Antivirus Scanning (CLI Procedure) 5700
Understanding POP3 Antivirus Scanning 5700
Understanding POP3 Antivirus Mail Message Replacement 5701
Understanding POP3 Antivirus Sender Notification 5701
Understanding POP3 Antivirus Subject Tagging 5701
Enabling POP3 Antivirus Scanning (CLI Procedure) 5702
Understanding IMAP Antivirus Scanning 5702
Understanding IMAP Antivirus Mail Message Replacement 5702
Understanding IMAP Antivirus Sender Notification 5703
Understanding IMAP Antivirus Subject Tagging 5703
Understanding IMAP Antivirus Scanning Limitations 5703
Enabling IMAP Antivirus Scanning (CLI Procedure) 5704

Chapter 244 Configuring Whitelists 5705

Understanding MIME Whitelists 5705
Example: Configuring MIME Whitelists to Bypass Antivirus Scanning 5706
Understanding URL Whitelists 5706
Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure) . . . 5707

Chapter 245 Configuring HTTP Trickling 5709

Understanding HTTP Trickling 5709
Configuring HTTP Trickling to Prevent Timeouts During Antivirus Scanning (CLI Procedure) 5709

Chapter 246 Configuring Notifications 5711

Understanding Protocol-Only Virus-Detected Notifications 5711
Configuring Protocol-Only Virus-Detected Notifications (CLI Procedure) 5711
Understanding E-Mail Virus-Detected Notifications 5712
Configuring E-Mail Virus-Detected Notifications (CLI Procedure) 5712

Copyright © 2016, Juniper Networks, Inc.

Understanding Custom Message Virus-Detected Notifications 5713
Configuring Custom Message Virus-Detected Notifications (CLI Procedure) . . 5713

Part 75 Configuring and Managing Sophos Antivirus Protection

Chapter 247 Configuring Sophos Antivirus Protection 5717

Sophos Antivirus Protection Overview 5717
Sophos Antivirus Features 5718
Understanding Sophos Antivirus Data File Update 5719
Comparison of Sophos Antivirus to Kaspersky Antivirus 5720
Sophos Antivirus Configuration Overview 5720
Example: Configuring Sophos Antivirus Custom Objects 5721
Example: Configuring Sophos Antivirus Feature Profile 5724
Example: Configuring Sophos Antivirus UTM Policies 5731
Example: Configuring Sophos Antivirus Firewall Security Policies 5732
Example: Configuring Sophos Antivirus Scanner with SSL Forward Proxy 5734
Managing Sophos Antivirus Data Files 5740

Part 76 Configuring and Monitoring Content Filtering

Chapter 248 Configuring Content Filtering 5745

Content Filtering Overview 5745
Understanding Content Filtering Protocol Support 5746
HTTP Support 5746
FTP Support 5747
E-Mail Support 5747
Specifying Content Filtering Protocols (CLI Procedure) 5747
Content Filtering Configuration Overview 5748
Example: Configuring Content Filtering Custom Objects 5749
Example: Configuring Content Filtering Feature Profiles 5751
Example: Configuring Content Filtering UTM Policies 5754
Example: Attaching Content Filtering UTM Policies to Security Policies 5756
Monitoring Content Filtering Configurations 5758

Part 77 Configuring Web Filtering

Chapter 249 Configuring Web Filtering 5763

Web Filtering Overview 5764
Enhanced Web Filtering Overview 5766
Understanding Enhanced Web Filtering Process 5767
Functional Requirements for Enhanced Web Filtering 5768
Example: Configuring Enhanced Web Filtering 5772
Understanding the Quarantine Action for Enhanced Web Filtering 5780
Example: Configuring Site Reputation Action for Enhanced Web Filtering 5781
Understanding Integrated Web Filtering 5787
Integrated Web Filtering Process 5787
Integrated Web Filtering Cache 5788
Integrated Web Filtering Profiles 5788
Profile Matching Precedence 5789
Example: Configuring Integrated Web Filtering 5789

Copyright © 2016, Juniper Networks, Inc.

Understanding Local Web Filtering 5797
User-Defined URL Categories 5798
Local Web Filtering Process 5798
Local Web Filtering Profiles 5798
Profile Matching Precedence 5799
Example: Configuring Local Web Filtering 5799
Understanding Redirect Web Filtering 5805
Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects 5806
Displaying Global SurfControl URL categories 5815
Monitoring Web Filtering Configurations 5815

Part 78 Configuration Statements and Operational Commands

Chapter 250 Configuration Statements 5819

Security Configuration Statement Hierarchy 5824
SMTP Configuration Statement Hierarchy 5825
[edit security policies] Hierarchy Level 5826
[edit security utm] Hierarchy Level 5830
action (Security UTM Web Filtering) 5838
address-blacklist 5838
address-whitelist 5839
admin-email 5839
administrator-email (Security Fallback Block) 5840
administrator-email (Security Virus Detection) 5840
allow-email (Security Fallback Block) 5841
allow-email (Security Virus Detection) 5841
application (Security Policies) 5842
application-proxy (Security UTM) 5843
anti-spam (Security Feature Profile) 5844
anti-spam (Security UTM Policy) 5844
anti-virus (Security Feature Profile) 5845
anti-virus (Security UTM Policy) 5849
block-command 5849
block-content-type 5850
block-extension 5850
block-message (Security UTM) 5851
block-mime 5851
cache 5852
category (Security Logging) 5853
category (Security Web Filtering) 5854
content-filtering (Security Feature Profile) 5860
content-filtering (Security UTM Policy) 5861
content-size 5862
content-size (Security Antivirus Sophos Engine) 5863
content-size-limit 5864
corrupt-file 5864
custom-block-message 5865
custom-message (Security Content Filtering) 5865

Copyright © 2016, Juniper Networks, Inc.

custom-message (Security Email Notify) 5866
custom-message (Security Fallback Block) 5866
custom-message (Security Fallback Non-Block) 5867
custom-message (Security Virus Detection) 5867
custom-message-subject (Security Email Notify) 5868
custom-message-subject (Security Fallback Block) 5868
custom-message-subject (Security Fallback Non-Block) 5869
custom-message-subject (Security Virus Detection) 5869
custom-objects 5870
custom-tag-string 5871
custom-url-category 5872
decompress-layer 5873
decompress-layer-limit 5873
default (Security Antivirus) 5874
default (Security Antivirus Sophos Engine) 5874
default (Security UTM) 5875
default (Security Web Filtering) 5876
display-host (Security Fallback Block) 5877
display-host (Security Virus Detection) 5877
download-profile (Security Antivirus FTP) 5878
download-profile (Security Content Filtering FTP) 5878
email-notify 5879
engine-not-ready 5879
engine-not-ready (Security Antivirus Sophos Engine) 5880
exception (Security Antivirus Mime Whitelist) 5880
exception (Security Content Filtering) 5881
fallback-block (Security Antivirus) 5881
fallback-non-block (Security Antivirus) 5882
fallback-options (Security Antivirus Juniper Express Engine) 5883
fallback-options (Security Antivirus Kaspersky Lab Engine) 5884
fallback-options (Security Antivirus Sophos Engine) 5885
fallback-settings (Security Web Filtering) 5886
fallback-settings (Security Web Filtering Juniper Local) 5886
fallback-settings (Security Web Filtering Websense Redirect) 5887
feature-profile 5888
filename-extension 5893
flag (SMTP) 5894
format (Security Log Stream) 5895
from-zone (Security Policies) 5896
ftp (UTM Policy Anti-Virus) 5898
ftp (UTM Policy Content Filtering) 5899
host (Security Web Filtering) 5899
http-profile (Security Antivirus) 5900
http-profile (Security Content Filtering) 5900
http-profile (Security Web Filtering) 5901
imap-profile (Security UTM Policy Antivirus) 5901
imap-profile (Security UTM Policy Content Filtering) 5902
intelligent-prescreening 5902
interval (Security Antivirus) 5903

Copyright © 2016, Juniper Networks, Inc.

ipc . . . 5904
juniper-enhanced . . . 5905
juniper-express-engine . . . 5906
juniper-local . . . 5908
limit (UTM Policy) . . . 5908
list (Security Antivirus Mime Whitelist) . . . 5909
list (Security Content Filtering Block Mime) . . . 5909
log (Security) . . . 5910
mime-pattern . . . 5912
mime-whitelist . . . 5913
no-autoupdate . . . 5914
no-intelligent-prescreening . . . 5915
no-notify-mail-recipient . . . 5916
no-notify-mail-sender (Security Content Filtering Notification Options) . . . 5916
no-notify-mail-sender (Security Fallback Block) . . . 5917
no-notify-mail-sender (Security Virus Detection) . . . 5917
no-sbl-default-server . . . 5918
notification-options (Security Antivirus) . . . 5919
notification-options (Security Content Filtering) . . . 5920
notify-mail-recipient . . . 5920
notify-mail-sender (Security Content Filtering Notification Options) . . . 5921
notify-mail-sender (Security Fallback Block) . . . 5921
notify-mail-sender (Security Virus Detection) . . . 5922
no-uri-check . . . 5922
out-of-resources . . . 5923
out-of-resources (Security Antivirus Sophos Engine) . . . 5924
over-limit . . . 5924
packet-filter . . . 5925
password (Security Antivirus) . . . 5926
password-file . . . 5926
pattern-update (Security Antivirus) . . . 5927
permit-command . . . 5928
policies . . . 5929
pop3-profile (Security UTM Policy Antivirus) . . . 5933
pop3-profile (Security UTM Policy Content Filtering) . . . 5934
port (Security Antivirus) . . . 5934
port (Security Web Filtering Server) . . . 5935
primary-server . . . 5935
profile (Security Antispam SBL) . . . 5936
profile (Security Antivirus Juniper Express Engine) . . . 5937
profile (Security Antivirus Kaspersky Lab Engine) . . . 5939
profile (Security Content Filtering) . . . 5940
profile (Security Sophos Engine Antivirus) . . . 5941
profile (Security Web Filtering Juniper Enhanced) . . . 5942
profile (Security Web Filtering Juniper Local) . . . 5943
profile (Security Web Filtering Surf Control Integrated) . . . 5944
profile (Security Web Filtering Websense Redirect) . . . 5945
protocol-command . . . 5946
proxy (Security Antivirus) . . . 5946

Copyright © 2016, Juniper Networks, Inc.

quarantine-message (Security UTM) 5947
sbl 5948
sbl-default-server 5948
scan-extension 5949
scan-mode 5949
scan-options (Security Antivirus Juniper Express Engine) 5950
scan-options (Security Antivirus Kaspersky Lab Engine) 5950
scan-options (Security Antivirus Sophos Engine) 5951
secondary-server 5951
server (Security Antivirus) 5952
server (Security Web Filtering) 5952
server-connectivity 5953
sessions-per-client 5954
site-reputation-action 5955
size (Security Web Filtering Cache) 5956
smtp-profile (Security UTM Policy Antispam) 5956
smtp-profile (Security UTM Policy Antivirus) 5957
smtp-profile (Security UTM Policy Content Filtering) 5957
sockets 5958
sophos-engine 5959
spam-action 5960
surf-control-integrated 5961
sxl-retry 5962
sxl-timeout 5962
timeout (Security Antivirus Fallback Options) 5963
timeout (Security Antivirus Fallback Options Sophos Engine) 5964
timeout (Security Antivirus Scan Options) 5964
timeout (Security Web Filtering) 5965
timeout (Security Web Filtering Cache) 5965
timeout (Security Web Filtering Fallback Settings) 5966
too-many-requests (Security Antivirus Fallback Options) 5967
too-many-requests (Security Antivirus Fallback Options Sophos Engine) . . . 5967
too-many-requests (Security Web Filtering Fallback Settings) 5968
to-zone (Security Policies) 5969
traceoptions (Security Antispam) 5971
traceoptions (Security Antivirus) 5972
traceoptions (Security Application Proxy) 5973
traceoptions (Security Content Filtering) 5974
traceoptions (Security UTM) 5975
traceoptions (Security Web Filtering) 5976
traceoptions (SMTP) 5977
traffic-options 5977
trickling 5978
type (Security Antivirus Feature Profile) 5979
type (Security Content Filtering Notification Options) 5979
type (Security Fallback Block) 5980
type (Security Virus Detection) 5981
type (Security Web Filtering) 5982
upload-profile (Security Antivirus FTP) 5982

Copyright © 2016, Juniper Networks, Inc.

upload-profile (Security Content Filtering FTP) . . . 5983
uri-check . . . 5983
url (Security Antivirus) . . . 5984
url-blacklist . . . 5984
url-pattern . . . 5985
url-whitelist (Security Antivirus) . . . 5985
url-whitelist (Security Web Filtering) . . . 5986
username (Security Antivirus) . . . 5986
utm . . . 5987
utm-policy . . . 5995
utm-policy (Application Services) . . . 5996
virus-detection (Security Antivirus) . . . 5997
web-filtering . . . 5998
websense-redirect . . . 5999

Chapter 251 Operational Commands . . . 6001

clear security utm anti-spam statistics . . . 6002
clear security utm antivirus statistics . . . 6003
clear security utm content-filtering statistics . . . 6004
clear security utm session . . . 6005
clear security utm web-filtering statistics . . . 6006
request security utm anti-virus juniper-express-engine . . . 6007
request security utm anti-virus kaspersky-lab-engine . . . 6008
request security utm anti-virus sophos-engine . . . 6009
request system license update . . . 6010
show configuration smtp . . . 6011
show groups junos-defaults . . . 6012
show security log . . . 6013
show security policies . . . 6016
show security utm anti-spam statistics . . . 6024
show security utm anti-spam status . . . 6025
show security utm anti-virus statistics . . . 6026
show security utm anti-virus status . . . 6028
show security utm content-filtering statistics . . . 6030
show security utm session . . . 6031
show security utm status . . . 6032
show security utm web-filtering statistics . . . 6033
show security utm web-filtering status . . . 6036

Guide 16 VPN Feature Guide for Security Devices

Part 79 Overview

Chapter 252 Introduction to IPsec VPNs . . . 6041

IPsec VPN Overview . . . 6041
IPsec VPN Topologies . . . 6042
Comparison of Policy-Based VPNs and Route-Based VPNs . . . 6042
Security Associations . . . 6043

Copyright © 2016, Juniper Networks, Inc.

IPsec Key Management 6044
Manual Key 6044
AutoKey IKE 6045
Diffie-Hellman Exchange 6045
IPsec Security Protocols 6046
AH Protocol 6046
ESP Protocol 6047
IPsec Tunnel Negotiation 6047
Understanding IKE and IPsec Packet Processing 6048
Packet Processing in Tunnel Mode 6049
IKE Packet Processing 6050
IPsec Packet Processing 6053
Understanding Phase 1 of IKE Tunnel Negotiation 6055
Main Mode 6056
Aggressive Mode 6057
Understanding Phase 2 of IKE Tunnel Negotiation 6057
Proxy IDs 6058
Perfect Forward Secrecy 6058
Replay Protection 6058
IPsec VPN with Autokey IKE Configuration Overview 6059
IPsec VPN with Manual Keys Configuration Overview 6060
Recommended Configuration Options for Site-to-Site VPN with Static IP Addresses 6060
Recommended Configuration Options for Site-to-Site or Dialup VPNs with Dynamic IP Addresses 6061
Configuring Remote IKE IDs for Site-to-Site VPNs 6062
Configuring IPsec VPN Using the VPN Wizard 6063
Understanding Suite B and PRIME Cryptographic Suites 6064

Chapter 253 Understanding VPN Tunnel Management 6067

Understanding Distributed VPNs in SRX Series Services Gateways 6067
Understanding VPN Support for Inserting Services Processing Cards 6068

Part 80 Configuring Route-Based IPsec VPNs

Chapter 254 Configuring Route-Based VPNs 6073

Understanding Route-Based IPsec VPNs 6073
Example: Configuring a Route-Based VPN 6074

Chapter 255 Configuring Hub-and-Spoke VPNs 6093

Understanding Hub-and-Spoke VPNs 6093
Example: Configuring a Hub-and-Spoke VPN 6094

Chapter 256 Configuring VPNs for IKEv2 6127

Understanding Internet Key Exchange Version 2 6127
Understanding IKEv2 Configuration Payload 6128
Example: Configuring a Route-Based VPN for IKEv2 6130
Understanding Pico Cell Provisioning 6146
Example: Configuring the SRX Series for Pico Cell Provisioning with IKEv2 Configuration Payload 6149

Copyright © 2016, Juniper Networks, Inc.

Chapter 257 Configuring Secure Tunnel Interface in a Virtual Router 6175
Understanding Virtual Router Support for Route-Based VPNs 6175
Understanding Virtual Router Limitations 6176
Example: Configuring an st0 Interface in a Virtual Router 6176
Chapter 258 Configuring Dual Stack Tunnels over an External Interface 6181
Understanding VPN Tunnel Modes 6181
Understanding Dual-Stack Tunnels over an External Interface 6183
Example: Configuring Dual-Stack Tunnels over an External Interface 6184
Chapter 259 Configuring Traffic Selectors in Route-Based VPNs 6195
Understanding Traffic Selectors in Route-Based VPNs 6195
Traffic Selector Configuration 6195
Traffic Selector Flexible Matches 6196
Multiple Tunnels for Traffic Selector Configuration 6197
Limitations 6198
Example: Configuring Traffic Selectors in a Route-Based VPN 6198
Understanding Auto Route Insertion 6213
Understanding Traffic Selectors and Overlapping IP Addresses 6213
Overlapping IP Addresses in Different VPNs Bound to the Same st0 Interface 6214
Overlapping IP Addresses in the Same VPN Bound to the Same st0 Interface 6214
Overlapping IP Addresses in Different VPNs Bound to Different st0 Interfaces 6215
Part 81 Configuring Policy-Based IPsec VPNs
Chapter 260 Configuring Policy-Based VPNs 6221
Understanding Policy-Based IPsec VPNs 6221
Example: Configuring a Policy-Based VPN 6222
Part 82 Configuring VPNs with NAT-T
Chapter 261 Configuring Route-Based and Policy-Based VPNs with NAT-T 6243
Understanding NAT-T 6243
Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device 6244
Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device 6272
Example: Configuring NAT-T with Dynamic Endpoint VPN 6298
Part 83 Configuring IPsec VPN Tunnels with Chassis Clusters
Chapter 262 Configuring IPsec VPN Tunnels with Chassis Clusters 6319
Understanding Dual Active-Backup IPsec VPN Chassis Clusters 6319
Understanding Loopback Interface for a High Availability VPN 6320
Example: Configuring Redundancy Groups for Loopback Interfaces 6321

Copyright © 2016, Juniper Networks, Inc.

Part 84 Configuring IPv6 IPsec VPNs

Chapter 263 Configuring IPv6 IPsec VPNs 6331

VPN Feature Support for IPv6 Addresses 6331
Understanding IPv6 IKE and IPsec Packet Processing 6335
IPv6 IKE Packet Processing 6335
IPv6 IPsec Packet Processing 6336
AH Protocol in IPv6 6337
ESP Protocol in IPv6 6337
IPv4 Options and IPv6 Extension Headers with AH and ESP 6338
Integrity Check Value Calculation in IPv6 6338
Header Construction in Tunnel Modes 6339
IPv6 IPsec Configuration Overview 6341
Example: Configuring an IPv6 IPsec Manual VPN 6341

Part 85 Configuring Public Key Infrastructure

Chapter 264 Managing Digital Certificates with PKI 6347

Understanding Certificates and PKI 6347
Certificate Signatures and Verification 6347
Public Key Infrastructure 6348
PKI Management and Implementation 6350
Internet Key Exchange 6351
Cryptographic Key Handling Overview 6351
Understanding CMPv2 and SCEP Certificate Enrollment 6352
Understanding Certificate Enrollment with CMPv2 6353
Certificate Enrollment and Reenrollment Messages 6353
End-Entity Certificate with Issuer CA Certificate 6353
End-Entity Certificate with CA Certificate Chain 6354
Digital Certificates Configuration Overview 6355
Enrolling Digital Certificates Online: Configuration Overview 6355
Manually Generating Digital Certificates: Configuration Overview 6356

Chapter 265 Configuring Digital Certificate Validation 6357

Understanding Digital Certificate Validation 6357
Policy Validation 6357
Policy OIDs Configured on SRX Series Devices 6358
No Policy OIDs Configured on SRX Series Devices 6358
Path Length Validation 6359
Key Usage 6360
EE Certificates 6360
CA Certificates 6360
Issuer and Subject Distinguished Name Validation 6361
Example: Improving Digital Certificate Validation by Configuring Policy OIDs on an SRX Series Device 6362

Chapter 266 Generating a Public-Private Key Pair 6367

Understanding Public Key Cryptography 6367
Example: Generating a Public-Private Key Pair 6368

Copyright © 2016, Juniper Networks, Inc.

Chapter 267 Configuring Certificate Authority Profiles 6369
Understanding Certificate Authority Profiles 6369
Example: Configuring a CA Profile 6369
Chapter 268 Configuring CA and Local Certificates 6371
Understanding Online CA Certificate Enrollment 6371
Understanding Local Certificate Requests 6371
Enrolling a CA Certificate Online Using SCEP 6372
Example: Enrolling a Local Certificate Online Using SCEP 6373
Example: Using SCEP to Automatically Renew a Local Certificate 6375
Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server 6376
Understanding Certificate Loading 6377
Example: Loading CA and Local Certificates Manually 6378
Deleting Certificates (CLI Procedure) 6379
Example: Configuring PKI 6380
Chapter 269 Managing Certificate Revocation 6409
Understanding Online Certificate Status Protocol 6409
Understanding Certificate Revocation Lists 6410
Comparison of Online Certificate Status Protocol and Certificate Revocation List 6411
Improving Security by Configuring OCSP for Certificate Revocation Status . . . 6412
Example: Manually Loading a CRL onto the Device 6427
Example: Configuring a Certificate Authority Profile with CRL Locations 6428
Example: Verifying Certificate Validity 6429
Deleting a Loaded CRL (CLI Procedure) 6431
Chapter 270 Generating Self-Signed Certificates 6433
Understanding Self-Signed Certificates 6433
Generating Self-Signed Certificates 6433
Automatically Generating Self-Signed Certificates 6434
Manually Generating Self-Signed Certificates 6434
Example: Manually Generating Self-Signed Certificates 6434
Using Automatically Generated Self-Signed Certificates (CLI Procedure) 6435
Chapter 271 Configuring a Device for Certificate Chains 6437
Understanding Certificate Chains 6437
Multilevel Hierarchy for Certificate Authentication 6437
Dynamic CRL Download and Checking 6439
Example: Configuring a Device for Peer Certificate Chain Validation 6440
Part 86 Configuring AutoVPN
Chapter 272 Configuring AutoVPN on Hub-and-Spoke Devices 6451
Understanding AutoVPN 6451
Secure Tunnel Modes 6451
Authentication 6452
Configuration and Management 6452
Understanding AutoVPN Limitations 6453

Copyright © 2016, Juniper Networks, Inc.

Understanding Spoke Authentication in AutoVPN Deployments 6454
Group IKE ID Configuration on the Hub 6454
Excluding a Spoke Connection 6456
AutoVPN Configuration Overview 6456
Example: Configuring Basic AutoVPN with iBGP 6457
Example: Configuring Basic AutoVPN with OSPF 6482
Example: Configuring AutoVPN with iBGP and ECMP 6506
Example: Configuring AutoVPN with iBGP and Active-Backup Tunnels 6531

Chapter 273 **Configuring Auto Discovery VPNs 6559**

Understanding Auto Discovery VPN 6559
ADVPN Protocol 6559
Establishing a Shortcut 6560
Shortcut Initiator and Responder Roles 6561
Shortcut Attributes 6561
Shortcut Termination 6562
ADVPN Configuration Limitations 6563
Understanding Traffic Routing with Shortcut Tunnels 6564
Example: Improving Network Resource Utilization with Auto Discovery VPN Dynamic Tunnels 6566
Enabling OSPF to Update Routes Quickly After ADVPN Shortcut Tunnels Are Established 6604

Chapter 274 **Configuring AutoVPN and Traffic Selectors 6607**

Understanding AutoVPN with Traffic Selectors 6607
Example: Forwarding Traffic Through an AutoVPN Tunnel with Traffic Selectors 6608
Example: Ensuring VPN Tunnel Availability with AutoVPN and Traffic Selectors 6624

Part 87 Configuring Group VPNs

Chapter 275 **Managing IPsec SA By Configuring Group VPNs 6649**

Group VPNv2 Overview 6649
Understanding the GDOI Protocol 6650
Understanding Group Servers and Members 6651
Understanding Group VPNv2 Limitations 6652
Group VPNv2 Configuration Overview 6652
Understanding IKE Phase 1 Configuration for Group VPNv2 6653
Understanding IPsec SA Configuration for Group VPNv2 6654
Understanding Group Configuration 6655
Understanding Traffic Steering 6655
Group Policies Configured on Group Servers 6656
IPsec Policies Configured on Group Members 6656
Fail-Close 6656
Exclude and Fail-Open Rules 6656
Priorities of IPsec Policies and Rules 6657
Example: Configuring a Group VPNv2 Server and Members 6657

Copyright © 2016, Juniper Networks, Inc.

Chapter 276 Configuring Server-Member Communication 6693
Understanding Server-Member Communication 6693
Understanding Group Key Operations 6694
Group Keys 6694
Rekey Messages 6694
Member Registration 6695
Example: Configuring Server-Member Communication for Unicast Rekey Messages 6695
Understanding the Recovery Probe Process 6696
Understanding Antireplay 6697
Chapter 277 Configuring Group VPNv2 Server Clusters 6699
Understanding Group VPNv2 Server Clusters 6699
Root-Server and Sub-Servers 6699
Group Member Registration with Server Clusters 6701
Dead Peer Detection 6702
Load Balancing 6702
Understanding Group VPNv2 Server Cluster Limitations 6703
Understanding Group VPNv2 Server Cluster Messages 6704
Cluster Exchanges 6704
Cluster-Init Exchanges 6704
Cluster-Update Messages 6705
Understanding Configuration Changes with Server Clusters 6706
Migrating a Standalone Group Server to a Server Cluster 6708
Example: Configuring a Group VPNv2 Server Cluster and Members 6709
Part 88 Monitoring and Improving VPN Traffic Performance
Chapter 278 Configuring VPN Monitoring Features 6777
Understanding VPN Alarms and Auditing 6777
Example: Setting an Audible Alert as Notification of a Security Alarm 6779
Example: Generating Security Alarms in Response to Potential Violations . . . 6780
Understanding VPN Monitoring and DPD 6782
Understanding Dead Peer Detection 6783
Understanding VPN Monitoring 6785
Understanding Global SPI and VPN Monitoring Features 6785
Example: Configuring Global SPI and VPN Monitoring Features 6786
Understanding Tunnel Events 6786
Chapter 279 Improving IPsec VPN Traffic Performance 6789
Understanding VPN Session Affinity 6789
Enabling VPN Session Affinity 6791
Accelerating the IPsec VPN Traffic Performance 6792
Part 89 Troubleshooting
Chapter 280 Tunnel Events 6797
Tunnel Events 6797

Copyright © 2016, Juniper Networks, Inc.

Part 90 Configuration Statements and Operational Commands

Chapter 281 Configuration Statements 6809

Security Configuration Statement Hierarchy 6814
Access Configuration Statement Hierarchy 6815
[edit security alarms] Hierarchy Level 6823
[edit security group-vpn] Hierarchy Level 6824
[edit security ike] Hierarchy Level 6827
[edit security ipsec] Hierarchy Level 6829
[edit security pki] Hierarchy Level 6831
address (Security Group VPN Server IKE Gateway) 6832
address (Security IKE Gateway) 6833
administrator 6833
advpn 6834
algorithm (Security) 6835
always-send 6836
authentication (IPsec SA for OSPF) 6837
authentication (Security IPsec) 6838
authentication-algorithm (Security Group VPN IKE) 6839
authentication-algorithm (Security Group VPN IPsec) 6839
authentication-algorithm (Security IKE) 6840
authentication-algorithm (Security IPsec) 6841
authentication-method 6842
authentication-method (Security Group VPN) 6843
auto-re-enrollment (Security) 6844
auxiliary-spi (IPsec SA for OSPF) 6845
bind-interface 6845
ca-identity (Security) 6846
ca-profile (Security PKI) 6847
ca-profile-name 6848
certificate 6848
certificate-id (Security) 6849
challenge-password (Security) 6850
connections-limit 6850
container 6851
crl (Security) 6852
cryptographic-self-test 6853
dead-peer-detection 6853
dead-peer-detection (Security Group VPN Server) 6854
decryption-failures 6854
description (Security Policies) 6855
destination-ip (Security IPsec) 6855
df-bit 6856
dh-group (Security IKE) 6857
dh-group (Security Group VPN IKE) 6858
disable (PKI) 6858
distinguished-name (Security) 6859
dynamic (Security) 6859
dynamic (Security Group VPN) 6860

Copyright © 2016, Juniper Networks, Inc.

encryption (IPsec SA for OSPF) . . . 6861
encryption (Security) . . . 6862
encryption-algorithm (Security Group VPN IKE) . . . 6863
encryption-algorithm (Security Group VPN IPsec) . . . 6864
encryption-algorithm (Security IKE) . . . 6865
encryption-algorithm (Security IPsec) . . . 6866
encryption-failures . . . 6867
enrollment (Security) . . . 6868
establish-tunnels . . . 6869
external-interface (Security IKE Gateway) . . . 6869
external-interface (Security Manual SA) . . . 6870
gateway (Security Group VPN Member IKE) . . . 6870
gateway (Security Group VPN Server IKE) . . . 6871
gateway (Security IKE) . . . 6872
gateway (Security IPsec VPN) . . . 6873
gateway (Security Manual SA) . . . 6873
general-ikeid . . . 6874
group (Security Group VPN) . . . 6875
group-vpn . . . 6877
hostname . . . 6880
idle-time . . . 6880
ike (Security) . . . 6881
ike (Security Group VPN Member) . . . 6883
ike (Security Group VPN Server) . . . 6885
ike (Security IPsec VPN) . . . 6886
ike-phase1-failures . . . 6887
ike-phase2-failures . . . 6888
ike-policy (Security Gateway) . . . 6888
ike-user-type . . . 6889
inet (Security Dynamic Peer) . . . 6889
inet6 (Security IKE Gateway) . . . 6890
install-interval . . . 6890
interval (Security IKE) . . . 6891
ipsec (Security) . . . 6892
ipsec (Security Group VPN Member) . . . 6895
ipsec (Security Group VPN Server) . . . 6897
ipsec-performance-acceleration (Security Flow) . . . 6897
ipsec-policy (Security) . . . 6898
ipsec-policy (Security Group VPN) . . . 6898
ipsec-sa (Security Group VPN) . . . 6899
ipsec-vpn (Security Flow) . . . 6900
key-generation-self-test . . . 6900
lifetime-kilobytes . . . 6901
lifetime-seconds (Security Group VPN) . . . 6901
lifetime-seconds (Security IKE) . . . 6902
lifetime-seconds (Security IPsec) . . . 6902
load-distribution . . . 6903
local (Security IPsec) . . . 6903
local-address . . . 6904

Copyright © 2016, Juniper Networks, Inc.

local-address (Security Group VPN Member) 6905
local-address (Security Group VPN Server) 6905
local-certificate (Security) 6906
local-identity 6907
local-identity (Security Group VPN) 6908
manual (Security IPsec) 6909
member (Security Group VPN) 6910
member-threshold (Security Group VPN) 6912
mode (Security Group VPN) 6913
mode (Security IKE Policy) 6914
nat-keepalive 6915
no-anti-replay (Security) 6915
no-nat-traversal 6916
non-cryptographic-self-test 6916
ocsp (Security PKI) 6917
optimized 6918
optimized (DPD) 6918
peer-certificate-type 6919
perfect-forward-secrecy (Security IPsec) 6920
pki 6921
pki-local-certificate 6922
policy (Security Group VPN IKE) 6923
policy (Security IKE) 6924
policy (Security IPsec) 6925
policy-oids 6926
pre-shared-key (Security IKE Policy) 6927
probe-idle-tunnel 6927
profile (Access) 6928
proposal (Security Group VPN Member IKE) 6931
proposal (Security Group VPN Server IKE) 6932
proposal (Security Group VPN Server IPsec) 6932
proposal (Security IKE) 6933
proposal (Security IPsec) 6934
proposals (Security Group VPN) 6934
proposals (Security IKE) 6935
proposals (Security IPsec) 6935
proposal-set (Security IKE) 6936
proposal-set (Security IPsec) 6939
protocol (IPsec SA for OSPF) 6941
protocol (Security IPsec) 6941
protocol (Security IPsec Manual SA) 6942
proxy-identity 6942
re-enroll-trigger-time-percentage (Security PKI) 6943
re-generate-keypair 6943
refresh-interval 6944
remote (Security IPsec) 6944
remote-identity 6945
remote-identity (Security Group VPN) 6946
replay-attacks 6947

Copyright © 2016, Juniper Networks, Inc.

respond-bad-spi . . . 6947
revocation-check (Security PKI) . . . 6948
routing-instance (Security Group VPN) . . . 6949
routing-instance (Security PKI) . . . 6949
security-association . . . 6950
server (Security Group VPN) . . . 6951
server-address (Security Group VPN Member) . . . 6953
server-cluster (Security Group VPN Server) . . . 6954
server-member-communication (Security Group VPN Server) . . . 6955
service (Security IPsec) . . . 6956
session-affinity . . . 6956
source-interface (Security) . . . 6957
spi (IPsec SA for OSPF) . . . 6957
spi (Security IPsec) . . . 6958
threshold (Security IKE Gateway) . . . 6958
traceoptions (Security Group VPN) . . . 6959
traceoptions (Security IKE) . . . 6962
traceoptions (Security IPsec) . . . 6964
traceoptions (Security PKI) . . . 6965
traffic-selector . . . 6966
trusted-ca (Security IKE Policy) . . . 6967
use-ocsp (Security PKI) . . . 6967
user-at-hostname . . . 6968
version (Security IKE Gateway) . . . 6968
vpn (Security) . . . 6969
vpn-monitor . . . 6970
vpn-monitor-options . . . 6971
wildcard . . . 6971
xauth . . . 6972

Chapter 282 Operational Commands . . . 6973

clear security group-vpn member group . . . 6975
clear security group-vpn member ike security-associations . . . 6976
clear security group-vpn member ipsec security-associations . . . 6977
clear security group-vpn member ipsec security-associations statistics . . . 6978
clear security group-vpn member ipsec statistics . . . 6979
clear security group-vpn server . . . 6980
clear security group-vpn server server-cluster statistics . . . 6981
clear security group-vpn server statistics . . . 6982
clear security ike respond-bad-spi-count . . . 6983
clear security ike security-associations . . . 6984
clear security ipsec security-associations . . . 6986
clear security ipsec statistics . . . 6987
clear security ipsec tunnel-events-statistics . . . 6988
clear security pki key-pair (Local Certificate) . . . 6989
clear security pki local-certificate (Device) . . . 6990
request security pki ca-certificate ca-profile-group load . . . 6991
request security pki ca-certificate enroll (Security) . . . 6993
request security pki ca-certificate load (Security) . . . 6994

Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate verify (Security) 6995
request security pki crl load (Security) 6996
request security pki generate-certificate-request (Security) 6997
request security pki generate-key-pair (Security) 6999
request security pki local-certificate enroll cmpv2 7000
request security pki local-certificate enroll scep 7002
request security pki local-certificate export 7004
request security pki local-certificate generate-self-signed (Security) 7005
request security pki local-certificate load 7007
request security pki local-certificate verify (Security) 7008
request security pki verify-integrity-status 7009
show security group-vpn member ike security-associations 7010
show security group-vpn member ipsec inactive-tunnels 7013
show security group-vpn member ipsec security-associations 7016
show security group-vpn member ipsec statistics 7020
show security group-vpn member kek security-associations 7022
show security group-vpn member policy 7025
show security group-vpn server ike security-associations 7027
show security group-vpn server ipsec security-associations 7031
show security group-vpn server kek security-associations 7034
show security group-vpn server registered-members 7037
show security group-vpn server server-cluster 7039
show security group-vpn server statistics 7042
show security ike active-peer 7043
show security ike pre-shared-key 7044
show security ike security-associations 7045
show security ike tunnel-map 7053
show security ipsec control-plane-security-associations 7056
show security ipsec inactive-tunnels 7058
show security ipsec next-hop-tunnels 7061
show security ipsec security-associations 7062
show security ipsec statistics 7074
show security ipsec traffic-selector 7077
show security ipsec tunnel-events-statistics 7079
show security pki ca-certificate (View) 7080
show security pki certificate-request (View) 7084
show security pki crl (View) 7086
show security pki local-certificate (View) 7089

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

List of Figures

Part 2 Configuring Data ALGs

Chapter 3 **Configuring the DNS ALG 13**

Figure 1: DNS Address Translation (Private to Public) 21
Figure 2: DNS Address Translation (Public to Private) 22

Chapter 6 **Configuring the PPTP ALG 45**

Figure 3: PPTP ALG Topology 48

Chapter 8 **Configuring the RSH ALG 75**

Figure 4: RSH ALG Topology 77

Chapter 9 **Configuring the RTSP ALG 89**

Figure 5: RTSP ALG Standard Mode 90
Figure 6: RTSP ALG Interleave Mode 90
Figure 7: RTSP ALG Conversation 94

Chapter 10 **Configuring the SQLNET ALG 103**

Figure 8: SQLNET ALG Topology 105

Chapter 11 **Configuring the TALK ALG 117**

Figure 9: TALK ALG Topology 119

Part 3 Configuring VoIP ALGs

Chapter 15 **Configuring the H.323 ALG 145**

Figure 10: H.323 ALG for VoIP Calls 146
Figure 11: H.323 Gatekeeper in the Private Zone 150
Figure 12: H.323 Gatekeeper in the External Zone 155
Figure 13: NAT with the H.323 ALG—Incoming Calls 161
Figure 14: NAT with the H.323 ALG—Outgoing Calls 169

Chapter 16 **Configuring the MGCP ALG 183**

Figure 15: Media Gateway in Subscriber Homes 190
Figure 16: Three-Zone ISP-Hosted Service 198

Chapter 17 **Configuring the SCCP ALG 217**

Figure 17: Call Setup and Teardown 221
Figure 18: Call Manager or TFTP Server in the Private Zone 225

Chapter 18 **Configuring the SIP ALG 241**

Figure 19: SIP ALG Call Setup 246
Figure 20: SIP NAT Scenario 1 262
Figure 21: SIP NAT Scenario 2 263

Copyright © 2016, Juniper Networks, Inc.

Figure 22: Using the SIP Registrar 267
Figure 23: Source NAT for Incoming SIP Calls 268
Figure 24: Source NAT Pool for Incoming SIP Calls 275
Figure 25: Static NAT for Incoming Calls 283
Figure 26: Configuring SIP Proxy in the Private Zone and NAT in a Public Zone 289
Figure 27: Three-Zone SIP Configuration with Proxy in the DMZ 295

Chapter 22 Understanding Application Identification 485

Figure 28: Mapping Sequence 487

Chapter 28 Configuring Encrypted Files Using SSL Proxy 531

Figure 29: SSL Inspection on an Existing SRX Series IDP Module 532
Figure 30: SSL Proxy on an Encrypted Payload 533
Figure 31: SSL Proxy Configuration Overview 541
Figure 32: Applying an SSL Proxy Profile to a Security Policy 547

Part 6 Configuring Screen Options to Protect Against Denial-of-Service Attacks

Chapter 36 Protecting Against Firewall DoS Attacks 861

Figure 33: Limiting Sessions Based on Source IP Address 863
Figure 34: Distributed DOS Attack 867

Chapter 37 Protecting Against Network DoS Attacks 873

Figure 35: SYN Flood Attack 874
Figure 36: Proxying SYN Segments 875
Figure 37: Device-Level SYN Flood Protection 880
Figure 38: Establishing a Connection with SYN Cookie Active 889
Figure 39: ICMP Flooding 894
Figure 40: UDP Flooding 897
Figure 41: Land Attack 900

Chapter 38 Protecting Against OS-Specific DoS Attacks 913

Figure 42: Ping of Death 914
Figure 43: Teardrop Attacks 916
Figure 44: Fragment Discrepancy 916
Figure 45: WinNuke Attack Indicators 918

Part 7 Configuring Reconnaissance Deterrence for Security Devices

Chapter 39 Protecting Against IP Sweep and Port Options 923

Figure 46: Address Sweep 924
Figure 47: Port Scan 927
Figure 48: UDP Port Scan 930
Figure 49: Routing Options 931

Chapter 40 Protecting Against System Probes and Flag Set 937

Figure 50: TCP Header with SYN and FIN Flags Set 938
Figure 51: TCP Header with FIN Flag Set 941
Figure 52: TCP Header with No Flags Set 943

Copyright © 2016, Juniper Networks, Inc.

Chapter 41 Protecting Against Attacker Evasion Techniques 947

Figure 53: SYN Flag Checking 949
Figure 54: IP Source Routing 955
Figure 55: Loose IP Source Route Option for Deception 956

Part 8 Configuring Suspicious Packet Attributes for Security Devices

Chapter 42 Protecting Against ICMP and SYN Fragment Attacks 963

Figure 56: Blocking ICMP Fragments 964
Figure 57: Blocking Large ICMP Packets 965
Figure 58: SYN Fragments 967

Chapter 43 Protecting Against IP Attacks 969

Figure 59: Incorrectly Formatted IP Options 970
Figure 60: Unknown Protocols 971
Figure 61: IP Packet Fragments 973
Figure 62: IPv6 Packet 973
Figure 63: Fragment Extension Header 973

Part 12 Configuring Address Books and Address Sets

Chapter 50 Configuring Address, Address Books, and Address Sets 1069

Figure 64: Applying Policies to Address Sets 1075
Figure 65: Configuring Addresses and Address Sets 1078

Part 13 Configuring Security Policies

Chapter 51 Enforcing Transit Traffic Rules by Configuring Security Policies 1085

Figure 66: Security Policy 1087
Figure 67: Permitting All Traffic 1095
Figure 68: Permitting Selected Traffic 1099

Chapter 53 Configuring Global Security Policy 1115

Figure 69: Multizone Global Policy Security Consideration 1116

Part 16 Overview

Chapter 65 Introduction to Class of Service 1401

Figure 70: CoS Across the Network 1403
Figure 71: Packet Flow Through Juniper Networks Device 1405

Part 17 Configuring Class of Service Components

Chapter 66 Assigning Service Levels with Classifiers 1411

Figure 72: Behavior Aggregate Classifier Scenario 1418

Chapter 67 Controlling Network Access with Traffic Policing 1427

Figure 73: Multifield Classifier Based on TCP Source Ports 1446
Figure 74: Multifield Classifier Scenario 1447

Chapter 72 Controlling Congestion with Drop Profiles 1519

Figure 75: Segmented and Interpolated Drop Profiles 1522

Copyright © 2016, Juniper Networks, Inc.

Chapter 75 Enabling Queuing for Tunnel Interfaces 1539
Figure 76: CoS Processing for Tunnel Traffic 1540
Figure 77: Configuring CoS Queuing for GRE Tunnels 1543

Part 18 Configuring Class of Service Scheduler Hierarchy

Chapter 77 Controlling Traffic by Configuring Scheduler Hierarchy 1557
Figure 78: Building a Scheduler Hierarchy 1562
Figure 79: Example 1 Handling Remaining Traffic with no Explicit Traffic Control Profile 1573
Figure 80: Example 2 Handling Remaining Traffic with an Interface Set 1574

Part 19 Configuring Class of Service for IPv6

Chapter 78 Configuring Class of Service for IPv6 Traffic 1581
Figure 81: Packet Flow Through an SRX Series Device 1581

Part 20 Configuring Class of Service for I/O Cards

Chapter 79 Configuring Class of Service for I/O Cards 1595
Figure 82: Hierarchical Schedulers and Priorities 1598

Part 22 Overview

Chapter 82 Introduction to Processing on Security Devices 1665
Figure 83: Traffic Flow for Flow-Based Processing 1666
Figure 84: First-Packet Processing 1673
Figure 85: Fast-Path Processing 1673
Figure 86: Session Creation: First-Packet Processing 1678
Figure 87: Packet Walk for Fast-Path Processing 1680

Part 24 Improving Flow-Based Performance

Chapter 86 Reducing Long Packet-Processing Latency by Express Path 1713
Figure 88: IOC3 Intra-PFE Express Path 1727
Figure 89: IOC3 Inter-PFE Express Path 1727
Figure 90: Inter-IOC3 Express Path 1728

Part 25 Managing Flow-Based Processing for IPv6

Chapter 89 Configuring IPv6 Dual-Stack 1773
Figure 91: DS-Lite NAT (IPv4-in-IPv6) 1774

Part 26 Monitoring Flow-Based Sessions

Chapter 91 Monitoring X2 Traffic By Configuring Mirror Filters 1793
Figure 92: SRX Series Device in an LTE Mobile Network 1794
Figure 93: Monitoring X2 Traffic 1794
Figure 94: Configuring Mirror Filters for X2 Traffic Monitoring 1796

Copyright © 2016, Juniper Networks, Inc.

Part 27 Configuring Packet-Based Forwarding

Chapter 92 Configuring Selective Stateless Packet-Based Services 1803

Figure 95: Traffic Flow for Packet-Based Forwarding 1803
Figure 96: Traffic Flow with Selective Stateless Packet-Based Services 1805
Figure 97: Intranet Traffic Using End-to-End Packet-Based Services 1809
Figure 98: Selective Stateless Packet-Based Services for Packet-Based Forwarding 1819

Part 29 Overview

Chapter 95 Introduction to General Packet Radio Service 2215

Figure 99: Gp and Gn Interfaces 2216
Figure 100: Gi Interface 2217

Part 30 Configuring GPRS Tunnel Protocol v1

Chapter 100 Configuring NAT for GTP 2265

Figure 101: NAT-PT Between an IPv4 Endpoint and an IPv6 Endpoint 2272

Chapter 101 Configuring GGSN 2279

Figure 102: GGSN Pooling Scenario 1 2280
Figure 103: Functionality : GGSN Pooling Scenario 1 2281
Figure 104: GGSN Pooling Scenario 2 2282
Figure 105: Functionality : GGSN Pooling Scenario 2 2283

Part 31 Configuring GPRS Tunnel Protocol v2

Chapter 102 Configuring GTPv2 2291

Figure 106: LTE Interfaces 2292

Part 32 Configuring Stream Control Transmission Protocol

Chapter 105 Configuring SCTP 2319

Figure 107: SCTP 4-way Handshake and TCP 3-way Handshake 2320
Figure 108: SCTP Multihoming with Two IPv4 Endpoints 2326
Figure 109: SCTP Multihoming with Two IPv6 Endpoints 2326
Figure 110: SCTP Packet Structure 2328
Figure 111: SCTP Firewall Implementation 2334

Part 34 Overview

Chapter 109 Configuring Interface Logical Properties 2451

Figure 112: Subnets in a Network 2454

Chapter 111 Configuring VLAN Tagging 2483

Figure 113: Typical LAN 2484
Figure 114: Typical VLAN 2484

Part 36 Configuring DSL Interfaces

Chapter 113 Configuring VDSL2 Interfaces 2501

Copyright © 2016, Juniper Networks, Inc.

Figure 115: Typical VDSL2 End-to-End Connectivity and Topology Diagram . . . 2503
Figure 116: Backward-Compatible ADSL Topology (ATM DSLAM) 2503
Figure 117: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario 2514

Part 37 Configuring Ethernet Interfaces

Chapter 114 **Performing Initial Configuration on Ethernet Interfaces 2541**

Figure 118: Ethernet Frame Format 2544

Chapter 118 **Configuring Ethernet OAM Link Fault Management 2603**

Figure 119: Ethernet LFM with SRX Series Devices 2606

Part 38 Configuring Interface Encapsulation

Chapter 120 **Interface Encapsulation Overview 2623**

Figure 120: Frame Relay Network 2624

Chapter 121 **Configuring Point-to-Point Protocol over Ethernet 2631**

Figure 121: PPPoE Session on the Ethernet Loop 2641

Part 41 Configuring Layer 2 Bridging and Transparent Mode

Chapter 128 **Configuring Interfaces 2905**

Figure 122: Architecture of Mixed Layer 2 and Layer 3 Mode 2911
Figure 123: Mixed Layer 2 and Layer 3 Mode 2912
Figure 124: Mixed Mode Topology 2915

Chapter 135 **Configuring Secure Wire 2955**

Figure 125: SRX Series In-Path Deployment with Secure Wire 2956
Figure 126: Secure Wire Access Mode Interfaces 2958
Figure 127: Secure Wire Trunk Mode Interfaces 2961
Figure 128: Secure Wire Aggregated Interfaces 2965
Figure 129: Secure Wire Redundant Ethernet Interfaces 2969
Figure 130: Secure Wire Redundant Ethernet Interface Child Links 2974

Part 44 Overview

Chapter 142 **Introduction to Logical Systems 3205**

Figure 131: Understanding Logical Systems 3206
Figure 132: Logical Systems, Their Virtual Routers, and Their Interfaces 3212

Part 45 Getting Started for Master Administrators

Chapter 145 **Configuring Device for Master Logical Systems 3241**

Figure 133: SRX Series Device Configured for Logical Systems 3243

Part 47 Configuring Routing and Interfaces Features

Chapter 149 **Configuring Master Logical System Routing and Interfaces 3341**

Figure 134: Configuring Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers 3344

Copyright © 2016, Juniper Networks, Inc.

Part 48 Configuring Logical Systems in Chassis Cluster

Chapter 151 Configuring Logical Systems When Device is in Chassis Cluster Mode . . 3369

Figure 135: Logical Systems in a Chassis Cluster 3373
Figure 136: Logical Systems in a Chassis Cluster (IPv6) 3406

Part 49 Configuring IPv6 for Logical Systems

Chapter 152 Configuring IPv6 Addresses for Logical Systems 3439

Figure 137: Configuring IPv6 Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers 3443

Part 53 Overview

Chapter 157 Introduction to MPLS 3743

Figure 138: Typical LSP Topology 3745

Part 54 Configuring Traffic Engineering

Chapter 158 Configuring MPLS Traffic Engineering and Signaling Protocols 3755

Figure 139: Typical LDP-Signaled LSP 3758
Figure 140: Typical RSVP-Signaled LSP with EROs 3763
Figure 141: Typical RSVP-Signaled LSP 3765
Figure 142: Point-to-Multipoint LSPs 3769
Figure 143: RSVP-Signaled Point-to-Multipoint LSP 3771

Part 55 Configuring MPLS VPNs

Chapter 159 Introduction to MPLS VPNs 3791

Figure 144: Typical VPN Topology 3792

Part 57 Configuring VPLS

Chapter 168 Introduction to VPLS 3841

Figure 145: Basic VPLS Topology 3842
Figure 146: Flooding a Packet with an Unknown Destination 3844

Chapter 171 Configuring Routing and Signaling Protocols 3867

Figure 147: VPLS Deployment Scenario 3874
Figure 148: Branch Office Circuit Cross Connect Termination 3875
Figure 149: Central Office Ingress (Head-End) Configuration with an SRX Series Device 3876
Figure 150: Central Office Ingress (Head-End) Configuration with an MX Series Device 3877
Figure 151: Configuring VPLS with BGP Signaling 3890

Part 59 Overview

Chapter 175 Introduction to Multicast 4019

Figure 152: Multicast Terminology in an IP Network 4023
Figure 153: Converting MAC Addresses to Multicast Addresses 4027

Copyright © 2016, Juniper Networks, Inc.

Part 60 Managing Group Membership

Chapter 177 Examples: Configuring MLD 4063

Figure 154: Routers Start Up on a Subnet 4064
Figure 155: Querier Router Is Determined 4064
Figure 156: General Query Message Is Issued 4065
Figure 157: Reports Are Received by the Querier Router 4065
Figure 158: Host Has No Interested Receivers and Sends a Done Message to Router 4065
Figure 159: Host Address Timer Expires and Address Is Removed from Multicast Address List 4066

Part 61 Configuring Protocol Independent Multicast

Chapter 180 Routing Content to Densely Clustered Receivers with PIM Dense Mode 4107

Figure 160: Multicast Traffic Flooded from the Source Using PIM Dense Mode 4108
Figure 161: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic 4109

Chapter 181 Routing Content to Larger, Sparser Groups with PIM Sparse Mode 4113

Figure 162: Rendezvous Point As Part of the RPT and SPT 4115
Figure 163: Join Suppression 4124
Figure 164: PIM Sparse Mode over an IPsec VPN 4127
Figure 165: Virtual Router Instance with Three Interfaces 4132
Figure 166: Extracting the Embedded RP IPv6 Address 4157
Figure 167: Building an RPT Between the RP and the Receiver 4168
Figure 168: PIM Register Message and PIM Join Message Exchanged 4169
Figure 169: Traffic Sent from the Source to the RP Router 4170
Figure 170: Traffic Sent from the RP Router Toward the Receiver 4170
Figure 171: Receiver DR Sends a PIM Join Message to the Source 4172
Figure 172: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router 4173
Figure 173: RP Router Receives PIM Prune Message 4173
Figure 174: RP Router Sends a PIM Prune Message to the Source DR 4174
Figure 175: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router 4174
Figure 176: PIM Assert Topology 4176

Chapter 182 Receiving Content Directly from the Source with SSM 4183

Figure 177: Receiver Announces Desire to Join Group G and Source S 4185
Figure 178: Router 3 (Last-Hop Router) Joins the Source Tree 4185
Figure 179: (S,G) State Is Built Between the Source and the Receiver 4186
Figure 180: Receiver Sends Messages to Join Group G and Source S 4187
Figure 181: Router 3 (Last-Hop Router) Joins the Source Tree 4188
Figure 182: (S,G) State Is Built Between the Source and the Receiver 4188
Figure 183: Simple RPF Topology 4188
Figure 184: Network on Which to Configure PIM SSM 4191

Chapter 183 Minimizing Routing State Information with Bidirectional PIM 4199

Copyright © 2016, Juniper Networks, Inc.

Figure 185: Example PIM Sparse-Mode Tree 4200
Figure 186: Example Bidirectional PIM Tree 4201
Figure 187: Bidirectional PIM with Statically Configured Rendezvous Points . . 4206

Chapter 185 Configuring PIM Options 4227

Figure 188: Nonstop Active Routing in PIM Domain 4229

Part 62 Configuring Multicast Routing Protocols

Chapter 186 Improving Multicast Reliability with PGM 4247

Figure 189: PGM Architecture and General Operation 4251

Chapter 187 Connecting Routing Domains Using MSDP 4253

Figure 190: MSDP in a VRF Instance Topology 4258
Figure 191: Source-Active Message Flooding 4265

Chapter 189 Facilitating Multicast Delivery Across Unicast-Only Networks with AMT 4273

Figure 192: Automatic Multicast Tunneling Connectivity 4274
Figure 193: AMT Gateway Topology 4282

Part 63 Configuring Multicast VPNs

Chapter 191 Configuring PIM Join Load Balancing 4299

Figure 194: PIM Join Load Balancing 4301

Chapter 192 Configuring Next-Generation Multicast VPNs 4303

Figure 195: PIM Join Load Balancing on Next-Generation MVPN 4306

Part 64 Configuring General Multicast Routing Options

Chapter 193 Preventing Routing Loops with Reverse Path Forwarding 4315

Figure 196: Multicast Routers and the RPF Check 4316
Figure 197: PIM RPF Selection 4325

Chapter 194 Enabling Multicast Between Layer 2 and Layer 3 Devices Using Snooping 4329

Figure 198: Networks Without IGMP Snooping Configured 4338
Figure 199: Networks with IGMP Snooping Configured 4339
Figure 200: VPLS Multihoming Topology 4348

Chapter 195 Configuring Multicast Routing Options 4353

Figure 201: Multicast with Subscriber VLANs 4362

Chapter 198 Overview 4839

Figure 202: NAT Rule Processing 4842

Chapter 200 Configuring Source NAT 4863

Figure 203: Source NAT Egress Interface Translation 4866
Figure 204: Source NAT Single Address Translation 4870
Figure 205: Source and Destination NAT Translations 4874
Figure 206: Source NAT with Multiple Translation Rules 4883

Chapter 201 Configuring Source NAT Pools 4895

Copyright © 2016, Juniper Networks, Inc.

Figure 207: Source NAT with Address Shifting 4903
Figure 208: Source NAT Multiple Addresses with PAT 4909
Figure 209: Source NAT Multiple Addresses Without PAT 4918

Chapter 202 Configuring Destination NAT 4923

Figure 210: Destination NAT Single Address Translation 4928
Figure 211: Destination NAT Address and Port Translation 4935
Figure 212: Destination NAT Subnet Translation 4940

Chapter 203 Configuring Static NAT 4947

Figure 213: Static NAT Single Address Translation 4950
Figure 214: Static NAT Subnet Translation 4955
Figure 215: Static NAT for Port Mapping 4960

Chapter 204 Configuring Persistent NAT and NAT64 4969

Figure 216: 464XLAT Architecture 4972
Figure 217: NAT64 Translation on the PLAT (SRX Series Device) 4973
Figure 218: Interface Persistent NAT Topology 4978

Chapter 205 Configuring NAT Hairpinning 4993

Figure 219: Persistent NAT Hairpinning 4995

Chapter 206 Configuring NAT for Multicast Flows 5001

Figure 220: NAT Translations for Multicast Flows 5003

Chapter 208 Configuring IPv6 Dual-Stack 5031

Figure 221: DS-Lite NAT (IPv4-in-IPv6) 5032

Part 66 Overview

Chapter 211 Introduction to User Authentication 5173

Figure 222: Three-Tiered User Firewall Features 5174

Part 67 Configuring Firewall User Authentication

Chapter 213 Configuring Encrypted Files Using SSL Proxy 5183

Figure 223: SSL Inspection on an Existing SRX Series IDP Module 5184
Figure 224: SSL Proxy on an Encrypted Payload 5185
Figure 225: SSL Proxy Configuration Overview 5193
Figure 226: Applying an SSL Proxy Profile to a Security Policy 5198

Chapter 214 Configuring Pass-Through Authentication 5205

Figure 227: Policy Lookup for a User 5206
Figure 228: Configuring Pass-Through Firewall Authentication 5208
Figure 229: Pass-Through Authentication Using HTTPS Traffic 5215

Chapter 215 Configuring Web Authentication 5221

Figure 230: Web Authentication Example 5222
Figure 231: Web Authentication Example 5224
Figure 232: Web Authentication Success Banner 5224
Figure 233: Web Authentication Using HTTPS Traffic 5231

Chapter 218 Customizing the Firewall Authentication Banner 5251

Copyright © 2016, Juniper Networks, Inc.

Figure 234: Banner Customization 5251

Part 68 Configuring Infranet Authentication

Chapter 219 **Configuring UAC in a Junos OS Environment 5257**

Figure 235: Integrating a Junos OS Security Device into a Unified Access Control Network 5258

Chapter 222 **Classifying Traffic with User Roles 5269**

Figure 236: Single Sign-On Support Topology 5272

Chapter 225 **Configuring Captive Portal 5299**

Figure 237: Enabling the Captive Portal Feature on a Junos OS Enforcer 5300

Part 69 Configuring Integrated User Firewall

Chapter 226 **Understanding Integrated User Firewalls 5309**

Figure 238: Scenario for Integrated User Firewall 5311

Part 77 Configuring Web Filtering

Chapter 249 **Configuring Web Filtering 5763**

Figure 239: Websense Redirect Architecture 5808

Part 79 Overview

Chapter 252 **Introduction to IPsec VPNs 6041**

Figure 240: Tunnel Mode 6049
Figure 241: Site-to-Site VPN in Tunnel Mode 6049
Figure 242: Dial-Up VPN in Tunnel Mode 6050
Figure 243: IKE Packet for Phases 1 and 2 6051
Figure 244: Generic ISAKMP Payload Header 6052
Figure 245: ISAKMP Header with Generic ISAKMP Payloads 6053
Figure 246: IPsec Packet—ESP in Tunnel Mode 6053
Figure 247: Outer IP Header (IP2) and ESP Header 6054
Figure 248: Inner IP Header (IP1) and TCP Header 6055

Part 80 Configuring Route-Based IPsec VPNs

Chapter 254 **Configuring Route-Based VPNs 6073**

Figure 249: Route-Based VPN Topology 6075

Chapter 255 **Configuring Hub-and-Spoke VPNs 6093**

Figure 250: Multiple Tunnels in a Hub-and-Spoke VPN Configuration 6093
Figure 251: Hub-and-Spoke VPN Topology 6095

Chapter 256 **Configuring VPNs for IKEv2 . . . . . 6127**

Figure 252: Typical Pico Cell Deployment Workflow 6148
Figure 253: SRX Series Support for Pico Cell Provisioning with IKEv2 Configuration Payload 6150

Chapter 258 **Configuring Dual Stack Tunnels over an External Interface 6181**

Copyright © 2016, Juniper Networks, Inc.

Figure 254: IPv4-in-IPv4 Tunnel 6181
Figure 255: IPv6-in-IPv6 Tunnel 6182
Figure 256: IPv6-in-IPv4 Tunnel 6182
Figure 257: IPv4-in-IPv6 Tunnel 6182
Figure 258: Dual-Stack Tunnels 6183
Figure 259: Dual-Stack Tunnel Example 6187

Chapter 259 Configuring Traffic Selectors in Route-Based VPNs 6195

Figure 260: Multiple Tunnels for Traffic Selector Configuration 6197
Figure 261: Traffic Selector Configuration Example 6200

Part 81 Configuring Policy-Based IPsec VPNs

Chapter 260 Configuring Policy-Based VPNs 6221

Figure 262: Policy-Based VPN Topology 6223

Part 82 Configuring VPNs with NAT-T

Chapter 261 Configuring Route-Based and Policy-Based VPNs with NAT-T 6243

Figure 263: Route-Based VPN Topology with Only the Responder Behind a NAT Device 6246
Figure 264: Policy-Based VPN Topology with Both an Initiator and a Responder Behind a NAT Device 6273
Figure 265: NAT-T with Dynamic Endpoint VPN 6299

Part 83 Configuring IPsec VPN Tunnels with Chassis Clusters

Chapter 262 Configuring IPsec VPN Tunnels with Chassis Clusters 6319

Figure 266: Active/Passive Chassis Cluster with IPsec VPN Tunnels 6319
Figure 267: Dual Active-Backup IPsec VPN Chassis Clusters 6319
Figure 268: Loopback Interface for Chassis Cluster VPN 6323

Part 84 Configuring IPv6 IPsec VPNs

Chapter 263 Configuring IPv6 IPsec VPNs 6331

Figure 269: IPv6 AH Tunnel Mode 6337
Figure 270: IPv6 ESP Tunnel Mode 6338

Part 85 Configuring Public Key Infrastructure

Chapter 264 Managing Digital Certificates with PKI 6347

Figure 271: Digital Signature Verification 6348
Figure 272: PKI Hierarchy of Trust—CA Domain 6349
Figure 273: Cross-Certification 6350
Figure 274: End-Entity Certificate with CA Certificate Chain 6354

Chapter 265 Configuring Digital Certificate Validation 6357

Figure 275: Policy Validation with requireExplicitPolicy Field 6358
Figure 276: Policy Validation with skipCerts Field 6359
Figure 277: Path Length Validation 6360
Figure 278: Issuer and Subject DN Validation 6361

Copyright © 2016, Juniper Networks, Inc.

Chapter 268 Configuring CA and Local Certificates 6371
Figure 279: Network Topology Diagram 6381
Chapter 269 Managing Certificate Revocation 6409
Figure 280: OCSP Configuration Example 6414
Chapter 271 Configuring a Device for Certificate Chains 6437
Figure 281: Multilevel Hierarchy for Certificate-Based Authentication 6438
Figure 282: Certificate Chain Example 6441

Part 86 Configuring AutoVPN

Chapter 272 Configuring AutoVPN on Hub-and-Spoke Devices 6451
Figure 283: Basic AutoVPN Deployment with iBGP 6460
Figure 284: Basic AutoVPN Deployment with OSPF 6485
Figure 285: AutoVPN Deployment with iBGP and ECMP 6509
Figure 286: AutoVPN Deployment with iBGP and Active-Backup Tunnels . . . 6534
Chapter 273 Configuring Auto Discovery VPNs 6559
Figure 287: Spoke-to-Spoke Traffic Passing Through Hub 6560
Figure 288: Spoke-to-Spoke Traffic Passing Through Shortcut 6561
Figure 289: Static Tunnels and Shortcut Tunnel Established in Hub-and-Spoke Network 6564
Figure 290: Traffic Path from Spoke C to Spoke A 6564
Figure 291: Traffic Path from Spoke C to Spoke A Through Shortcut Tunnels . . 6565
Figure 292: Traffic Path from Spoke C to Spoke A Through Shortcut Tunnel . . 6565
Figure 293: AutoVPN Deployment with ADVPN 6569
Chapter 274 Configuring AutoVPN and Traffic Selectors 6607
Figure 294: AutoVPN with Traffic Selectors 6611
Figure 295: Georedundant IPsec VPN Gateways to eNodeB Devices 6626

Part 87 Configuring Group VPNs

Chapter 275 Managing IPsec SA By Configuring Group VPNs 6649
Figure 296: Point-to-Point SAs 6649
Figure 297: Shared SAs 6650
Figure 298: Group VPNv2 Server with SRX or vSRX and MX Series Members . . 6659
Chapter 277 Configuring Group VPNv2 Server Clusters 6699
Figure 299: Group VPNv2 Server Cluster 6700
Figure 300: Group VPNv2 Server Cluster Messages 6704
Figure 301: Group VPNv2 Server Cluster with SRX or vSRX and MX Series Members 6712

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

List of Tables

About the Documentation cxliii

Table 1: Notice Icons cxlv
Table 2: Text and Syntax Conventions cxlv

Part 2 Configuring Data ALGs

Chapter 7 Configuring the RPC ALG 61

Table 3: Predefined Sun RPC Services 66
Table 4: Predefined MS-RPC services 73

Chapter 9 Configuring the RTSP ALG 89

Table 5: RTSP Payload IP NAT 94

Part 3 Configuring VoIP ALGs

Chapter 16 Configuring the MGCP ALG 183

Table 6: MGCP Commands 186
Table 7: Three-Zone ISP-Host Service 198

Chapter 17 Configuring the SCCP ALG 217

Table 8: Station to Call Manager Messages 221
Table 9: Call Manager to Station Messages 222
Table 10: Call Manager 4.0 Messages and Post Sccp 6.2 222
Table 11: Call Manager to Station 222

Chapter 18 Configuring the SIP ALG 241

Table 12: Requesting Messages with NAT Table 260
Table 13: SIP Responses 264

Part 4 Configuration Statements and Operational Commands

Chapter 19 Configuration Statements 309

Table 14: Port Supported by Services Interfaces 348
Table 15: Category Names 408

Chapter 20 Operational Commands 415

Table 16: show chassis cluster data-plane statistics Output Fields 424
Table 17: show chassis cluster statistics Output Fields 426
Table 18: show security alg h323 counters Output Fields 430
Table 19: show security alg msrpc Output Fields 434
Table 20: show security alg sccp calls Output Fields 436
Table 21: show security alg sccp counters Output Fields 437
Table 22: show security alg sip calls Output Fields 439

Copyright © 2016, Juniper Networks, Inc.

Table 23: show security alg sip counters Output Fields 441
Table 24: show security alg sip rate Output Fields 445
Table 25: show security flow gate Output Fields 449
Table 26: show security flow session Output Fields 455
Table 27: show security flow session application Output Fields 460
Table 28: show security flow session resource-manager Output Fields 463
Table 29: show security resource-manager group Output Fields 468
Table 30: show security resource-manager resource Output Fields 471
Table 31: show security resource-manager summary Output Fields 474
Table 32: show security zones Output Fields 475
Table 33: show security zones type Output Fields 478

Chapter 28 Configuring Encrypted Files Using SSL Proxy 531

Table 34: Supported SSL Cipher List 533
Table 35: SSL Proxy Logs 537
Table 36: SSL Proxy Log Prefixes 537
Table 37: Trace Levels 552
Table 38: Supported Flags in Trace 553

Chapter 29 Configuring Application Firewall 555

Table 39: Application Firewall Actions 560

Chapter 31 Configuring Application QoS 583

Table 40: Standard CoS Aliases and Bit Values 585

Chapter 33 Operational Commands 743

Table 41: show class-of-service application-traffic-control counter Output Fields 769
Table 42: show class-of-service application-traffic-control statistics rate-limiter Output Fields 771
Table 43: show class-of-service application-traffic-control statistics rule Output Fields 773
Table 44: show security application-firewall rule-set Output Fields 774
Table 45: show security application-firewall rule-set logical-system Output Fields 778
Table 46: show security application-tracking counters 780
Table 47: show security flow session Output Fields 783
Table 48: show security flow session application-firewall extensive Output Fields 788
Table 49: show security pki ca-certificate Output Fields 793
Table 50: show security pki local-certificate Output Fields 797
Table 51: show security policies Output Fields 803
Table 52: show services application-identification application Output Fields . . 810
Table 53: show services application-identification application-system-cache Output Fields 813
Table 54: show services application-identification counter Output Fields 816
Table 55: show services application-identification group Output Fields 819
Table 56: show services application-identification statistics applications Output Fields 821

Copyright © 2016, Juniper Networks, Inc.

Table 57: show services application-identification statistics application-groups Output Fields 823
Table 58: show services application-identification status Output Fields 825
Table 59: show services ssl proxy statistics Output Fields 829

Part 5 Overview

Chapter 34 **Introduction to Attack Detection and Prevention 835**

Table 60: IPv6 Extension Headers and Type Values 836
Table 61: Bad Option Extension Header Screening Criteria 838
Table 62: IPv6 Packet Header Screening Criteria 839
Table 63: Statistics-Based Screen Options 840
Table 64: Signature-Based Screen Options 841
Table 65: Statistics-Based Screen Options 845

Part 6 Configuring Screen Options to Protect Against Denial-of-Service Attacks

Chapter 37 **Protecting Against Network DoS Attacks 873**

Table 66: SYN Flood Protection Parameters 881

Part 7 Configuring Reconnaissance Deterrence for Security Devices

Chapter 39 **Protecting Against IP Sweep and Port Options 923**

Table 67: IP Options and Attributes 931

Part 9 Configuration Statements and Operational Commands

Chapter 45 **Operational Commands 1021**

Table 68: show security screen ids-option Output Fields 1027
Table 69: show security screen statistics Output Fields 1033

Part 11 Configuring Security Zones and Interfaces

Chapter 48 **Managing Inbound Traffic for Security Zones 1055**

Table 70: Supported Inbound System Protocols 1058
Table 71: System Services for Host Inbound Traffic 1061
Table 72: Protocols for Host Inbound Traffic 1062

Part 12 Configuring Address Books and Address Sets

Chapter 50 **Configuring Address, Address Books, and Address Sets 1069**

Table 73: Available Addresses Displayed in the CLI 1074

Part 13 Configuring Security Policies

Chapter 55 **Configuring User Role Firewall Security Policies 1127**

Table 74: Trust Zone to Untrust Zone Policy Sequence 1129
Table 75: UIT Authentication Details 1129
Table 76: Trust Zone to Untrust Zone Policy Sequence 1130
Table 77: User Role Firewall Policies 1139

Copyright © 2016, Juniper Networks, Inc.

Table 78: User Role Firewall Usage 1145
Table 79: User Role Firewall Usage 1147

Chapter 57 Monitoring and Troubleshooting Security Policies 1153

Table 80: Policy Limitations for High-End SRX Series Devices 1155

Part 14 Configuring Security Policy Applications

Chapter 61 Setting Policy Application Timeout 1181

Table 81: Protocol-Based Default Timeout 1181

Chapter 62 Understanding Predefined Policy Applications 1185

Table 82: Predefined Applications 1185
Table 83: Predefined Microsoft Applications 1187
Table 84: Dynamic Routing Protocols 1188
Table 85: Supported Streaming Video Applications 1189
Table 86: RPC ALG Applications 1189
Table 87: Supported Applications 1190
Table 88: Predefined IP-Related Applications 1191
Table 89: Predefined Internet-Messaging Applications 1191
Table 90: Predefined Management Applications 1192
Table 91: Predefined Mail Applications 1193
Table 92: Predefined UNIX Applications 1194
Table 93: Predefined Miscellaneous Applications 1194
Table 94: ICMP Messages 1195
Table 95: Message Descriptions 1200

Part 15 Configuration Statements and Operational Commands

Chapter 63 Configuration Statements 1205

Table 96: Port Supported by Services Interfaces 1245
Table 97: Category Names 1327

Chapter 64 Operational Commands 1339

Table 98: show security alarms 1350
Table 99: show security firewall-authentication users address Output Fields . . 1353
Table 100: show security firewall-authentication users auth-type Output Fields 1356
Table 101: show security flow session application Output Fields 1359
Table 102: show security match-policies Output Fields 1363
Table 103: show security policies Output Fields 1368
Table 104: show security policies hit-count Output Fields 1376
Table 105: show security policies unknown-source-identity Output Fields 1378
Table 106: show security shadow-policies logical-system Output Fields 1380
Table 107: show security user-identification local-authentication-table Output Fields 1381
Table 108: show security user-identification role-provision all Output Fields . . 1383
Table 109: show security user-identification source-identity-provision all Output Fields 1384
Table 110: show security user-identification user-provision all Output Fields . . 1385
Table 111: show security zones Output Fields 1386

Copyright © 2016, Juniper Networks, Inc.

Table 112: show security zones type Output Fields 1389
Table 113: show system services dns-proxy 1392
Table 114: show system services dynamic-dns 1395

Part 16 Overview

Chapter 65 Introduction to Class of Service 1401

Table 115: Supported Junos OS CoS Components 1404
Table 116: Reasons to Configure Class of Service (CoS) 1408

Part 17 Configuring Class of Service Components

Chapter 66 Assigning Service Levels with Classifiers 1411

Table 117: BA Classification 1412
Table 118: MF Classification 1413
Table 119: Default IP Precedence Classifier 1414
Table 120: Default Behavior Aggregate Classification 1415
Table 121: Sample Behavior Aggregate Classification Forwarding Classes and Queues 1416
Table 122: Sample ba-classifier Loss Priority Assignments 1418

Chapter 67 Controlling Network Access with Traffic Policing 1427

Table 123: Two-Color Policer Configuration and Application Overview 1435
Table 124: Simple Filter Match Conditions 1444

Chapter 68 Controlling Output Queues with Forwarding Classes 1453

Table 125: Default Forwarding Class Queue Assignments 1454
Table 126: Sample Output Queue Assignments for mf-classifier Forwarding Queues 1460

Chapter 69 Altering Outgoing Packets Headers with Rewrite Rules 1467

Table 127: Sample rewrite-dscps Rewrite Rules to Replace DSCPs 1469

Chapter 70 Defining Output Queue Properties with Schedulers 1473

Table 128: Sample Transmission Scheduling 1479
Table 129: Shaping Rates and WFQ Weights 1480
Table 130: Example Shaping Rates and WFQ Weights 1481
Table 131: Rounding Configured Weights to Hardware Weights 1482
Table 132: Allocating Weights with PIR and CIR on Logical Interfaces 1483
Table 133: Example of Shared Bandwidth Among Logical Interfaces 1484
Table 134: First Example of Bandwidth Sharing 1484
Table 135: Second Example of Bandwidth Sharing 1485
Table 136: Final Example of Bandwidth Sharing 1485
Table 137: Sample Schedulers 1486
Table 138: Maximum Available Delay Buffer Time by Channelized Interface and Rate 1490
Table 139: Delay Buffer Size Allocation Methods 1491
Table 140: Delay Buffer Allocation Method and Queue Buffer 1492
Table 141: Interface Delay Times Enabled By q-pic-large-buffer 1499
Table 142: Sample diffserv-cos-map Scheduler Mapping 1501

Chapter 72 Controlling Congestion with Drop Profiles 1519

Copyright © 2016, Juniper Networks, Inc.

Table 143: Sample RED Drop Profiles 1520
Table 144: Configuring RED Drop Profiles for Assured Forwarding Congestion Control 1521
Table 145: Sample RED Drop Profiles 1523

Chapter 76 **Naming Components with Code-Point Aliases 1549**

Table 146: Standard CoS Aliases and Bit Values 1551

Part 18 Configuring Class of Service Scheduler Hierarchy

Chapter 77 **Controlling Traffic by Configuring Scheduler Hierarchy 1557**

Table 147: Hierarchical Scheduler Nodes 1557

Part 19 Configuring Class of Service for IPv6

Chapter 78 **Configuring Class of Service for IPv6 Traffic 1581**

Table 148: Default IPv6 BA Classifier Mapping 1583

Part 20 Configuring Class of Service for I/O Cards

Chapter 79 **Configuring Class of Service for I/O Cards 1595**

Table 149: Internal Node Queue Priority for PIR-Only Mode 1595
Table 150: Internal Node Queue Priority for CIR Mode 1596
Table 151: Queue Priority 1597
Table 152: Packet Forwarding Engine Properties within 40x1GE IOC and 4x10GE IOC 1599
Table 153: Shaper Accuracy of 1-Gbps Ethernet at the Logical Interface Level . . 1602
Table 154: Shaper Accuracy of 10-Gbps Ethernet at the Logical Interface Level 1603
Table 155: Shaper Accuracy of 1-Gbps Ethernet at the Interface Set Level 1604
Table 156: Shaper Accuracy of 10-Gbps Ethernet at the Interface Set Level . . . 1604
Table 157: Shaper Accuracy of 1-Gbps Ethernet at the Physical Port Level 1604
Table 158: Shaper Accuracy of 10-Gbps Ethernet at the Physical Port Level . . 1604
Table 159: Junos Priorities Mapped to IOC Hardware Priorities 1605
Table 160: Forwarding Class Samples 1609
Table 161: Scheduler Samples 1609

Part 21 Configuration Statements and Operational Commands

Chapter 81 **Operational Commands 1653**

Table 162: show class-of-service application-traffic-control counter Output Fields 1654
Table 163: show class-of-service application-traffic-control statistics rate-limiter Output Fields 1656
Table 164: show class-of-service application-traffic-control statistics rule Output Fields 1658
Table 165: show class-of-service forwarding-class Output Fields 1659

Part 22 Overview

Chapter 82 **Introduction to Processing on Security Devices 1665**

Copyright © 2016, Juniper Networks, Inc.

Table 166: Combo Mode Processing 1670

Part 24 Improving Flow-Based Performance

Chapter 84 Expanding Session Capacity by Device 1703

Table 167: Maximum Central Point Session Increases 1703

Chapter 86 Reducing Long Packet-Processing Latency by Express Path 1713

Table 168: Session Cache Installation Bars 1715
Table 169: Session Cache Table Utilization Bits Status 1716
Table 170: Express Path Support on SRX Series Device Cards 1718
Table 171: Total Number of Sessions per Wing in Network Processor Express Path Configuration Mode 1720

Part 25 Managing Flow-Based Processing for IPv6

Chapter 87 Enabling IPv6 Flow-Based Processing 1751

Table 172: Device Status Upon Configuration Change 1757

Chapter 88 Managing IPv6 Packets 1763

Table 173: IPv6 Extension Headers 1766
Table 174: ICMPv6 Packets That Junos OS Handles Differently from Other ICMPv6 Packets 1770

Chapter 89 Configuring IPv6 Dual-Stack 1773

Table 175: Softwire Initiator and Softwire Concentrator Capacity 1775

Part 26 Monitoring Flow-Based Sessions

Chapter 90 Monitoring Security Flow Sessions 1781

Table 176: Session Create Log Fields 1787
Table 177: Session Close Log Fields 1788
Table 178: Session Deny Log Fields 1791

Chapter 91 Monitoring X2 Traffic By Configuring Mirror Filters 1793

Table 179: X2 Traffic Terminology 1795

Part 28 Configuration Statements and Operational Commands

Chapter 93 Configuration Statements 1831

Table 180: Ports Supported by Services Interfaces 1847
Table 181: Device Status Upon Configuration Change 1867
Table 182: Ports Supported by Services Interfaces 1880
Table 183: Session Capacity and Resulting Throughput 1894

Chapter 94 Operational Commands 1895

Table 184: show chassis environment Output Fields 1928
Table 185: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary 1932
Table 186: show chassis fpc Output Fields 1934
Table 187: show chassis fpc errors Output Fields 1940
Table 188: show chassis hardware Output Fields 1942
Table 189: show chassis pic Output Fields 1969

Copyright © 2016, Juniper Networks, Inc.

Table 190: show chassis power Output Fields 1971
Table 191: show chassis power sequence Output Fields 1973
Table 192: show firewall Output Fields 1974
Table 193: Aggregated Ethernet show interfaces Output Fields 1976
Table 194: show interfaces Output Fields 1989
Table 195: show interfaces diagnostics optics Output Fields 2017
Table 196: show interfaces flow-statistics Output Fields 2022
Table 197: Flow Error Statistics (Packet Drop Statistics for the Flow Module) . . 2023
Table 198: show interfaces <swfab0 | swfab1> Output Fields 2028
Table 199: show monitor security flow Output Fields 2029
Table 200: show security flow cp-session Output Fields 2031
Table 201: show security flow cp-session destination-port Output Fields 2034
Table 202: show security flow cp-session destination-prefix Output Fields . . 2036
Table 203: show security flow cp-session family Output Fields 2038
Table 204: show security flow cp-session protocol Output Fields 2041
Table 205: show security flow cp-session source-port Output Fields 2043
Table 206: show security flow cp-session source-prefix Output Fields 2045
Table 207: show security flow gate Output Fields 2047
Table 208: show security flow ip-action Output Fields 2052
Table 209: show security flow gate brief node Output Fields 2059
Table 210: show security flow gate destination-port Output Fields 2065
Table 211: show security flow gate destination-prefix Output Fields 2068
Table 212: show security flow gate protocol Output Fields 2072
Table 213: show security flow gate summary node Output Fields 2074
Table 214: show security flow session Output Fields 2081
Table 215: show security flow session brief node Output Fields 2085
Table 216: show security flow session destination-port Output Fields 2089
Table 217: show security flow session destination-prefix Output Fields 2093
Table 218: show security flow session extensive node Output Fields 2097
Table 219: show security flow session family Output Fields 2103
Table 220: show security flow session interface Output Fields 2108
Table 221: show security flow session nat Output Fields 2112
Table 222: show security flow session policy-id Output Fields 2115
Table 223: show security flow session protocol Output Fields 2119
Table 224: show security flow session resource-manager Output Fields 2123
Table 225: show security flow session services-offload Output Fields 2128
Table 226: show security flow session session-identifier Output Fields 2132
Table 227: show security flow session source-port Output Fields 2136
Table 228: show security flow session source-prefix Output Fields 2140
Table 229: show security flow session summary Output Fields 2144
Table 230: show security flow session summary node Output Fields 2146
Table 231: show security flow session summary services-offload Output Fields 2153
Table 232: show security flow session tunnel Output Fields 2156
Table 233: show security flow statistics Output Fields 2162
Table 234: show security flow status Output Fields 2164
Table 235: show security forward-options mirror-filter 2167
Table 236: show security policies Output Fields 2172
Table 237: show security policies hit-count Output Fields 2180

Copyright © 2016, Juniper Networks, Inc.

Table 238: show security resource-manager group Output Fields 2182
Table 239: show security resource-manager resource Output Fields 2185
Table 240: show security resource-manager settings Output Fields 2188
Table 241: show security resource-manager summary Output Fields 2190
Table 242: show security screen ids-option Output Fields 2191
Table 243: show security screen statistics Output Fields 2197
Table 244: show security zones Output Fields 2205
Table 245: show security zones type Output Fields 2208

Part 30 Configuring GPRS Tunnel Protocol v1

Chapter 98 Configuring GTP Message Filtering 2233

Table 246: GTP Messages 2235

Chapter 99 Configuring GTP Information Elements 2247

Table 247: Supported Information Elements 2252
Table 248: Supported Information Elements 2257
Table 249: Supported Information Elements 2259
Table 250: Supported Information Elements 2260

Chapter 100 Configuring NAT for GTP 2265

Table 251: Configuring NAT-PT Details Between IPv4 and IPv6 Endpoints 2272

Part 31 Configuring GPRS Tunnel Protocol v2

Chapter 103 Configuring GTPv2 Message Filtering 2303

Table 252: GTPv2 Messages 2304

Part 32 Configuring Stream Control Transmission Protocol

Chapter 105 Configuring SCTP 2319

Table 253: Common Header Fields 2328
Table 254: Data Chunk Fields 2328

Part 33 Configuration Statements and Operational Commands

Chapter 107 Operational Commands 2409

Table 255: show security gprs gtp counters all Output Fields 2419
Table 256: show security gprs gtp counters path-rate-limit Output Fields 2425
Table 257: show security gprs sctp association 2427
Table 258: show security gprs sctp counters 2429

Part 34 Overview

Chapter 108 Introduction to Interfaces 2441

Table 259: Network Interfaces 2442
Table 260: Configurable Services Interfaces 2443
Table 261: Non-Configurable Services Interfaces 2445
Table 262: Special Interfaces 2445
Table 263: Network Interface Names 2446

Chapter 109 Configuring Interface Logical Properties 2451

Copyright © 2016, Juniper Networks, Inc.

Table 264: Device Status Upon Configuration Change 2461

Chapter 110 Understanding Interface Physical Properties 2475

Table 265: Interface Physical Properties 2475
Table 266: MTU Values for the SRX Series Services Gateways PIMs 2479

Chapter 111 Configuring VLAN Tagging 2483

Table 267: VLAN ID Range by Interface Type Supported on the SRX Series Devices 2485
Table 268: Flexible VLANs 2485

Part 36 Configuring DSL Interfaces

Chapter 113 Configuring VDSL2 Interfaces 2501

Table 269: VDSL2 Annex A and Annex B Features 2504
Table 270: VDLS2 Operating Mode Backward Compatibility with ADSL 2505
Table 271: Supported Profiles on the VDSL2 Interfaces 2506

Part 37 Configuring Ethernet Interfaces

Chapter 114 Performing Initial Configuration on Ethernet Interfaces 2541

Table 272: Collision Backoff Algorithm Rounds 2542

Chapter 118 Configuring Ethernet OAM Link Fault Management 2603

Table 273: Supported Interface Modes 2604

Chapter 119 Configuring Power over Ethernet 2611

Table 274: PoE Specifications for the SRX320 and SRX340 Devices 2611
Table 275: SRX Series Devices PoE Specifications 2612

Part 39 Configuration Statements and Operational Commands

Chapter 125 Operational Commands 2759

Table 276: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary 2771
Table 277: show chassis fpc Output Fields 2773
Table 278: show chassis hardware Output Fields 2779
Table 279: show interfaces Output Fields 2806
Table 280: show ethernet-switching table Output Fields 2808
Table 281: show igmp-snooping route Output Fields 2813
Table 282: show interfaces Output Fields 2818
Table 283: show interfaces diagnostics optics Output Fields 2846
Table 284: show interfaces flow-statistics Output Fields 2851
Table 285: Flow Error Statistics (Packet Drop Statistics for the Flow Module) 2852
Table 286: show interfaces queue Output Fields 2856
Table 287: show ipv6 neighbors Output Fields 2861
Table 288: show lacp interfaces Output Fields 2863
Table 289: show lacp statistics interfaces Output Fields 2867
Table 290: show oam ethernet link-fault-management Output Fields 2868
Table 291: show poe controller Output Fields 2873
Table 292: show pppoe interfaces Output Fields 2874

Copyright © 2016, Juniper Networks, Inc.

Table 293: show pppoe statistics Output Fields 2878
Table 294: show poe telemetries interface Output Fields 2880

Part 41 Configuring Layer 2 Bridging and Transparent Mode

Chapter 127 **Configuring Bridging and Transparent Mode 2897**

Table 295: Security Features Supported in Transparent Mode 2898
Table 296: MAC Addresses Default Limits 2900
Table 297: Enhanced Layer 2 Configuration Statement Changes 2902
Table 298: Enhanced Layer 2 Operational Command Changes 2903

Chapter 128 **Configuring Interfaces 2905**

Table 299: Ethernet Physical Interface and Supported Family Types 2911
Table 300: Security Features Supported in Mixed Mode (Layer 2 and Layer 3) 2913
Table 301: Layer 2 and Layer 3 Parameters 2915

Chapter 134 **Configuring IPv6 Flows 2947**

Table 302: Device Status Upon Configuration Change 2949
Table 303: IPv6 Transparent Mode Configuration for IPv6 Flows 2951

Part 42 Configuring Layer 2 Next generation Switching Mode

Chapter 136 **Configuring Switching Modes 2983**

Table 304: Supported Devices and Ports for Switching Features 2984

Chapter 137 **Configuring VLANs 2989**

Table 305: VLAN Configuration Details 2989

Chapter 138 **Configuring Link Aggregation Control Protocol 2993**

Table 306: LACP (Link Aggregation Control Protocol) Configuration 2994
Table 307: Details of Aggregation 2995
Table 308: Aggregated Ethernet Interface Options 2995
Table 309: Edit VLAN Options 2996

Part 43 Configuration Statements and Operational Commands

Chapter 141 **Operational Commands 3095**

Table 310: show interfaces Output Fields 3101
Table 311: show ethernet-switching table Output Fields 3103
Table 312: show igmp-snooping route Output Fields 3108
Table 313: show igmp-snooping vlans 3110
Table 314: show interfaces Output Fields 3115
Table 315: show oam ethernet connectivity-fault-management adjacencies Output Fields 3143
Table 316: show oam ethernet connectivity-fault-management forwarding-state Output Fields 3144
Table 317: show oam ethernet connectivity-fault-management interfaces Output Fields 3146
Table 318: show oam ethernet connectivity-fault-management mep-database Output Fields 3148

Copyright © 2016, Juniper Networks, Inc.

Table 319: show oam ethernet connectivity-fault-management mep-statistics Output Fields . . . 3152
Table 320: show oam ethernet connectivity-fault-management mip Output Fields . . . 3155
Table 321: show oam ethernet connectivity-fault-management path-database Output Fields . . . 3156
Table 322: show oam ethernet connectivity-fault-management routes Output Fields . . . 3158
Table 323: show oam ethernet link-fault-management Output Fields . . . 3160
Table 324: show security flow gate family Output Fields . . . 3165
Table 325: show security flow ip-action Output Fields . . . 3168
Table 326: show security flow session family Output Fields . . . 3175
Table 327: show security flow statistics Output Fields . . . 3180
Table 328: show security flow status Output Fields . . . 3182
Table 329: show security forward-options secure-wire Output Fields . . . 3185
Table 330: show security policies Output Fields . . . 3188
Table 331: show security zones Output Fields . . . 3195

Part 44 Overview

Chapter 144 Understanding User Logical Systems . . . 3225

Table 332: ls-marketing-dept Logical System Configuration . . . 3229
Table 333: ls-accounting-dept Logical System Configuration . . . 3229

Part 46 Configuring Security Features

Chapter 146 Configuring Master Logical System Security Profiles . . . 3253

Table 334: Security Profiles Used for Reserved Resource Assessments . . . 3257
Table 335: Reserved Resource Allocation Assessment Across Logical Systems . . . 3257

Chapter 147 Configuring Master Logical System Security Features . . . 3267

Table 336: Access Profile Configuration . . . 3269
Table 337: root-logical-system Security Feature Configuration . . . 3272
Table 338: IDP Configuration for the Master Logical System . . . 3281
Table 339: Logical System VPN Tunnel Configuration . . . 3295

Chapter 148 Configuring User Logical System Security Features . . . 3299

Table 340: User Logical System Zone and Address Book Configuration . . . 3301
Table 341: User Logical System Screen Options Configuration . . . 3305
Table 342: User Logical System Security Policies Configuration . . . 3308
Table 343: User Logical System Firewall Authentication Configuration . . . 3314
Table 344: User Logical System Route-Based VPN Configuration . . . 3335

Part 47 Configuring Routing and Interfaces Features

Chapter 150 Configuring User Logical System Routing, Interfaces, and NAT Features . . . 3355

Table 345: User Logical System Static NAT Configuration . . . 3357
Table 346: User Logical System Interface and Routing Instance Configuration . . . 3360

Copyright © 2016, Juniper Networks, Inc.

Part 49 Configuring IPv6 for Logical Systems

Chapter 152 Configuring IPv6 Addresses for Logical Systems 3439

Table 347: User Logical System Zone and Address Book Configuration 3450
Table 348: User Logical System Security Policies Configuration 3453

Part 50 Configuring System Resources Allocation

Chapter 153 System Resources Allocation (Master Administrators Only) 3461

Table 349: CPU Utilization Scenario 1 3464
Table 350: CPU Utilization Scenario 2 3464
Table 351: CPU Utilization Scenario 3 3464
Table 352: Logical Systems, Security Profiles, and Reserved CPU Quotas 3466

Part 52 Configuration Statements and Operational Commands

Chapter 156 Operational Commands 3637

Table 353: show chassis cluster status Output Fields 3644
Table 354: show security application-firewall rule-set Output Fields 3651
Table 355: show security application-firewall rule-set logical-system Output Fields 3655
Table 356: show security application-tracking counters 3657
Table 357: show security dns-cache Output Fields 3660
Table 358: show security firewall-authentication history Output Fields 3662
Table 359: show security firewall-authentication users Output Fields 3664
Table 360: show security flow session Output Fields 3668
Table 361: show security idp logical-system policy-association Output Fields . . 3672
Table 362: show security ike security-associations Output Fields 3674
Table 363: show security ipsec security-associations 3682
Table 364: show security match-policies Output Fields 3694
Table 365: show security nat destination rule Output Fields 3698
Table 366: show security nat destination summary Output Fields 3701
Table 367: show security nat source rule Output Fields 3703
Table 368: show security nat source summary Output Fields 3707
Table 369: show security nat static rule Output Fields 3709
Table 370: show security policies Output Fields 3714
Table 371: show security screen statistics Output Fields 3722
Table 372: show system security-profile Output Fields 3730
Table 373: show security zones Output Fields 3735

Part 58 Configuration Statements and Operational Commands

Chapter 174 Operational Commands 3985

Table 374: show bgp neighbor Output Fields 3986
Table 375: show interfaces flow-statistics Output Fields 3990
Table 376: Flow Error Statistics (Packet Drop Statistics for the Flow Module) . . 3991
Table 377: show security flow status Output Fields 3996
Table 378: show security ipsec security-associations 4000
Table 379: show security ipsec statistics Output Fields 4011

Copyright © 2016, Juniper Networks, Inc.

Part 59 Overview

Chapter 175 Introduction to Multicast 4019

Table 380: Multicast Routing Protocols Compared 4030

Part 60 Managing Group Membership

Chapter 176 Configuring IGMP 4037

Table 381: IGMP Event Messages 4056

Chapter 177 Examples: Configuring MLD 4063

Table 382: MLD Event Messages 4082

Part 61 Configuring Protocol Independent Multicast

Chapter 181 Routing Content to Larger, Sparser Groups with PIM Sparse Mode 4113

Table 383: Tunnel PIC Requirements for IPv4 and IPv6 Multicast 4118
Table 384: Local RP and Auto-RP Message Types 4152
Table 385: PIM Join Filter Match Conditions 4163

Chapter 182 Receiving Content Directly from the Source with SSM 4183

Table 386: ASM and SSM Terminology 4184

Part 62 Configuring Multicast Routing Protocols

Chapter 187 Connecting Routing Domains Using MSDP 4253

Table 387: MSDP Source-Active Message Filter Match Conditions 4256
Table 388: Source-Active Message Flooding Explanation 4265

Part 65 Configuration Statements and Operational Commands

Chapter 197 Operational Commands 4631

Table 389: show amt statistics Output Fields 4658
Table 390: show amt summary Output Fields 4661
Table 391: show amt tunnel Output Fields 4663
Table 392: show dvmrp interfaces Output Fields 4666
Table 393: show dvmrp neighbors Output Fields 4668
Table 394: show dvmrp prefix Output Fields 4670
Table 395: show dvmrp prunes Output Fields 4672
Table 396: show igmp group Output Fields 4674
Table 397: show igmp interface Output Fields 4678
Table 398: show igmp snooping interface Output Fields 4682
Table 399: show igmp snooping membership Output Fields 4685
Table 400: show igmp snooping statistics Output Fields 4689
Table 401: show igmp statistics Output Fields 4692
Table 402: show mld group Output Fields 4695
Table 403: show mld interface Output Fields 4699
Table 404: show mld statistics Output Fields 4702
Table 405: show msdp Output Fields 4705
Table 406: show msdp source Output Fields 4708
Table 407: show msdp source-active Output Fields 4710

Copyright © 2016, Juniper Networks, Inc.

Table 408: show msdp statistics Output Fields 4711
Table 409: show multicast backup-pe-groups Output Fields 4715
Table 410: show multicast flow-map Output Fields 4717
Table 411: show multicast interface Output Fields 4719
Table 412: show multicast pim-to-igmp-proxy Output Fields 4721
Table 413: show multicast pim-to-mld-proxy Output Fields 4723
Table 414: show multicast route Output Fields 4726
Table 415: show multicast rpf Output Fields 4732
Table 416: show multicast scope Output Fields 4735
Table 417: show multicast sessions Output Fields 4737
Table 418: show multicast snooping route Output Fields 4741
Table 419: show multicast snooping statistics Output Fields 4743
Table 420: show multicast usage Output Fields 4747
Table 421: show pgm negative-acknowledgments Output Fields 4749
Table 422: show pgm source-path-messages Output Fields 4751
Table 423: show pgm statistics Output Fields 4752
Table 424: show pim bidirectional df-election Output Fields 4755
Table 425: show pim bidirectional df-election interface Output Fields 4758
Table 426: show pim bootstrap Output Fields 4761
Table 427: show pim interfaces Output Fields 4763
Table 428: show pim join Output Fields 4768
Table 429: show pim neighbors Output Fields 4788
Table 430: show pim rps Output Fields 4792
Table 431: show pim source Output Fields 4799
Table 432: show pim statistics Output Fields 4802
Table 433: show policy Output Fields 4811
Table 434: show route Output Fields 4814
Table 435: show sap listen Output Fields 4834

Chapter 198 Overview 4839

Table 436: Number of Rules on SRX Series Devices 4842
Table 437: Number of Rules and Rule Sets 4843

Chapter 199 Configuring General NAT Options 4845

Table 438: Summary of Key Incoming Table Output Fields 4860
Table 439: Summary of Key Interface NAT Output Fields 4860

Chapter 200 Configuring Source NAT 4863

Table 440: Source NAT Monitoring Page 4889

Chapter 202 Configuring Destination NAT 4923

Table 441: Interfaces, Zones, Server, and IP Address Information 4928
Table 442: Summary of Key Destination NAT Output Fields 4944

Chapter 203 Configuring Static NAT 4947

Table 443: Summary of Key Static NAT Output Fields 4965

Chapter 204 Configuring Persistent NAT and NAT64 4969

Table 444: NAT Feature Compatibility with the Address Persistent Feature . . 4973
Table 445: Interfaces, Zones, Servers, and IP Address Information 4979

Chapter 205 Configuring NAT Hairpinning 4993

Copyright © 2016, Juniper Networks, Inc.

Table 446: Persistent NAT Binding Table 4995

Chapter 208 **Configuring IPv6 Dual-Stack 5031**

Table 447: Softwire Initiator and Softwire Concentrator Capacity 5033

Chapter 210 **Operational Commands 5119**

Table 448: show security nat destination pool Output Fields 5127
Table 449: show security nat destination rule Output Fields 5129
Table 450: show security nat destination rule-application Output Fields 5132
Table 451: show security nat destination summary Output Fields 5134
Table 452: show security nat incoming-table Output Fields 5136
Table 453: show security nat interface-nat-ports Output Fields 5138
Table 454: show security nat resource-usage source-pool Output Fields 5141
Table 455: show security nat source deterministic Output Fields 5145
Table 456: show security nat source paired-address Output Fields 5146
Table 457: show security nat source persistent–nat–table Output Fields 5148
Table 458: show security nat source pool Output Fields 5151
Table 459: show security nat source port-block Output Fields 5155
Table 460: show security nat source rule Output Fields 5157
Table 461: show security nat source rule-application Output Fields 5161
Table 462: show security nat source summary Output Fields 5163
Table 463: show security nat static rule Output Fields 5165

Part 66 Overview

Chapter 211 **Introduction to User Authentication 5173**

Table 464: Comparison of User Firewall Features 5174

Part 67 Configuring Firewall User Authentication

Chapter 213 **Configuring Encrypted Files Using SSL Proxy 5183**

Table 465: Supported SSL Cipher List 5185
Table 466: SSL Proxy Logs 5189
Table 467: SSL Proxy Log Prefixes 5189
Table 468: Trace Levels 5202
Table 469: Supported Flags in Trace 5202

Part 68 Configuring Infranet Authentication

Chapter 225 **Configuring Captive Portal 5299**

Table 470: Redirect URL String Options 5302

Part 69 Configuring Integrated User Firewall

Chapter 226 **Understanding Integrated User Firewalls 5309**

Table 471: Active Directory Authentication Table Support by SRX Series Devices 5313
Table 472: Events Triggering Active Directory Authentication Table Updates . . . 5315

Chapter 227 **Managing Event Logs 5329**

Copyright © 2016, Juniper Networks, Inc.

Table 473: Probe Responses and Associated Active Directory Authentication Table Actions 5333

Part 70 Configuration Statements and Operational Commands

Chapter 229 Operational Commands 5511

Table 474: show network-access requests pending Output Fields 5529
Table 475: show network-access requests statistics Output Fields 5531
Table 476: show network-access securid-node-secret-file Output Fields 5532
Table 477: show security firewall-authentication history Output Fields 5533
Table 478: show security firewall-authentication history address Output Fields 5535
Table 479: show security firewall-authentication history identifier Output Fields 5538
Table 480: show security firewall-authentication users Output Fields 5541
Table 481: show security firewall-authentication users address Output Fields 5543
Table 482: show security firewall-authentication users identifier Output Fields 5546
Table 483: show security policies Output Fields 5550
Table 484: show services unified-access-control counters Output Fields 5559
Table 485: show services unified-access-control roles Output Fields 5563
Table 486: show services user-identification active-directory-access active-directory-authentication-table all extensive Output Fields 5566
Table 487: show services user-identification active-directory-access domain-controller Output Fields 5569
Table 488: show services user-identification active-directory-access statistics ip-user-mapping Output Fields 5572
Table 489: show services user-identification active-directory-access statistics ip-user-probe Output Fields 5573
Table 490: show services user-identification active-directory-access statistics user-group-mapping Output Fields 5573
Table 491: show services user-identification active-directory-access user-group-mapping group Output Fields 5575
Table 492: show services user-identification active-directory-access user-group-mapping status Output Fields 5576
Table 493: show services user-identification active-directory-access user-group-mapping user Output Fields 5576

Part 71 Overview

Chapter 231 Managing UTM Licensing 5587

Table 494: UTM Feature Subscription Service License Requirements 5587

Part 74 Configuring Full Antivirus Protection and Pattern Updates

Chapter 243 Configuring Application Protocol Scanning 5695

Table 495: Supported Profile-based Settings By Protocol 5695

Copyright © 2016, Juniper Networks, Inc.

Part 78 Configuration Statements and Operational Commands

Chapter 250 Configuration Statements 5819

Table 496: List of categories predefined by Websense 5855

Chapter 251 Operational Commands 6001

Table 497: show configuration smtp 6011
Table 498: show security log Output Fields 6014
Table 499: show security policies Output Fields 6017

Part 79 Overview

Chapter 252 Introduction to IPsec VPNs 6041

Table 500: Comparison Between Policy-Based VPNs and Route-Based VPNs 6042
Table 501: Recommended Configuration for Site-to-Site VPN with Static IP Addresses 6061
Table 502: Recommended Configuration for Site-to-Site or Dialup VPNs with Dynamic IP Addresses 6062

Chapter 253 Understanding VPN Tunnel Management 6067

Table 503: Load Balancing Across SPUs 6068

Part 80 Configuring Route-Based IPsec VPNs

Chapter 254 Configuring Route-Based VPNs 6073

Table 504: Interface, Static Route, Security Zone, and Address Book Information 6076
Table 505: IKE Phase 1 Configuration Parameters 6076
Table 506: IPsec Phase 2 Configuration Parameters 6077
Table 507: Security Policy Configuration Parameters 6077
Table 508: TCP-MSS Configuration Parameters 6077

Chapter 255 Configuring Hub-and-Spoke VPNs 6093

Table 509: Interface, Security Zone, and Address Book Information 6095
Table 510: IKE Phase 1 Configuration Parameters 6097
Table 511: IPsec Phase 2 Configuration Parameters 6098
Table 512: Security Policy Configuration Parameters 6099
Table 513: TCP-MSS Configuration Parameters 6100

Chapter 256 Configuring VPNs for IKEv2 6127

Table 514: IKEv2 Configuration Attributes 6129
Table 515: Interface, Static Route, Security Zone, and Address Book Information 6130
Table 516: IKE Phase 1 Configuration Parameters 6131
Table 517: IPsec Phase 2 Configuration Parameters 6131
Table 518: Security Policy Configuration Parameters 6132
Table 519: TCP-MSS Configuration Parameters 6132
Table 520: Phase 1 and Phase 2 Options for the SRX Series 6151

Chapter 258 Configuring Dual Stack Tunnels over an External Interface 6181

Table 521: Phase 1 Options for Dual-Stack Tunnel Configuration 6185

Copyright © 2016, Juniper Networks, Inc.

Table 522: Phase 2 Options for Dual-Stack Tunnel Configuration 6186

Chapter 259 **Configuring Traffic Selectors in Route-Based VPNs 6195**

Table 523: Traffic Selector Configurations 6199
Table 524: Phase 1 Options for Traffic Selector Configurations 6199
Table 525: Phase 2 Options for Traffic Selector Configurations 6200

Part 81 Configuring Policy-Based IPsec VPNs

Chapter 260 **Configuring Policy-Based VPNs 6221**

Table 526: Interface, Security Zone, and Address Book Information 6224
Table 527: IKE Phase 1 Configuration Parameters 6224
Table 528: IPsec Phase 2 Configuration Parameters 6225
Table 529: Security Policy Configuration Parameters 6225
Table 530: TCP-MSS Configuration Parameters 6226

Part 82 Configuring VPNs with NAT-T

Chapter 261 **Configuring Route-Based and Policy-Based VPNs with NAT-T 6243**

Table 531: Interface, Routing Options, Zones, and Security Policies for the Initiator 6247
Table 532: IKE Phase 1 Configuration Parameters for the Initiator 6247
Table 533: IPsec Phase 2 Configuration Parameters for the Initiator 6248
Table 534: Interface, Routing Options, Zones, and Security Policies for the Responder 6248
Table 535: IKE Phase 1 Configuration Parameters for the Responder 6249
Table 536: IPsec Phase 2 Configuration Parameters for the Responder 6249
Table 537: Interface, Routing Options, and Security Zones for the Initiator 6274
Table 538: IKE Phase 1 Configuration Parameters for the Initiator 6274
Table 539: IPsec Phase 2 Configuration Parameters for the Initiator 6275
Table 540: Security Policy Configuration Parameters for the Initiator 6275
Table 541: Interface, Routing Options, and Security Zones for the Responder . . 6275
Table 542: IKE Phase 1 Configuration Parameters for the Responder 6276
Table 543: IPsec Phase 2 Configuration Parameters for the Responder 6276
Table 544: Security Policy Configuration Parameters for the Responder 6277

Part 84 Configuring IPv6 IPsec VPNs

Chapter 263 **Configuring IPv6 IPsec VPNs 6331**

Table 545: IPv6 Address Support in VPN Features 6331
Table 546: ISAKMP ID Types and Their Values 6335
Table 547: Support for IPv4 Options or IPv6 Extension Headers 6338
Table 548: IPv6 Header Construction for IPv6-in-IPv6 and IPv4-in-IPv6 Tunnel Modes 6339
Table 549: IPv4 Header Construction for IPv6-in-IPv4 and IPv4-in-IPv4 Tunnel Modes 6340

Part 85 Configuring Public Key Infrastructure

Chapter 264 **Managing Digital Certificates with PKI 6347**

Copyright © 2016, Juniper Networks, Inc.

Table 550: Comparison of CMPv2 and SCEP Certificate Enrollment 6352

Chapter 269 **Managing Certificate Revocation 6409**

Table 551: Phase 1 Options for OCSP Configuration Example 6412
Table 552: Phase 2 Options for OCSP Configuration Example 6413

Part 86 Configuring AutoVPN

Chapter 272 **Configuring AutoVPN on Hub-and-Spoke Devices 6451**

Table 553: Comparison Between AutoVPN Point-to-Point and Point-to-Multipoint Secure Tunnel Modes 6452
Table 554: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke Configurations 6458
Table 555: AutoVPN Configuration for Hub and All Spokes 6458
Table 556: Comparison Between the Spoke Configurations 6459
Table 557: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke Basic OSPF Configurations 6483
Table 558: AutoVPN Basic OSPF Configuration for Hub and All Spokes 6484
Table 559: Comparison Between the Basic OSPF Spoke Configurations 6484
Table 560: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke iBGP ECMP Configurations 6507
Table 561: AutoVPN iBGP ECMP Configuration for Hub and Spoke 1 6508
Table 562: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke iBGP Active-Backup Tunnel Configurations 6532
Table 563: AutoVPN IBGP Active-Backup Tunnel Configuration for Hub and Spoke 1 6533

Chapter 273 **Configuring Auto Discovery VPNs 6559**

Table 564: Shortcut Parameters 6561
Table 565: Phase 1 and Phase 2 Options for AutoVPN Hub and Spokes for ADVPN Example 6567
Table 566: IKE Gateway Configuration for ADVPN Example 6568

Chapter 274 **Configuring AutoVPN and Traffic Selectors 6607**

Table 567: Phase 1 and Phase 2 Options for AutoVPN Hubs and Spokes with Traffic Selectors 6609
Table 568: Phase 1 and Phase 2 Options for Georedundant AutoVPN Hubs . . 6625

Part 87 Configuring Group VPNs

Chapter 277 **Configuring Group VPNv2 Server Clusters 6699**

Table 569: Effects of Configuration Changes on Group VPNv2 Servers 6706
Table 570: Effects of Group VPNv2 Server Cluster Configuration Changes . . . 6708

Part 89 Troubleshooting

Chapter 280 **Tunnel Events 6797**

Table 571: IPsec VPN Tunnel Events 6797

Part 90 Configuration Statements and Operational Commands

Chapter 282 **Operational Commands 6973**

Copyright © 2016, Juniper Networks, Inc.

Table 572: show security group-vpn member ike security-associations Output Fields 7010
Table 573: show security group-vpn member ipsec inactive-tunnels Output Fields 7013
Table 574: show security group-vpn member ipsec security-associations 7016
Table 575: show security group-vpn member ipsec statistics Output Fields . . 7020
Table 576: show security group-vpn member kek security-associations 7022
Table 577: show security group-vpn member policy Output Fields 7025
Table 578: show security group-vpn server ike security-associations Output Fields 7027
Table 579: show security group-vpn server ipsec security-associations 7031
Table 580: show security group-vpn server kek security-associations Output Fields 7034
Table 581: show security group—vpn server registered-members Output Fields 7037
Table 582: show security group-vpn server server-cluster Output Fields 7039
Table 583: show security group-vpn server statistics Output Fields 7042
Table 584: show security ike security-associations Output Fields 7046
Table 585: show security ike tunnel-map Output Fields 7053
Table 586: show security ipsec control-plane-security-associations Output Fields 7056
Table 587: show security ipsec inactive-tunnels Output Fields 7059
Table 588: show security ipsec next-hop-tunnels Output Fields 7061
Table 589: show security ipsec security-associations 7063
Table 590: show security ipsec statistics Output Fields 7074
Table 591: show security ipsec traffic-selector Output Fields 7078
Table 592: show security pki ca-certificate Output Fields 7080
Table 593: show security pki certificate-request Output Fields 7084
Table 594: show security pki crl Output Fields 7086
Table 595: show security pki local-certificate Output Fields 7089

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

About the Documentation

- Documentation and Release Notes on page cxliii
- Using the Examples in This Manual on page cxliii
- Documentation Conventions on page cxlv
- Documentation Feedback on page cxlvii
- Requesting Technical Support on page cxlvii

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Copyright © 2016, Juniper Networks, Inc.

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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


Copyright © 2016, Juniper Networks, Inc.

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

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

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 defines notice icons used in this guide.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`

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

Convention	Description	Examples
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** ***metric*****>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Copyright © 2016, Juniper Networks, Inc.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/

Copyright © 2016, Juniper Networks, Inc.

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

Copyright © 2016, Juniper Networks, Inc.

Application Layer Gateways Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

PART 1

Overview

- Introduction to ALGs on page 3

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 1

Introduction to ALGs

- ALG Overview on page 3
- Understanding Custom ALG Services on page 4
- Understanding IPv6 ALG Support for ICMP on page 5

ALG Overview

An Application Layer Gateway (ALG) is a software component that is designed to manage specific protocols such as Session Initiation Protocol (SIP) or FTP on Juniper Networks devices running Junos OS. The ALG module is responsible for Application-Layer aware packet processing.

ALG functionality can be triggered either by a service or application configured in the security policy:

- A service is an object that identifies an application protocol using Layer 4 information (such as standard and accepted TCP and UDP port numbers) for an application service (such as Telnet, FTP, SMTP, and HTTP).
- An application specifies the Layer 7 application that maps to a Layer 4 service.

A predefined service already has a mapping to a Layer 7 application. However, for custom services, you must link the service to an application explicitly, especially if you want the policy to apply an ALG.

ALGs for packets destined to well-known ports are triggered by service type. The ALG intercepts and analyzes the specified traffic, allocates resources, and defines dynamic policies to permit the traffic to pass securely through the device:

1. When a packet arrives at the device, the flow module forwards the packet according to the security rule set in the policy.
2. If a policy is found to permit the packet, the associated service type or application type is assigned and a session is created for this type of traffic.
3. If a session is found for the packet, no policy rule match is needed. The ALG module is triggered if that particular service or application type requires the supported ALG processing.

Copyright © 2016, Juniper Networks, Inc.

The ALG also inspects the packet for embedded IP address and port information in the packet payload, and performs Network Address Translation (NAT) processing if necessary. The ALG also opens a gate for the IP address and port number to permit data exchange for the session. The control session and data session can be coupled to have the same timeout value, or they can be independent.

ALGs are supported on chassis clusters.

The ALG message buffer optimization is enhanced to reduce high memory consumption. A message buffer is allocated only when the packet is ready to process. The buffer is freed after the packet completes ALG handling, including modifying the payload, performing NAT, opening a pinhole for a new connection between a client and a server, and transferring data between a client and a server located on opposite sides of a Juniper Networks device.

The maximum size of the jbuf is 9 Kb. If the message buffer size is more than 9 Kb, the entire message cannot be transferred to the ALG packet handler. This causes subsequent packets in the session to bypass ALG handling, resulting in a transaction failure.

Related Documentation

- Understanding Custom ALG Services on page 4
- Understanding VoIP ALG Types on page 141
- Understanding Data ALG Types on page 11
- Understanding H.323 ALG on page 145
- Understanding the SIP ALG on page 241
- Understanding SCCP ALGs on page 217
- Understanding the MGCP ALG on page 183
- Understanding RPC ALGs on page 61

Understanding Custom ALG Services

By default, ALGs are bound to predefined services. For example, the FTP ALG is bound to junos-ftp, the RTSP ALG is bound to junos-rtsp, and so on.

A predefined service already has a mapping to a Layer 7 application. However, for custom services, you must link the service to an application explicitly, especially if you want the policy to apply an ALG.

When you apply predefined services to your policy, traffic matching the service will be sent to its corresponding ALG for further processing. However, under some circumstances, the customer needs to define custom services in order to achieve the following:

- Utilize the ALG handler to process special traffic, with customer-specified protocols, destination ports and so on.
- Permit traffic but bypass ALG processing, when traffic matches predefined services that bind with ALG.
- Add more applications to the current ALG's application set.

Copyright © 2016, Juniper Networks, Inc.

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

The three usages of custom services are illustrated below, considering MS-RPC ALG as an example:

- **Utilize the ALG handler to process special traffic**:

  ```
  [edit]
  user@host# set applications application customer-msrpc application-protocol ms-rpc
  user@host# set applications application customer-msrpc protocol tcp
  user@host# set applications application customer-msrpc destination-port 6000
  ```

  Traffic with TCP destination port 6000 will be sent to MS-RPC ALG for further processing.

- **Permit traffic but bypass ALG processing**:

  ```
  [edit]
  user@host# set applications application customer-ignore application-protocol ignore
  user@host# set applications application customer-ignore protocol tcp
  user@host# set applications application customer-ignore destination-port 135
  ```

  MS-RPC ALG will be ignored by traffic with TCP destination port 135.

- **Add more applications to the current ALG's application set**—To add applications such as MS-RPC or SUNRPC services, which are not predefined on SRX Series devices:

  ```
  [edit]
  user@host# set applications application customer-msrpc term t1 protocol tcp
  user@host# set applications application customer-msrpc term t1 uuid
    e3514235-4b06-11d1-ab04-00c04fc2dcd2
  ```

  MS-RPC data traffic with TCP, uuid e3514235-4b06-11d1-ab04-00c04fc2dcd2, will be permitted, when custom-msrpc is applied to the policy along with other predefined junos-ms-rpc** applications.

Related Documentation

- Understanding VoIP ALG Types on page 141
- ALG Overview on page 3
- Understanding Microsoft RPC Services on page 71
- Understanding RPC ALGs on page 61

Understanding IPv6 ALG Support for ICMP

The Internet Control Message Protocol (ICMP) Application Layer Gateway (ALG) is one of the ALG's that handle ICMP traffic.

IPv6 nodes use the ICMPv6 protocol to report errors encountered in processing packets and to perform other Internet-layer functions such as diagnostics. ICMPv6 is an integral part of IPv6 and must be fully implemented by every IPv6 node; therefore the ALG layer is always enabled for ICMPv6.

Copyright © 2016, Juniper Networks, Inc.

ICMP Error Messages

ICMPv6 messages are grouped into two classes:

- ICMPv6 error messages
  - Destination unreachable
  - Packet too big
  - Time exceeded
  - Parameter problem
- ICMPv6 informational (or ping) messages
  - Echo request
  - Echo reply

The ICMP ALG monitors all these messages, and then does the following :

- Closes the session
- Modifies the payload

The ICMP ALG closes a session if it meets the following conditions:

- Receives echo reply message.
- Receives a destination unreachable error message and has not received any replies yet.

NOTE: The ICMP ALG checks if the session has received any replies from destination node. If it has received any reply , the destination should be reachable and the ICMP error message is not credible, therefore it does not close the session. This is to avoid hackers from sniffing the TCP/UDP packet and forging an ICMP destination unreachable packet to kill the session.

ICMP ALG Functionality

ICMP ALG behaves differently in various modes.

ICMP ALG functionality in NAT mode:

1. Close the session.
2. Modify the identifier, the sequence number or both of the echo request.
3. Resume the original identifier and sequence number for the echo reply.
4. NAT translates the embedded IPv6 packet for theICMPv6 error message.

Copyright © 2016, Juniper Networks, Inc.

ICMP ALG functionality in NAT-PT support mode:

1. Close the session.
2. Translate the ICMPv4 ping message to the ICMPv6 ping message.
3. Translate the ICMPv6 ping message to the ICMPv4 ping message.
4. Translate the ICMPv4 error message to the ICMPv6 error message and translate its embedded IPv4 packet to an IPv6 packet.
5. Translate the ICMPv6 error message to the ICMPv4 error message and translate its embedded IPv6 packet to an IPv4 packet .

Related Documentation

- Understanding How SRX Series Devices Handle Packet Fragmentation for IPv6 Flows on page 1769
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

PART 2

Configuring Data ALGs

- Data ALG Types on page 11
- Configuring the DNS ALG on page 13
- Configuring the FTP ALG on page 25
- Configuring the IKE and ESP ALG on page 33
- Configuring the PPTP ALG on page 45
- Configuring the RPC ALG on page 61
- Configuring the RSH ALG on page 75
- Configuring the RTSP ALG on page 89
- Configuring the SQLNET ALG on page 103
- Configuring the TALK ALG on page 117
- Configuring the TFTP ALG on page 131

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 2

Data ALG Types

- Understanding Data ALG Types on page 11

Understanding Data ALG Types

Junos OS supports the following data ALGs:

- DNS—Provides an ALG for the Domain Name System. The DNS ALG monitors DNS query and reply packets and closes session if the DNS flag indicates the packet is a reply message.
- DDNS—Dynamic DNS (DDNS) is an addition to the DNS standard. DDNS updates a DNS server with new or changed records for IP addresses without the need for human intervention. Unlike DNS that only works with static IP addresses, DDNS is also designed to support dynamic IP addresses, such as those assigned by a DHCP server. DDNS is a good option for home networks, which often receive dynamic public IP addresses from their Internet provider that occasionally changes.
- FTP—Provides an ALG for the File Transfer Protocol (FTP).The FTP ALG monitors PORT, PASV, and 227 commands. It performs NAT on the IP, port, or both in the message and gate opening on the device as necessary.
- IKE and ESP ALG—Monitors IKE traffic between the client and the server and permits only one IKE Phase 2 message exchange between any given client/server pair, not just one exchange between any client and any server.

  Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) traffic is exchanged between the clients and the server. However, if the clients do not support NAT-Traversal (NAT-T) and if the device assigns the same NAT-generated IP address to two or more clients, the device will be unable to distinguish and route return traffic properly.

NOTE: If the user wants to support both NAT-T-capable and non-NAT-T-capable clients, then some additional configurations are required. If there are NAT-T capable clients, the user must enable the source NAT address persistence.

Copyright © 2016, Juniper Networks, Inc.

- TFTP—Provides an ALG for the Trivial File Transfer Protocol (TFTP). The TFTP ALG processes TFTP packets that initiate the request and opens a gate to allow return packets from the reverse direction to the port that sends the request.
- PPTP—Provides an ALG for the Point-to-Point Tunneling Protocol (PPTP). The PPTP is a Layer 2 protocol that tunnels PPP data across TCP/IP networks. The PPTP client is freely available on Windows systems and it is also popularly applied on Linux systems, and is widely deployed for building Virtual Private Networks (VPNs).
- MS-RPC—Provides an ALG for the Microsoft Remote Procedure Call.
- SUNRPC—Provides an ALG for the SUN Remote Procedure Call.
- RSH—Provides an ALG for the Remote Shell (RSH). The RSH ALG handles TCP packets destined for port 514 and processes the RSH port command. The RSH ALG performs NAT on the port in the port command and opens gates as necessary.
- SQL—Provides an ALG for the Structured Query Language (SQL). The SQLNET ALG processes SQL TNS response frame from the server side. It parses the packet and looks for the (HOST=ipaddress), (PORT=port) pattern and performs NAT and gate opening on the client side for the TCP data channel.
- TALK—Provides an ALG for the TALK Protocol. The TALK protocol uses UDP port 517 and port 518 for control channel connections. The talk program consists of a server and a client. The server handles client notifications and helps to establish talk sessions. There are two types of talk servers: ntalk and talkd. The TALK ALG processes packets of both ntalk and talkd formats. It also performs NAT and gate opening as necessary.
- RTSP—Provides an ALG for the Real Time Streaming Protocol (RTSP). RTSP is a standard protocol for streaming media applications. It controls the delivery of data with real-time properties such as audio and video.

For information about enabling and configuring each of these ALGs through J-Web, select the **Configure>Security>ALG** page in the J-Web user interface and click **Help**.

Related Documentation

- ALG Overview on page 3
- Understanding Custom ALG Services on page 4

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 3

Configuring the DNS ALG

- DNS ALG Overview on page 13
- Understanding the IPv6 DNS ALG for Routing, NAT, and NAT-PT on page 14
- Example: Configuring the DNS ALG on page 16
- Understanding DNS and DDNS Doctoring on page 19
- Disabling DNS and DDNS Doctoring on page 23

DNS ALG Overview

The DNS Application Layer Gateway (ALG) service provides an application-level gateway for use with DNS clients. The DNS ALG service allows a client to access multiple DNS servers in different networks and provides routing to and from those servers. It also supports flexible address translation of the DNS query and response packets. These functions allow the DNS client to query many different domains from a single DNS server instance on the client side of the network.

The DNS server listens through UDP port 53 for incoming queries from DNS resolvers. A resolver communicates with DNS servers by sending DNS queries and handling DNS responses.

NOTE: **The default port for DNS ALG is port 53.**

The DNS ALG performs the following functions:

- Monitors DNS query and reply packets and closes the session when the DNS reply is received
- Performs DNS doctoring
- Performs the IPv4 and IPv6 address transformations
- Modifies the DNS payload in NAT mode

Dynamic DNS (DDNS) support is now available in addition to the DNS standard. The Domain Name System (DNS) was originally designed to support queries of a static configured database and the data was expected to change.

Copyright © 2016, Juniper Networks, Inc.

The main difference between DNS and DDNS is in the message format of the header section and the update message.

DDNS messages are processed differently when compared to DNS messages. Message parsing is rewritten for DDNS. DDNS does NAT and NAT-PT in the query part of the message and DNS does NAT and NAT-PT in the response part of the message.

NOTE: The DNS ALG supports all the new formats and new functionality.

Related Documentation

- *DNS Overview*
- *DNSSEC Overview*

Understanding the IPv6 DNS ALG for Routing, NAT, and NAT-PT

Domain Name System (DNS) is the part of the ALG that handles DNS traffic, monitors DNS query and reply packets, and closes the session if the DNS flag indicates the packet is a reply message.

The DNS ALG supports IPv4 in route mode for Junos OS Release 10.0 and earlier releases. In Junos OS Release 10.4, this feature implements IPv6 support on the DNS ALG for routing, Network Address Translation (NAT), and Network Address Translation-Protocol Translation (NAT-PT).

When the DNS ALG receives a DNS query from the DNS client, a security check is done on the DNS packet. When the DNS ALG receives a DNS reply from the DNS server, a similar security check is done, and then the session for the DNS traffic closes.

IPv6 DNS ALG Traffic in NAT mode

IPv6 NAT provides address translation between IPv4 and IPv6 addressed network devices. It also provides address translation between IPv6 hosts. NAT between IPv6 hosts is done in a similar manner and for similar purposes as IPv4 NAT.

When the DNS traffic works in NAT mode, the DNS ALG translates the public address in a DNS reply to a private address when the DNS client is on private network, and similarly translates a private address to a public address when the DNS client is on a public network.

In Junos OS Release 10.4 IPv6 NAT supports:

- Source NAT translations
- Destination NAT mappings
- Static NAT mappings

NOTE: The IPv6 DNS ALG NAT supports only static NAT mapping.

Copyright © 2016, Juniper Networks, Inc.

IPv6 DNS ALG Traffic in NAT-PT mode

IPv6 NAT-PT provides address allocation and protocol translation between IPv4 and IPv6 addressed network devices. The translation process is based on the Stateless IP/ICMP Translation (SIIT) method; however, the state and the context of each communication is retained during the session lifetime. IPv6 NAT-PT supports Internet Control Message Protocol (ICMP), Transmission Control Protocol (TCP), and User Datagram Protocol (UDP) packets.

IPv6 NAT-PT supports the following types of NAT-PT:

- Traditional NAT-PT
- Bidirectional NAT-PT

A DNS-based mechanism dynamically maps IPv6 addresses to IPv4-only servers. NAT-PT uses the DNS ALG to transparently do the translations.

For example, a company using an internal IPv6 network needs to be able to communicate with external IPv4 servers that do not have IPv6 addresses.

To support the dynamic address binding, a DNS should be used for name resolution. The IPv4 host looks up the name of the IPv6 node in its local configured IPv4 DNS server, which then passes the query to the IPv6 DNS server through an SRX Series device using NAT-PT.

When DNS traffic works in NAT-PT mode, the DNS ALG translates the IP address in a DNS reply packet between the IPv4 address and the IPv6 address when the DNS client is in an IPv6 network and the server is in an IPv4 network, and vice versa.

NOTE: **In NAT-PT mode, only IPV4 to IPV6 addresses translation is supported in the DNS ALG. To support NAT-PT mode in a DNS ALG, the NAT module should support NAT-PT.**

When the DNS ALG receives a DNS query from the DNS client, the DNS ALG performs the following security and sanity checks on the DNS packets:

- Enforces the maximum DNS message length (the default is 512 bytes and the maximum length is 8KB)
- Enforces a domain-name length of 255 bytes and a label length of 63 bytes
- Verifies the integrity of the domain-name referred to by the pointer if compression pointers are encountered in the DNS message
- Checks to see if a compression pointer loop exists

Similar sanity checks are done when the DNS ALG receives a DNS reply from the DNS Server, after which the session for this DNS traffic gets closed.

Related Documentation

- DNS ALG Overview on page 13

Copyright © 2016, Juniper Networks, Inc.

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013

Example: Configuring the DNS ALG

This example shows how to configure the DNS ALG to pass through DNS traffic with a static NAT pool on Juniper Networks devices.

- Requirements on page 16
- Overview on page 16
- Configuration on page 16
- Verification on page 18

Requirements

Before you begin:

- Configure static NAT pool for all IP address.
- Understand the concepts behind ALG for DNS. See "DNS ALG Overview" on page 13.

Overview

In this example, the ALG for DNS is configured to monitor and allow DNS traffic to be exchanged between the clients and the server located on opposite sides of a Juniper Networks device.

This example shows how to configure a static NAT pool and rule set, and associate the DNS ALG to a policy.

Configuration

- Configuring a NAT Static Pool and Rule Set on page 16
- Configuring and Printing the DNS Trace on page 18

Configuring a NAT Static Pool and Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone untrust
set security nat static rule-set rs1 rule r1 match destination-address 1.1.1.1/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 2.2.2.2/32
set security policies from-zone untrust to-zone trust policy u2t match source-address
    any
set security policies from-zone untrust to-zone trust policy u2t match destination-address
    any
set security policies from-zone untrust to-zone trust policy u2t match application
    junos-dns-udp
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone untrust to-zone trust policy u2t then permit
```

Step-by-Step Procedure

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

To configure a static NAT pool:

1. Create a NAT static rule set.

```
[edit ]
user@host# set security nat static rule-set rs1 from zone untrust
user@host# set security nat static rule-set rs1 rule r1 match destination-address
  1.1.1.1/32
user@host# set security nat static rule-set rs1 rule r1 then static-nat prefix 2.2.2.2/32
```

2. Associate the NAT Traversal (NAT-T) application using a policy.

```
[edit]
user@host# set security policies from-zone untrust to-zone trust policy u2t match
  source-address any
user@host# set security policies from-zone untrust to-zone trust policy u2t match
  destination-address any
user@host# set security policies from-zone untrust to-zone trust policy u2t match
  application junos-dns-udp
user@host# set security policies from-zone untrust to-zone trust policy u2t then
  permit
```

Results

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

```
user@host# show security nat
static {
  rule-set rs1 {
    from zone untrust;
    rule r1 {
      match {
        destination-address 1.1.1.1/32;
      }
      then {
        static-nat {
          prefix {
            2.2.2.2/32;
          }
        }
      }
    }
    }
  }
```

```
[edit]
user@host# show security policies
from-zone untrust to-zone trust {
  policy u2t {
    match {
```

Copyright © 2016, Juniper Networks, Inc.

```
        source-address any;
        destination-address any;
        application [ junos-dns-udp];
      }
      then {
        permit;
      }
    }
  }
  default-policy {
    permit-all;
  }
```

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

Configuring and Printing the DNS Trace

Purpose Print the DNS trace file.

Action From configuration mode, enter the following command.

```
set security alg traceoptions file alglog
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
set security alg dns traceoptions flag all
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying DNS ALG Custom Applications on page 18
- Verifying DNS ALG on page 18

Verifying DNS ALG Custom Applications

Purpose Verify that the custom applications to support the DNS ALG are enabled.

Action From operational mode, enter the **show security policies** command.

```
user@host> show security policies
```

```
Default policy: permit-all
From zone: untrust, To zone: trust
Policy: u2t, State: enabled, Index: 6, Scope Policy: 0, Sequence number: 1
Source addresses: any
Destination addresses: any
Applications: junos-dns-udp
Action: permit
```

Meaning The sample output shows that the custom applications to support the DNS ALG is enabled.

Verifying DNS ALG

Purpose Verify that DNS ALG is enabled.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status

DNS      : Enabled
FTP      : Enabled
H323     : Enabled
```

Meaning The output shows the DNS ALG status as follows:

- Enabled—Shows the DNS ALG is enabled.
- Disabled—Shows the DNS ALG is disabled.

Related Documentation

- DNS ALG Overview on page 13

Understanding DNS and DDNS Doctoring

Junos OS for SRX Series devices provides Domain Name System (DNS) support. The DNS ALG monitors DNS query and reply packets and closes the session if the DNS flag indicates that the packet is a reply message. To configure the DNS ALG, use the **edit security alg dns** statement at the **[edit security alg]** hierarchy level.

DNS provides name-to-address mapping within a routing class, whereas Network Address Translation (NAT) attempts to provide transparent routing between hosts in disparate address realms of the same routing class. As a result, NAT can cause some DNS problems the DNG ALG must handle through a process called DNS doctoring.

The same doctoring feature applies to the dynamic domain name system (DDNS). For DDNS in NAT mode, you also can do the IP translation in the DDNS update.

To resolve the problems introduced by NAT, DNS and DDNS ALG functionality has been extended to support static NAT, allowing the problems to be resolved through DNS doctoring.

NOTE: The DNS ALG must be enabled on the devices to perform DNS doctoring. With the DNS ALG enabled on SRX5600 and SRX5800 devices, DNS doctoring is enabled by default.

Copyright © 2016, Juniper Networks, Inc.

The restoring and doctoring process is performed in two parts:

- **Packet sanity check**

DNS packet format

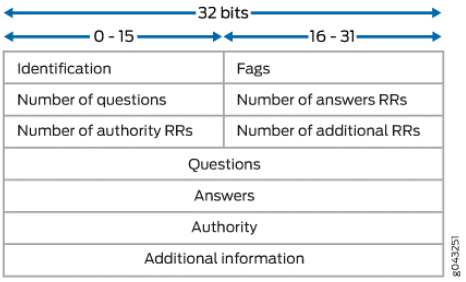

For the DNS packet, the DNS ALG check fields are questions, answers, authority, and additional information. The DNS ALG drops the packet if the number of questions is more than 1, the domain name is more than 255 bytes, or the label length is more than 63 bytes.

DDNS packet format

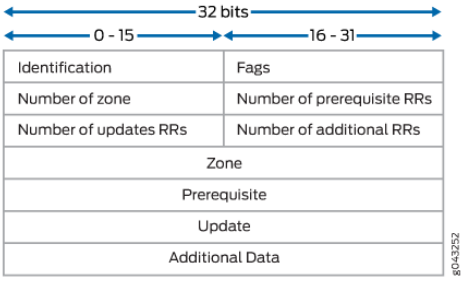

For the DDNS packet, the DNS ALG check fields are zone, prerequisite, update, and additional data. The DNS ALG drops the packet if the number of zones is more than 1, the domain name is more than 255 bytes, or the label length is more than 63 bytes.

For both DNS and DDNS, the DNS ALG drops the packet that does not comply with the standards.

Copyright © 2016, Juniper Networks, Inc.

- NAT

Figure 1 shows how DNS translates a private address to a public address.

Figure 1: DNS Address Translation (Private to Public)

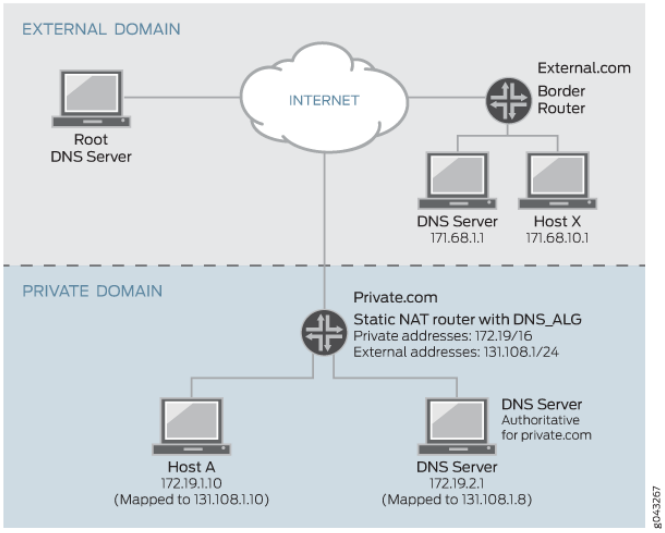

When host X in external.com wants to resolve host A's address through DNS and if the DNS ALG does not support NAT, it takes a private address such as 172.19.1.10, which is invalid to host X. The private address is translated to public address 131.108.1.10 through the DNS ALG.

Copyright © 2016, Juniper Networks, Inc.

Figure 2 shows how DNS translates a public address to a private address.

Figure 2: DNS Address Translation (Public to Private)

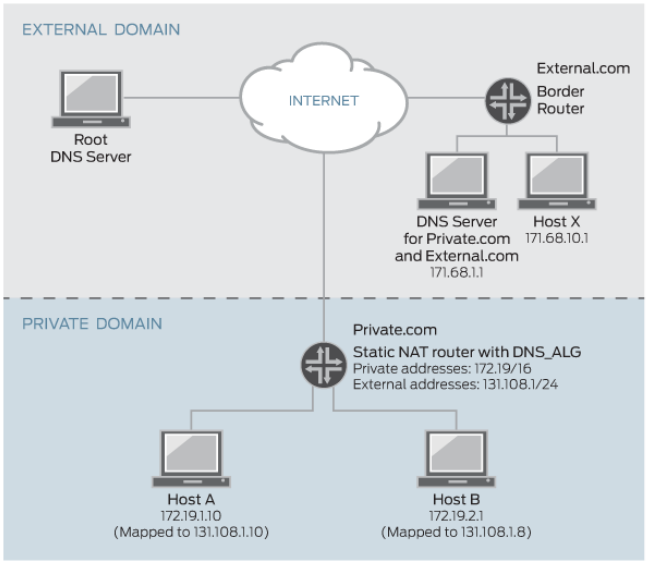

When host A in private.com wants to resolve host B's address through DNS and if the DNS ALG does not support NAT, it takes a public address from the DNS server in external.com, such as 131.108.1.8. If Host A sends traffic to host B with public address 131.108.1.8, which is invalid to host B in the private domain. Hence, the public address in the DNS query A-record is translated to private address 172.19.2.1 through the DNS ALG.

NOTE: The DNS ALG can translate the first 32 A-records in a single DNS reply. A-records after the first 32 records are not handled. Also note that the DNS ALG supports only IPv4 addresses and does not support VPN tunnels.

Related Documentation

- *DNS Overview*
- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014
- Disabling DNS and DDNS Doctoring on page 23

Copyright © 2016, Juniper Networks, Inc.

Disabling DNS and DDNS Doctoring

The DNS and DDNS doctoring feature is enabled by default. You can disable DNS and DDNS doctoring with the CLI.

To disable DNS and DDNS doctoring:

1. Disable all the doctoring features by specifying the **none** configuration option.

   This command disables all the doctoring features.

   user@host# set security alg dns doctoring none

2. Disable the NAT feature and retain the sanity-check feature by specifying the **sanity-check** configuration option.

   This option disables the NAT feature and retains the sanity-check feature.

   user@host# set security alg dns doctoring sanity-check

3. If you are finished configuring the device, commit the configuration.

4. From configuration mode in the CLI, enter the **show security alg dns doctoring** command to verify the configuration.

Related Documentation

- *DNS Overview*
- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 4

Configuring the FTP ALG

- FTP ALG Overview on page 25
- Understanding FTP Commands on page 26
- Example: Configuring the FTP ALG on page 28
- Understanding the IPv6 FTP ALG for Routing on page 31

FTP ALG Overview

The File Transfer Protocol (FTP) is a widely and commonly used method of exchanging files over IP networks. In addition to the main control connection, data connections are also made for any data transfer between the client and the server; and the host, port, and direction are negotiated through the control channel.

For active mode FTP, the Junos OS stateful firewall service scans the client-to-server application data for the PORT command, which provides the IP address and port number to which the server connects. For passive-mode FTP, the Junos OS stateful firewall service scans the client-to-server application data for the PASV command and then scans the server-to-client responses for the 227 response, which contains the IP address and port number to which the client connects.

FTP represents the addresses and port numbers in ASCII. As a result, when addresses and ports are rewritten, the TCP sequence number might be changed, and thereafter the NAT service needs to maintain this delta in SEQ and ACK numbers by performing sequence NAT on all subsequent packets.

The FTP ALG supports the following:

- Automatically allocates data ports and firewall permissions for dynamic data connection
- Monitors the control connection in both active and passive modes
- Rewrites the control packets with the appropriate NAT address and port information
- Network Address Translation, Protocol Translation (NAT-PT)
- Transport Layer Security (TLS) as the security mechanism

Related Documentation

- Example: Configuring the FTP ALG on page 28

Copyright © 2016, Juniper Networks, Inc.

Understanding FTP Commands

The FTP ALG monitors commands and responses on the FTP control channel for syntactical correctness and opens corresponding pinholes to permit data channel connections to be established. In Junos OS Release 10.4, the FTP ALG supported IPv4 routing and NAT mode, and IPv6 routing mode only. In Junos OS Release 11.2 and later releases, the FTP ALG also supports IPv6 NAT and NAT-PT modes.

PORT Command

The PORT command is used in active FTP mode. The PORT command specifies the address and the port number to which a server should connect. When you use this command, the argument is a concatenation of a 32-bit Internet host address and a 16-bit TCP port address. The address information is broken into 8-bit fields, and the value of each field is transmitted as a decimal number (in character string representation). The fields are separated by commas.

The following is a sample PORT command, where h1 is the highest order 8-bit of the Internet host address:

PORT h1,h2,h3,h4,p1,p2

PASV Command

The PASV command requests a server to listen on a data port that is not the default data port of the server and to wait for a connection, rather than initiating another connection. The response to the PASV command includes the host and port address the server is listening on.

Extended FTP Commands

Extended FTP commands provide a method by which FTP can communicate the data connection endpoint information for network protocols other than IPv4. Extended FTP commands are specified in RFC 2428. In RFC 2428, the extended FTP commands EPRT and EPSV, replace the FTP commands PORT and PASV, respectively.

EPRT Command

The EPRT command allows for the specification of an extended address for the data connection. The extended address must consist of the network protocol as well as the network and transport addresses.

Copyright © 2016, Juniper Networks, Inc.

The format of EPRT is:

EPRT<space><d><net-prt><d><net-addr><d><tcp-port><d>

Parameter	Description
net-prt	An address family number defined by IANA.
net-addr	A protocol-specific string of the network address.
tcp-port	A TCP port number on which the host is listening for data connection.
Delimiter	The delimiter character must be one of the ASCII characters in range 33 to 126 inclusive. The character "|" (ASCII 124) is recommended.

The following command shows how to specify the server to use an IPv4 address to open a data connection to host 132.235.1.2 on TCP port 6275:

EPRT |1|132.235.1.2|6275|

The following command shows how to specify the server to use an IPv6 network protocol and a network address to open a TCP data connection on port 5282:

EPRT |2|1080::8:800:200C:417A|5282|

In this mode, FTP ALG focuses only on the EPRT command; it extracts the IPv6 address and port from the EPRT command and opens the pinhole.

EPSV mode

The EPSV command requests that a server listen on a data port and wait for a connection. The response to this command includes only the TCP port number of the listening connection.

An example response string is as follows:

```
Entering Extended Passive Mode (|||6446|)
```

NOTE: **The response code for entering passive mode using an extended address must be 229. You should extract the TCP port in 229 payloads and use it to open the pinhole.**

Related Documentation

- FTP ALG Overview on page 25
- Example: Configuring the FTP ALG on page 28

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the FTP ALG

This example shows how to configure the NAT-PT for FTP ALG.

- Requirements on page 28
- Overview on page 28
- Configuration on page 28
- Verification on page 30

Requirements

Before you begin:

- Configure proxy ARP for all IP addresses in the source NAT pool.
- Understand the concepts behind ALG for FTP. See "FTP ALG Overview" on page 25.

Overview

In this example, the ALG for FTP is configured to monitor and allow FTP traffic to be exchanged between the clients and the server located on opposite sides of a Juniper Networks device.

This example shows how to configure the NAT-PT for FTP ALG.

Configuration

Configuring a NAT Source Pool, NAT Static Pool and Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone untrust
set security nat source rule-set rs-source to zone trust
set security nat source rule-set rs-source rule src-nat match source-address 3333::130/128
set security nat source rule-set rs-source rule src-nat match destination-address
  40.0.0.211/32
set security nat source rule-set rs-source rule src-nat then source-nat interface
set security nat static rule-set rs2 from zone untrust
set security nat static rule-set rs2 rule r2 match destination-address 4444::141/128
set security nat static rule-set rs2 rule r2 then static-nat prefix 40.0.0.211/32
```

Step-by-Step Procedure

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

To configure a source NAT pool:

1. Create a source NAT, static NAT, and interface NAT rule set.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security nat source rule-set rs-source from zone untrust
user@host# set security nat source rule-set rs-source to zone trust
user@host# set security nat source rule-set rs-source rule src-nat match
  source-address 3333::130/128
user@host# set security nat source rule-set rs-source rule src-nat match
  destination-address 40.0.0.211/32
user@host# set security nat source rule-set rs-source rule src-nat then source-nat
  interface
user@host# set security nat static rule-set rs2 from zone untrust
user@host# set security nat static rule-set rs2 rule r2 match destination-address
  4444::141/128
user@host# set security nat static rule-set rs2 rule r2 then static-nat prefix
  40.0.0.211/32
```

2. Associate the NAT-PT application using a policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy ftp-basic
  match source-address any
user@host# set security policies from-zone trust to-zone untrust policy ftp-basic
  match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy ftp-basic
  match application junos-ftp
user@host# set security policies from-zone trust to-zone untrust policy ftp-basic
  then permit
```

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

```
user@host# show security nat
static {
  rule-set rs2 {
    from zone untrust;
    rule r2 {
      match {
        destination-address 4444::141/128;
      }
      then {
        static-nat {
          prefix {
            40.0.0.211/32
          }
        }
      }
    }
    }
  }

[edit]
user@host# show security policies
from-zone untrust to-zone trust {
  policy ftp-basic {
    match {
      source-address any;
      destination-address any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        application [ junos-ping junos-mgcp junos-ftp junos-rsh junos-h323 ];
      }
      then {
        permit;
      }
    }
  }
  default-policy {
    permit-all;
  }
```

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

Configuring FTP ALG Security Extension

Purpose Set the security alg ftp extension

Action From configuration mode, enter the following command.

set security alg ftp ftps-extension

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the NAT Source Pool, NAT Static Pool Rule Set on page 30
- Verifying FTP ALGs on page 30

Verifying the NAT Source Pool, NAT Static Pool Rule Set

Purpose Verify that the NAT source pool and rule set used to support the FTP ALG are working properly.

Action From operational mode, enter the **show configuration security nat** command.

Verifying FTP ALGs

Purpose Verify that FTP ALG is enabled.

Action From the operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
FTP      : Enabled
```

Meaning The output shows the FTP ALG status as follows:

- Enabled—Shows the FTP ALG is enabled.
- Disabled—Shows the FTP ALG is disabled.

NOTE: The FTP ALG is enabled by default.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- FTP ALG Overview on page 25

Understanding the IPv6 FTP ALG for Routing

File Transfer Protocol (FTP) is the part of the ALG that handles FTP traffic. The PORT/PASV requests and corresponding 200/227 responses in FTP are used to announce the TCP port, which the host listens to for the FTP data connection.

EPRT/EPSV/229 commands are used for these requests and responses. FTP ALG supports EPRT/EPSV/229 already, but only for IPv4 addresses.

In Junos OS Release 10.4, EPRT/EPSV/229 commands have been updated to support both IPv4 and IPv6 addresses.

FTP ALG uses preallocated objcache to store its session cookies. When both IPv4 and IPv6 addresses are supported on FTP ALG, the session cookie structure will enlarge by 256 bits (32 bytes) to store IPv6 address.

FTP ALG Support for IPv6

The FTP ALG monitors commands and responses on the FTP control channel for syntactical correctness and opens corresponding pinholes to permit data channel connections to be established. In Junos OS Release 10.4, the FTP ALG supported IPv4 routing, IPv6 routing, and NAT mode only. In Junos OS Release 11.2 and later releases, the FTP ALG also supports IPv6 NAT and NAT-PT modes.

EPRT mode

The EPRT command allows for the specification of an extended address for the data connection. The extended address must consist of the network protocol as well as the network and transport addresses.

The format of EPRT is:

EPRT<space><d><net-prt><d><net-addr><d><tcp-port><d>

- <net-prt>: An address family number defined by IANA
- <net-addr>: A protocol specific string of the network address
- <tcp-port>: A TCP port number

The following are sample EPRT commands for IPv6:

EPRT |2|1080::8:800:200C:417A|5282|

In this mode, FTP ALG focuses only on the EPRT command; it extracts the IPv6 address and port from the EPRT command and opens the pinhole.

Copyright © 2016, Juniper Networks, Inc.

EPSV mode

The EPSV command requests that a server be listening on a data port and waiting for a connection. The response to this command includes only the TCP port number of the listening connection.

An example response string is follows:

```
Entering Extended Passive Mode (|||6446|)
```

NOTE: The response code for entering passive mode using an extended address must be 229. You should extract the TCP port in 229 payloads and use it to open the pinhole.

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 5

Configuring the IKE and ESP ALG

- Understanding the ALG for IKE and ESP on page 33
- Understanding IKE and ESP ALG Operation on page 34
- Example: Configuring the IKE and ESP ALG on page 35
- Example: Enabling the IKE and ESP ALG and Setting Timeouts on page 41

Understanding the ALG for IKE and ESP

An SRX Series device can be used solely as a Network Address Translation (NAT) device when placed between VPN clients on the private side of the NAT gateway and the virtual private network (VPN) gateways on the public side.

Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) traffic is exchanged between the clients and the server. However, if the clients do not support NAT-Traversal (NAT-T) and if the device assigns the same NAT-generated IP address to two or more clients, the device will be unable to distinguish and route return traffic properly.

NOTE: If the user wants to support both NAT-T-capable and non-NAT-T-capable clients, then some additional configurations are required. If there are NAT-T capable clients, the user must enable the source NAT address persistence.

The ALG for IKE and ESP monitors IKE traffic between the client and the server and permits only one IKE Phase 2 message exchange between any given client/server pair, not just one exchange between any client and any server.

ALG for IKE and ESP traffic has been created and NAT has been enhanced to implement the following:

- To enable the SRX Series devices to pass IKE and ESP traffic with a source NAT pool
- To allow the device to be configured to return the same NAT-generated IP address for the same IP address without NAT ("address-persistent NAT"). As a result, the device is able to associate a client's outgoing IKE traffic with its return traffic from the server, especially when the IKE session times out and needs to be reestablished.

Copyright © 2016, Juniper Networks, Inc.

- The resulting ESP traffic between the client and the server is also allowed, especially in the direction from the server to the client.
- The return ESP traffic matches the following:
  - The server IP address as source IP
  - The client IP address as destination IP

NOTE: In SRX1500, SRX5600, and SRX5800 devices, IKE negotiations involving NAT traversal do not work if the IKE peer is behind a NAT device that will change the source IP address of the IKE packets during the negotiation. For example, if the NAT device is configured with DIP, it changes the source IP because the IKE protocol switches the UDP port from 500 to 4500.

Related Documentation

- Example: Configuring the IKE and ESP ALG on page 35

Understanding IKE and ESP ALG Operation

Application Layer Gateway (ALG) for Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) traffic has the following behavior:

- An IKE and ESP ALG monitors IKE traffic between the client and the server, and it permits only one IKE Phase 2 message exchange between the client and the server at any given time.
- For a Phase 2 message:
  - If a Phase 2 message exchange between the client and server does not happen, the IKE ALG gates are opened for the relevant ESP traffic from the client to the server and from the server to the client.
  - If both IKE ALG gates are not opened successfully, or if the Phase 2 message exchange already took place, then the Phase 2 message is dropped.
- When ESP traffic hits the IKE ALG gates, sessions are created to capture subsequent ESP traffic, and to perform the proper NATing (that is, the source IP address translation from the client to the server traffic and the destination IP address translation from the server to the client traffic).
- When the ESP traffic does not hit either one or both of the gates, then the gates naturally time out.
- Once the IKE ALG gates are collapsed or timed out, another IKE Phase 2 message exchange is permitted.
- IKE NAT-T traffic on floating port 4500 is processed in an IKE ALG. To support a mixture of NAT-T-capable and non-capable clients, you need to enable source NAT address persistent.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- ALG Overview on page 3
- Introduction to NAT on page 4839
- Understanding the ALG for IKE and ESP on page 33
- Example: Configuring the IKE and ESP ALG on page 35
- Example: Enabling the IKE and ESP ALG and Setting Timeouts on page 41

Example: Configuring the IKE and ESP ALG

This example shows how to configure the IKE and ESP ALG to pass through IKE and ESP traffic with a source NAT pool on Juniper Networks devices.

- Requirements on page 35
- Overview on page 35
- Configuration on page 35
- Verification on page 40

Requirements

Before you begin:

- Configure proxy ARP for all IP addresses in the source NAT pool.
- Understand the concepts behind IKE and ESP ALG. See "Understanding IKE and ESP ALG Operation" on page 34.

Overview

In this example, the ALG for IKE and ESP is configured to monitor and allow IKE and ESP traffic to be exchanged between the clients and the server located on opposite sides of a Juniper Networks device.

This example shows how to configure a source NAT pool and rule set, configure a custom application to support the IKE and ESP ALG, and associate this ALG to a policy.

If you want to support a mixture of NAT-traversal (NAT-T) capable clients and noncapable clients, you must enable persistent source NAT translation (so that once a particular source NAT is associated with a given IP address, subsequent source NAT translations use the same IP address). You also must configure a custom IKE NAT traversal application to support the encapsulation of IKE and ESP in UDP port 4500. This configuration enables IKE and ESP to pass through the NAT-enabled device.

Configuration

- Configuring a NAT Source Pool and Rule Set on page 36
- Configuring a Custom Application and Associating it to a Policy on page 37
- Configuring IKE and ESP ALG Support for Both NAT-T Capable and Noncapable Clients on page 38

Copyright © 2016, Juniper Networks, Inc.

Configuring a NAT Source Pool and Rule Set

CLI Quick Configuration

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

```
set security nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
set security zones security-zone green address-book address sa1 1.1.1.0/24
set security zones security-zone red address-book address da1 2.2.2.0/24
set security nat source rule-set rs1 from zone green
set security nat source rule-set rs1 to zone red
set security nat source rule-set rs1 rule r1 match source-address 1.1.1.0/24
set security nat source rule-set rs1 rule r1 match destination-address 2.2.2.0/24
set security nat source rule-set rs1 rule r1 then source-nat pool pool1
```

Step-by-Step Procedure

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

To configure a source NAT pool:

1. Create a NAT source pool.

```
[edit ]
user@host# set security nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
```

2. Configure security zone address book entries.

```
[edit]
user@host# set security zones security-zone green address-book address sa1
  1.1.1.0/24
user@host# set security zones security-zone red address-book address da1 2.2.2.0/24
```

3. Create a NAT source rule set.

```
[edit security nat source rule-set rs1]
user@host# set from zone green
user@host# set to zone red
user@host# set rule r1 match source-address 1.1.1.0/24
user@host# set rule r1 match destination-address 2.2.2.0/24
user@host# set rule r1 then source-nat pool pool1
```

Results

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

```
user@host# show security nat
source {
  pool pool1 {
    address {
      10.10.10.1/32 to 10.10.10.10/32;
    }
  }
  rule-set rs1 {
    from zone green;
```


Copyright © 2016, Juniper Networks, Inc.

```
    to zone red;
    rule r1 {
      match {
        source-address 1.1.1.0/24;
        destination-address 2.2.2.0/24;
      }
      then {
        source-nat {
          pool {
            pool1;
          }
        }
      }
    }
  }
```

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

Configuring a Custom Application and Associating it to a Policy

CLI Quick Configuration

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

```
set applications application custom-ike-alg source-port 500 destination-port 500 protocol
  udp application-protocol ike-esp-nat
set security policies from-zone green to-zone red policy pol1 match destination-address
  da1
set security policies from-zone green to-zone red policy pol1 match application
  custom-ike-alg
set security policies from-zone green to-zone red policy pol1 then permit
```

Step-by-Step Procedure

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

To configure a custom application and associate it to a policy:

1. Configure a custom application.

   ```
   [edit]
   user@host# set applications application custom-ike-alg source-port 500
     destination-port 500 protocol udp application-protocol ike-esp-nat
   ```

2. Associate the custom application to a policy.

   ```
   [edit security policies from-zone green to-zone red policy pol1]
   user@host# set match source-address sa1
   user@host# set match destination-address da1
   user@host# set match application custom-ike-alg
   user@host# set then permit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show applications** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show applications
application custom-ike-alg {
   application-protocol ike-esp-nat;
   protocol udp;
   source-port 500;
   destination-port 500;
}
```

```
[edit]
user@host# show security zones
security-zone Trust {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
      ge-0/0/1.0;
   }
}
security-zone green {
   address-book {
      address sa1 1.1.1.0/24;
   }
}
security-zone red {
   address-book {
      address da1 2.2.2.0/24;
   }
}
```

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

Configuring IKE and ESP ALG Support for Both NAT-T Capable and Noncapable Clients

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

```
set security nat source address-persistent
set applications application custom-ike-natt protocol udp source-port 4500
   destination-port 4500
set security policies from-zone green to-zone red policy pol1 match source-address sa1
set security policies from-zone green to-zone red policy pol1 match destination-address
   da1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone green to-zone red policy pol1 match application
  custom-ike-natt
set security policies from-zone green to-zone red policy pol1 then permit
```

Step-by-Step Procedure

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

To configure IKE and ESP ALG support for both NAT-T capable and noncapable clients:

1. Globally enable persistent source NAT translation.

   ```
   [edit]
   user@host# set security nat source address-persistent
   ```

2. Configure the IKE NAT-T application.

   ```
   [edit]
   user@host# set applications application custom-ike-natt protocol udp source-port
     4500 destination-port 4500
   ```

3. Associate the NAT-T application using a policy.

   ```
   [edit security policies from-zone green to-zone red policy pol1]
   user@host# set match source-address sa1
   user@host# set match destination-address da1
   user@host# set match application custom-ike-natt
   user@host# set then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  address-persistent;
}
```

```
[edit]
user@host# show security policies
from-zone green to-zone red {
  policy pol1 {
    match {
      source-address sa1;
      destination-address da1;
      application [ custom-ike-alg custom-ike-natt ];
    }
    then {
      permit;
    }
  }
}
default-policy {
  permit-all;
}
```

Copyright © 2016, Juniper Networks, Inc.

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying IKE and ESP ALG Custom Applications on page 40
- Verifying the Security Polices of ALG on page 40

Verifying IKE and ESP ALG Custom Applications

Purpose Verify that the custom applications to support the IKE and ESP ALG are enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status

ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Enabled
  MGCP     : Enabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Enabled
```

Meaning The output shows the ALG status as follows:

- Enabled—Shows the ALG is enabled.
- Disabled—Shows the ALG is disabled.

Verifying the Security Polices of ALG

Purpose Verify that the application custom IKE ALG and application custom IKE NATT are set.

Action From operational mode, enter the **show security policies** command.

```
user@host> show security policies

Default policy: permit-all
From zone: green, To zone: red
Policy: pol1, State: enabled, Index: 7, Scope Policy: 0, Sequence number: 1
Source addresses: sa1
Destination addresses: da1
Applications: custom-ike-alg, custom-ike-natt
Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

Meaning The sample output shows that custom IKE ALG and custom IKE NATT applications are set.

Related Documentation

- ALG Overview on page 3
- Understanding the ALG for IKE and ESP on page 33
- Example: Enabling the IKE and ESP ALG and Setting Timeouts on page 41

Example: Enabling the IKE and ESP ALG and Setting Timeouts

This example shows how to enable the IKE and ESP ALG and set the timeout values to allow time for the ALG to process ALG state information, ESP gates, and ESP sessions.

- Requirements on page 41
- Overview on page 41
- Configuration on page 41
- Verification on page 42

Requirements

Understand the concepts behind ALG for IKE and ESP. See "Understanding IKE and ESP ALG Operation" on page 34.

Overview

The IKE and ESP ALG processes all traffic specified in any policy to which the ALG is attached. In this example, you configure the **set security alg ike-esp-nat enable** statement so the current default IPsec pass-through behavior is disabled for all IPsec pass-through traffic, regardless of policy.

You then set the timeout values to allow time for the IKE and ESP ALG to process ALG state information, ESP gates, and ESP sessions. In this example, you set the timeout of ALG state information. The timeout range is 180 through 86400 seconds. The default timeout is 14400 seconds. You then set the timeout of the ESP gates created after an IKE Phase 2 exchange has completed. The timeout range is 2 through 30 seconds. The default timeout is 5 seconds. Finally, you set the idle timeout of the ESP sessions created from the IPsec gates. If no traffic hits the session, it is aged out after this period of time. The timeout range is 60 through 2400 seconds. The default timeout is 1800 seconds.

Configuration

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

```
set security alg ike-esp-nat enable
set security alg ike-esp-nat esp-gate-timeout 20
set security alg ike-esp-nat esp-session-timeout 2400
set security alg ike-esp-nat state-timeout 360
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To enable the IKE and ESP ALG and set the timeout values:

1. Enable the IKE and ESP ALG.

   ```
   [edit]
   user@host# set security alg ike-esp-nat enable
   ```

2. Set the timeout for the ALG state information.

   ```
   [edit security alg ike-esp-nat]
   user@host# set state-timeout 360
   ```

3. Set the timeout for the ESP gates created after an IKE Phase 2 exchange has completed.

   ```
   [edit security alg ike-esp-nat]
   user@host# set esp-gate-timeout 20
   ```

4. Set the idle timeout for the ESP sessions created from the IPsec gates.

   ```
   [edit security alg ike-esp-nat]
   user@host# set esp-session-timeout 2400
   ```

Results

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

```
[edit]
user@host# show security alg
ike-esp-nat {
   enable;
   state-timeout 360;
   esp-gate-timeout 20;
   esp-session-timeout 2400;
}
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

Verifying the ALG for IKE and ESP and Timeout Settings

Purpose

Verify that the ALG for IKE and ESP is enabled and the timeout settings for this feature are correct.

Action

From operational mode, enter the **show security alg ike-esp-nat** command.

Related Documentation

- ALG Overview on page 3
- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

- Understanding the ALG for IKE and ESP on page 33
- Understanding IKE and ESP ALG Operation on page 34
- Example: Configuring the IKE and ESP ALG on page 35

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 6

Configuring the PPTP ALG

- Understanding the PPTP ALG on page 45
- Understanding IPv6 Support for the PPTP ALG on page 45
- Example: Configuring the PPTP ALG on page 46

Understanding the PPTP ALG

The Point-to-Point Tunneling Protocol (PPTP) ALG is used for tunneling Point-to-Point Protocol (PPP) packets over an IP network. The PPTP ALG is often used to implement a client/server architecture, a PPTP network server, and a PPTP access concentrator.

The PPTP ALG processes PPTP packets, performs Network Address Translation (NAT), open pinholes for new data connections between a client and a server, and transfers data between a client and a server located on opposite sides of a Juniper Networks device.

Related Documentation

- Example: Configuring the PPTP ALG on page 46

Understanding IPv6 Support for the PPTP ALG

IPv6 is supported on the PPTP ALG.

The PPTP ALG supports IPv6 data packets. The PPTP ALG uses TCP port 1723 to connect and disconnect a client and a server.

To support IPv6, the PPTP ALG parses both IPv4 and IPv6 PPTP packets, performs NAT, and then opens a pinhole for the data tunnel. The flow module supports IPv6 to parse the GRE packet and use the GRE call ID as fake port information to search the session table and gate table.

The PPTP ALG with IPv6 support does not support NAT-PT and NAT64, because PPP packets are compressed with Microsoft Point-to-Point Encryption (MPPE) protocol after the tunnel is set up; therefore translation of the IP header in the PPP package cannot be handled.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The PPTP ALG can support NAT64 in a specific scenario in which translation of the IP header in the PPP package is not required—that is, if the PPTP client works in dual-stack mode in the IPv6 network and server in the IPv4 network.

- The PPTP ALG with IPv6 support has the following limitation:
  - Because PPP packets are compressed with Microsoft Point-to-Point Encryption (MPPE) protocol after the tunnel is set up, translation of the IP header in the PPP package cannot be handled; therefore, to make sure PPTP connection works well, the PPTP client must be able to work in dual stack mode. So that an IPv6 PPTP client can accept an IPv4 address for PPP tunnel interface, by which it can communicate with the IPv4 PPTP server without IP address translation for PPP packets.

Related Documentation

- Understanding the PPTP ALG on page 45

Example: Configuring the PPTP ALG

The PPTP ALG processes PPTP packets, performs NAT, and open pinholes for new data connections between a client and a server.

This example shows how to configure the PPTP ALG in route or NAT mode. The configuration allows PPTP traffic to pass through a device, transferring data between a client and a server located on opposite sides of a Juniper Networks device.

- Requirements on page 46
- Overview on page 47
- Configuration on page 48
- Verification on page 56

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Two PCs (client and server)

Before you begin:

- Understand the concepts behind ALGs. See "ALG Overview" on page 3.
- Understand the basics of PPTP ALG. See "Understanding the PPTP ALG" on page 45.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, first you configure network interfaces on the device, create security zones and assign interfaces to the zones, and configure a policy to allow PPTP traffic to go through an SRX Series device.

Then you create a static NAT rule set rs1 with a rule r1 to match with the destination address 30.5.2.120/32, and you create a static NAT prefix with address 10.5.1.120/32.

Next you create a source NAT pool src-p1 with a source rule set src-rs1 to translate packets from zone trust to zone untrust. For matching packets, the source address is translated to an IP address in the src-p1 pool.

Then you create a destination NAT pool des-p1 with a destination rule set des-rs1 to translate packets from zone trust to destination address 30.5.1.120/32. For matching packets, the destination address is translated to an IP address in the des-p1 pool. Finally, you configure PPTP ALG trace options.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 3 shows the PPTP ALG topology.

Figure 3: PPTP ALG Topology

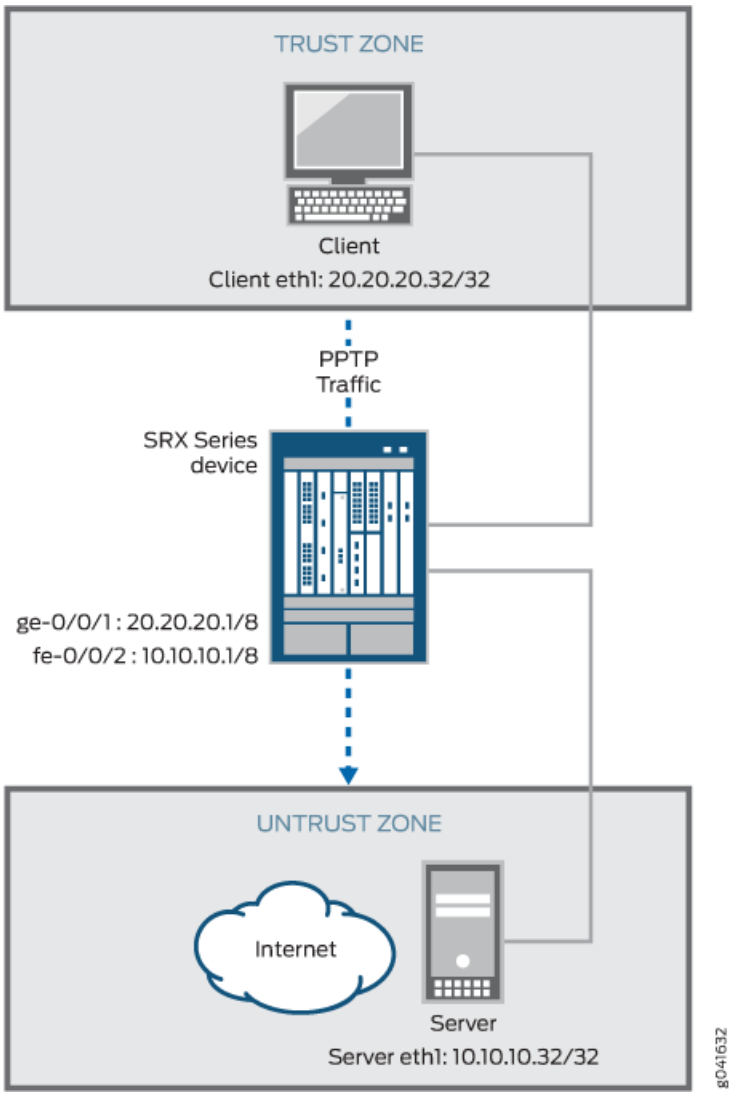

Configuration

To configure the PPTP ALG, perform these tasks:

- Configuring a Route Mode on page 49
- Configuring a Static NAT Rule Set on page 51
- Configuring a Source NAT Pool and Rule Set on page 52
- Configuring a Destination NAT Pool and Rule Set on page 54
- Configuring PPTP ALG trace options on page 55

Copyright © 2016, Juniper Networks, Inc.

Configuring a Route Mode

CLI Quick Configuration

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

```
set interfaces ge-0/0/1 unit 0 family inet address 20.20.20.1/8
set interfaces fe-0/0/2 unit 0 family inet address 10.10.10.1/8
set security zones security-zone trust interfaces ge-0/0/1 host-inbound-traffic system-services
 all
set security zones security-zone trust interfaces ge-0/0/1 host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/2 host-inbound-traffic system-services
 all
set security zones security-zone untrust interfaces fe-0/0/2 host-inbound-traffic protocols all
set security policies from-zone trust to-zone untrust policy pptp match source-address any
set security policies from-zone trust to-zone untrust policy pptp match destination-address any
set security policies from-zone trust to-zone untrust policy pptp match application junos-pptp
set security policies from-zone trust to-zone untrust policy pptp then permit
```

Step-by-Step Procedure

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

To configure route mode:

1. Configure interfaces.

   ```
   [edit interfaces]
   user@host#set ge-0/0/1 unit 0 family inet address 20.20.20.1/8
   user@host#set fe-0/0/2 unit 0 family inet address 10.10.10.1/8
   ```

2. Configure zones and assign interfaces to the zones.

   ```
   [edit security zones security-zone trust]
   user@host#set interfaces ge-0/0/1 host-inbound-traffic system-services all
   user@host#set interfaces ge-0/0/1 host-inbound-traffic protocols all
   [edit security zones security-zone untrust]
   user@host#set interfaces fe-0/0/2 host-inbound-traffic system-services all
   user@host#set interfaces fe-0/0/2 host-inbound-traffic protocols all
   ```

3. Configure a PPTP policy that allows PPTP traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone untrust to-zone trust]
   user@host#set policy pptp match source-address any
   user@host#set policy pptp match destination-address any
   user@host#set policy pptp match application junos-pptp
   user@host#set policy pptp then permit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
user@host# show interfaces
...
   ge-0/0/1 {
   unit 0 {
      family inet {
         address 20.20.20.1/8;
      }
   }
}
fe-0/0/2 {
   unit 0 {
      family inet {
         address 10.10.10.1/8;
      }
   }
}
...
[edit]
user@host# show security zones
security-zone trust {
   ....
      interfaces {
      ge-0/0/1 {
         host-inbound-traffic {
            system-services {
               all;
            }
            protocols {
               all;
            }
         }
      }
   }
}
...
security-zone untrust {
   interfaces {
      fe-0/0/2 {
         host-inbound-traffic {
            system-services {
               all;
            }
            protocols {
               all;
            }
         }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}

[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy pptp {
    match {
      source-address any;
      destination-address any;
      application junos-pptp;
    }
    then {
      permit;
    }
  }
}
```

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

Configuring a Static NAT Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone trust
set security nat static rule-set rs1 rule r1 match destination-address 30.5.2.120/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 10.5.1.120/32
```

Step-by-Step Procedure

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

To configure a static NAT rule set:

1. Create a static NAT rule set.

   ```
   [edit security nat static rule-set rs1]
   user@host#set from zone trust
   ```

2. Define the rule to match with the destination address.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 match destination-address 30.5.2.120/32
   ```

3. Define the static NAT prefix for the device.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 then static-nat prefix 10.5.1.120/32
   ```

Results

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

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
    from zone trust;
    rule r1 {
      match {
        destination-address 30.5.2.120/32;
      }
      then {
        static-nat {
          prefix {
            10.5.1.120/32;
          }
        }
      }
    }
  }
}
```

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

Configuring a Source NAT Pool and Rule Set

CLI Quick Configuration

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

```
set security nat source pool src-p1 address 30.5.1.120/32
set security nat source rule-set src-rs1 from zone trust
set security nat source rule-set src-rs1 to zone untrust
set security nat source rule-set src-rs1 rule src-r1 match source-address 20.5.1.120/32
set security nat source rule-set src-rs1 rule src-r1 match destination-address 10.5.2.120/32
set security nat source rule-set src-rs1 rule src-r1 then source-nat pool src-p1
```

Step-by-Step Procedure

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

To configure a source NAT pool and rule set:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host#set pool src-p1 address 30.5.1.120/32
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source ]
   user@host# set rule-set src-rs1 from zone trust
   user@host# set rule-set src-rs1 to zone untrust
   ```

3. Configure a rule that matches packets and translates the source address to an address in the source pool.

   ```
   [edit security nat source]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set rule-set src-rs1 rule src-r1 match source-address 20.5.1.120/32
```

4. Configure a rule that matches packets and translates the destination address to an address in the source pool.

```
[edit security nat source]
user@host# set rule-set src-rs1 rule src-r1 match destination-address 10.5.2.120/32
```

5. Configure a source NAT pool in the rule.

```
[edit security nat source]
user@host# set rule-set src-rs1 rule src-r1 then source-nat pool src-p1
```

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

```
[edit]
user@host# show security nat
source {
  pool src-p1 {
    address {
      30.5.1.120/32;
    }
  }
  rule-set src-rs1 {
    from zone trust;
    to zone untrust;
    rule src-r1 {
      match {
        source-address 20.5.1.120/32;
        destination-address 10.5.2.120/32;
      }
      then {
        source-nat {
          pool {
            src-p1;
          }
        }
      }
    }
  }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Configuring a Destination NAT Pool and Rule Set

CLI Quick Configuration

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

```
set security nat destination pool des-p1 address 10.5.1.120/32
set security nat destination rule-set des-rs1 from zone trust
set security nat destination rule-set des-rs1 rule des-r1 match source-address 20.5.1.120/32
set security nat destination rule-set des-rs1 rule des-r1 match destination-address 30.5.1.120/32
set security nat destination rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
```

Step-by-Step Procedure

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

To configure a destination NAT pool and rule set:

1. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host#set pool des-p1 address 10.5.1.120/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 from zone trust
   ```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match source-address 20.5.1.120/32
   ```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match destination-address 30.5.1.120/32
   ```

5. Configure a source NAT pool in the rule.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
   ```

Results

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

```
[edit]
user@host# show security nat
destination {
  pool des-p1 {
    address {
      10.5.1.120/32;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    rule-set des-rs1 {
      from zone trust;
      rule des-r1 {
        match {
          source-address 20.5.1.120/32;
          destination-address 30.5.1.120/32;
        }
        then {
          destination-nat {
            pool {
              des-p1;
            }
          }
        }
      }
    }
  }
```

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

Configuring PPTP ALG trace options

CLI Quick Configuration

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

```
set security alg pptp traceoptions flag all
set security alg traceoptions file trace
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
```

Step-by-Step Procedure

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

To configure PPTP ALG trace options:

1. Enable PPTP ALG trace options.

   ```
   [edit security alg]
   user@host#set pptp traceoptions flag all
   ```

2. Configure a filename to receive output from the tracing operation.

   ```
   [edit security alg]
   user@host#set traceoptions file trace
   ```

3. Specify the maximum trace file size.

   ```
   [edit security alg]
   user@host#set traceoptions file size 1g
   ```

4. Specify the level of tracing output.

   ```
   [edit security alg]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host#set traceoptions level verbose
```

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

```
[edit]
user@host# show security alg
traceoptions {
   file trace size 1g;
   level verbose;
   }
pptp traceoptions flag all;
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the PPTP ALG Control Session on page 57
- Verifying the PPTP ALG Flow Gate Information on page 57
- Verifying PPTP ALG on page 58
- Verifying the PPTP Resource Manager Group on page 58
- Verifying the PPTP Resource Information on page 58

Copyright © 2016, Juniper Networks, Inc.

Verifying the PPTP ALG Control Session

Purpose Verify that the PPTP control session is created and all the PPTP control and data sessions are created.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
 SSession ID: 57, Policy name: pptp, Timeout: 1787
Resource information : PPTP ALG, 1, 0
In: 20.20.20.32/3905 --> 10.10.10.32/1723;tcp, If: ge-0/0/1.0 Pkts: 6, Bytes: 584
Out: 10.10.10.32/1723 --> 20.20.20.32/3905;tcp, If: fe-0/0/2.0 Pkts: 4, Bytes:
352

Session ID: 58, Policy name: pptp, Timeout: 1799
In: 20.20.20.32/0 --> 10.10.10.32/256;gre, If: ge-0/0/1.0
Out: 10.10.10.32/256 --> 20.20.20.32/65001;gre, If: fe-0/0/2.0

Session ID: 59, Policy name: pptp, Timeout: 1787
In: .10.10.10.32/0 --> 20.20.20.32/260;gre, If: ge-0/0/1.0
Out: 20.20.20.32/260 --> 10.10.10.32/65000;gre, If: fe-0/0/2.0
```

Meaning

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **Policy name**—Policy name that permitted the traffic.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and the source interface for this session is ge-0/0/1.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and destination interface for this session is fe-0/0/2.0).

Verifying the PPTP ALG Flow Gate Information

Purpose Verify that the flow gate is opened for TCP data channel connection.

Action From operational mode, enter the **show security flow gate** command.

```
user@host>show security flow gate

  Hole: 20.0.172.24-20.0.172.24/0-0->21.0.172.38-21.0.172.38/25750-25750
  Translated: 2015::172:24/65000->2005::172:108/360
  Protocol: gre
  Application: PPTP ALG/69
  Age: 118 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 12-1-1

  Hole: 2005::172:108-0-0->2015::172:24-2432-2432
  Translated: 21.0.172.38/65001->20.0.172.24/2432
  Protocol: gre
  Application: PPTP ALG/69
```

Copyright © 2016, Juniper Networks, Inc.

```
  Age: 120 seconds
  Flags: 0x8080
  Zone: untrust
  Reference count: 1
  Resource: 12-1-2

Valid gates: 2
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 2
```

Verifying PPTP ALG

Purpose Verify that the PPTP ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host>show security alg status
ALG Status :
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the PPTP ALG status as follows:

- Enabled—Shows the PPTP ALG is enabled.
- Disabled—Shows the PPTP ALG is disabled.

Verifying the PPTP Resource Manager Group

Purpose Verify the total number of resource manager groups and active groups that are used by the PPTP ALG.

Action From operational mode, enter the **show security resource-manager group active** command.

```
user@host>show security resource-manager group active
Group ID 1: Application - PPTP ALG
       Total groups 19763, active groups 1
```

Verifying the PPTP Resource Information

Purpose Verify the total number of resources and active resources that are used by the PPTP ALG.

Action From operational mode, enter the **show security resource-manager resource active** command.

```
user@host>show security resource-manager resource active
Resource ID 2: Group ID - 1, Application - PPTP ALG
```


Copyright © 2016, Juniper Networks, Inc.

```
Resource ID 1: Group ID - 1, Application - PPTP ALG
Total Resources 93286, active resources 2
```

Related Documentation

- Understanding the PPTP ALG on page 45

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 7

Configuring the RPC ALG

- Understanding RPC ALGs on page 61
- Understanding Sun RPC ALGs on page 62
- Enabling Sun RPC ALGs (CLI Procedure) on page 63
- Enabling Sun RPC ALGs (J-Web Procedure) on page 63
- Customizing Sun RPC Applications (CLI Procedure) on page 64
- Understanding Sun RPC Services on page 64
- Understanding Microsoft RPC ALGs on page 67
- Configuring the Microsoft RPC ALG on page 68
- Enabling Microsoft RPC ALGs (CLI Procedure) on page 70
- Enabling Microsoft RPC ALGs (J-Web Procedure) on page 70
- Verifying the Microsoft RPC ALG Tables on page 71
- Understanding Microsoft RPC Services on page 71
- Customizing Microsoft RPC Applications (CLI Procedure) on page 73

Understanding RPC ALGs

Junos OS supports basic Remote Procedure Call Application Layer Gateway (RPC ALG) services. RPC is a protocol that allows an application running in one address space to access the resources of applications running in another address space as if the resources were local to the first address space. The RPC ALG is responsible for RPC packet processing.

The RPC ALG in Junos OS supports the following services and features:

- Sun Microsystems RPC Open Network Computing (ONC)
- Microsoft RPC Distributed Computing Environment (DCE)
- Dynamic port negotiation
- Ability to allow and deny specific RPC services
- Static Network Address Translation (NAT) and source NAT (with no port translation)
- RPC applications in security policies

Copyright © 2016, Juniper Networks, Inc.

Use the RPC ALG if you need to run RPC-based applications such as NFS or Microsoft Outlook. The RPC ALG functionality is enabled by default.

Related Documentation

- ALG Overview on page 3
- Understanding Sun RPC ALGs on page 62
- Understanding Microsoft RPC ALGs on page 67

Understanding Sun RPC ALGs

Sun Microsystems Remote Procedure Call (Sun RPC)—also known as Open Network Computing Remote Procedure Call (ONC RPC)—provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

Junos OS supports the Sun RPC as a predefined service and allows and denies traffic based on a security policy you configure. The Application Layer Gateway (ALG) provides the functionality for Juniper Networks devices to handle the dynamic transport address negotiation mechanism of the Sun RPC and to ensure program number-based security policy enforcement. You can define a security policy to permit or deny all RPC requests, or to permit or deny by specific program number. The ALG also supports route mode and Network Address Translation (NAT) mode for incoming and outgoing requests.

When an application or a PC client calls a remote service, it needs to find the transport address of the service. In the case of TCP/UDP, the address is a port number. A typical procedure for this case is as follows:

1. The client sends the GETPORT message to the RPCBIND service on the remote machine. The GETPORT message contains the program number, and version and procedure number of the remote service it is attempting to call.
2. The RPCBIND service replies with a port number.
3. The client calls the remote service using the port number returned.
4. The remote service replies to the client.

A client also can use the CALLIT message to call the remote service directly, without determining the port number of the service. In this case, the procedure is as follows:

1. The client sends a CALLIT message to the RPCBIND service on the remote machine. The CALLIT message contains the program number and the version and procedure number of the remote service it attempting to call.
2. RPCBIND calls the service for the client.
3. RCPBIND replies to the client if the call has been successful. The reply contains the call result and the service's port number.

Copyright © 2016, Juniper Networks, Inc.

The Sun RPC ALG dynamically allocates new mapping entries instead of using a default size (512 entries). It also offers a flexible time-based RPC mapping entry that removes the mapping entry (auto-clean) without affecting the associated active RPC sessions, including both control session and data session.

You can define the Sun RPC mapping entry ageout value using the **set security alg sunrpc map-entry-timeout value** command. The ageout value ranges from 1 hour to 72 hours, and the default value is 32 hours.

If the Sun RPC ALG service does not trigger the control negotiation even after 72 hours, the maximum RPC ALG mapping entry value times out and the new data connection to the service fails.

Related Documentation

- Understanding RPC ALGs on page 61
- Enabling Sun RPC ALGs (J-Web Procedure) on page 63
- Enabling Sun RPC ALGs (CLI Procedure) on page 63
- Understanding Sun RPC Services on page 64
- Understanding Microsoft RPC ALGs on page 67

Enabling Sun RPC ALGs (CLI Procedure)

The Sun RPC ALG is enabled by default and requires no configuration.

To disable the Sun RPC ALG, enter the following command:

```
user@host# set security alg sunrpc disable
```

To reenable the Sun RPC ALG, enter the following command:

```
user@host# delete security alg sunrpc
```

Related Documentation

- Understanding Sun RPC ALGs on page 62
- Enabling Sun RPC ALGs (J-Web Procedure) on page 63

Enabling Sun RPC ALGs (J-Web Procedure)

The Sun RPC ALG is enabled by default and requires no configuration.

To disable or reenable the RPC ALG:

1. Select **Configure>Security>ALG**.
2. Select the **Enable SUNRPC** check box.
3. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.

Related Documentation

- Understanding Sun RPC ALGs on page 62
- Enabling Sun RPC ALGs (CLI Procedure) on page 63

Copyright © 2016, Juniper Networks, Inc.

Customizing Sun RPC Applications (CLI Procedure)

All Sun RPC applications can be customized by using a predefined application set.

For example, an application can be customized to open the control session only and not allow any data sessions:

```
application-set junos-sun-rpc {
  application junos-sun-rpc-tcp;
  application junos-sun-rpc-udp;
}
```

In the following example, the predefined application set allows data sessions only. It will not work without the control session:

```
application-set junos-sun-rpc-portmap {
  application junos-sun-rpc-portmap-tcp;
  application junos-sun-rpc-portmap-udp;
}
```

To customize all Sun RPC applications with predefined application sets, use both application sets in the policy:

```
application-set [junos-sun-rpc junos-sun-rpc-portmap]
```

NOTE: MS RPC applications are customized in the same way as SUN RPC applications.

Related Documentation

- Understanding Sun RPC ALGs on page 62
- Customizing Microsoft RPC Applications (CLI Procedure) on page 73

Understanding Sun RPC Services

Sun RPC, also known as Open Network computing remote procedure call (ONC RPC), provides a way for a program running on one host to call procedures in a program running on another host. Sun RPC services are defined by a program identifier. The program identifier is independent of any transport address, and most of the Sun RPC sessions are initiated through TCP or UDP port 111. Each host links the required RPC service to a dynamic TCP or UDP port that is negotiated over the port 111 control channel, allowing the client to connect to either TCP or UDP port 111.

Predefined Sun Microsystems remote procedure call (Sun RPC) services include:

- **junos-sun-rpc-tcp**
- **junos-sun-rpc-udp**

Copyright © 2016, Juniper Networks, Inc.

The Sun RPC ALG can be applied by using the following methods:

- ALG default application—Use one of the following predefined applications for control and data connections in your policy:
  - **junos-sun-rpc-any-tcp**
  - **junos-sun-rpc-any-udp**
  - **junos-sun-rpc-mountd-tcp**
  - **junos-sun-rpc-mountd-udp**
  - **junos-sun-rpc-nfs-tcp**
  - **junos-sun-rpc-nfs-udp**
  - **junos-sun-rpc-nlockmgr-tcp**
  - **junos-sun-rpc-nlockmgr-udp**
  - **junos-sun-rpc-portmap-tcp**
  - **junos-sun-rpc-portmap-udp**
  - **junos-sun-rpc-rquotad-tcp**
  - **junos-sun-rpc-rquotad-udp**
  - **junos-sun-rpc-ruserd-tcp**
  - **junos-sun-rpc-ruserd-udp**
  - **junos-sun-rpc-sadmind-tcp**
  - **junos-sun-rpc-sadmind-udp**
  - **junos-sun-rpc-sprayd-tcp**
  - **junos-sun-rpc-sprayd-udp**
  - **junos-sun-rpc-status-tcp**
  - **junos-sun-rpc-status-udp**
  - **junos-sun-rpc-walld-tcp**
  - **junos-sun-rpc-walld-udp**
  - **junos-sun-rpc-ypbind-tcp**
  - **junos-sun-rpc-ypbind-udp**
  - **junos-sun-rpc-ypserv-tcp**
  - **junos-sun-rpc-ypserv-udp**
- Default control application—Use the predefined control through **junos-sun-rpc**:
  - Create an application for data (**USER_DEFINED_DATA**). You can make a set of your own data (for example, **my_rpc_application_set**) and use it in the policy.

Copyright © 2016, Juniper Networks, Inc.

- ALG default application set—Use the predefined application set for control and customized data application in the policy:
  - **junos-sun-rpc** (for control sessions)
  - **junos-sun-rpc-any**
  - **junos-sun-rpc-mountd**
  - **junos-sun-rpc-nfs**
  - **junos-sun-rpc-nfs-access**
  - **junos-sun-rpc-nlockmgr**
  - **junos-sun-rpc-portmap** (for data sessions)
  - **junos-sun-rpc-rquotad**
  - **junos-sun-rpc-ruserd**
  - **junos-sun-rpc-sadmind**
  - **junos-sun-rpc-sprayd**
  - **junos-sun-rpc-status**
  - **junos-sun-rpc-walld**
  - **junos-sun-rpc-ypbind**
  - **junos-sun-rpc-ypserv**
- Custom control and custom data application—Use a customized application:
  - Create an application for control (**USER_DEFINED_CONTROL**) and data (**USER_DEFINED_DATA**).
  - In the policy, use the user-defined application set for a control and customized data application:
    - **USER_DEFINED_CONTROL**
    - **USER_DEFINED_DATA**

Table 3 lists predefined Sun RPC services, a program identifier associated with each service, and a description of each service.

Table 3: Predefined Sun RPC Services

Service	Program ID	Description
PORTMAP	100000	Sun RPC Portmapper protocol is a TCP or UDP port-based service that includes TCP or UDP port 111.
NFS	100003	Sun RPC Network File System.
MOUNT	100005	Sun RPC mount process.

Copyright © 2016, Juniper Networks, Inc.

Table 3: Predefined Sun RPC Services *(continued)*

Service	Program ID	Description
YPBIND	100007	Sun RPC Yellow Page Bind service.
STATUS	100024	Sun RPC status.

Related Documentation

- Understanding Sun RPC ALGs on page 62
- Customizing Sun RPC Applications (CLI Procedure) on page 64
- Understanding Microsoft RPC Services on page 71

Understanding Microsoft RPC ALGs

Microsoft Remote Procedure Call (MS-RPC) is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun RPC, MS-RPC provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service program's universal unique identifier (UUID). The specific UUID is mapped to a transport address.

Junos OS devices running Junos OS support MS-RPC as a predefined service and allow and deny traffic based on a policy you configure. The Application Layer Gateway (ALG) provides the functionality for Juniper Networks devices to handle the dynamic transport address negotiation mechanism of the MS-RPC, and to ensure UUID-based security policy enforcement. You can define a security policy to permit or deny all RPC requests, or to permit or deny by specific UUID number. The ALG also supports route mode and Network Address Translation (NAT) mode for incoming and outgoing requests.

When both the MS-RPC client and MS-RPC server are 64 bit capable (such as MS Exchange 2008), they negotiate to use NDR64 transfer syntax during the network communication. when you use NDR64, the interface parameters should be encoded according to NDR64 syntax, because the packet format for NDR64 is different from the packet format for NDR20 (32 bit version).

In MS-RPC, there is a remote activation interface of the DCOM Remote Protocol called ISystemActivator (also known as IRemoteSCMActivator). It is used by the Windows Management Instrumentation Command-line (WMIC), Internet Information Services (IIS), and many other applications that are used extensively.

The MS-RPC ALG dynamically allocates new mapping entries instead of using a default size (512 entries). It also offers a flexible time-based RPC mapping entry that removes the mapping entry (auto-clean) without affecting the associated active RPC sessions, including both control session and data session.

You can define the MS-RPC mapping entry ageout value using the **set security alg msrpc map-entry-timeout value** command. The ageout value ranges from 1 hour to 72 hours, and the default value is 32 hours.

Copyright © 2016, Juniper Networks, Inc.

If the MS-RPC ALG service does not trigger the control negotiation even after 72 hours, the maximum RPC ALG mapping entry value times out and the new data connection to the service fails.

Related Documentation

- Understanding RPC ALGs on page 61
- Enabling Microsoft RPC ALGs (J-Web Procedure) on page 70
- Enabling Microsoft RPC ALGs (CLI Procedure) on page 70
- Understanding Microsoft RPC Services on page 71
- Understanding Sun RPC ALGs on page 62
- Verifying the Microsoft RPC ALG Tables on page 71

Configuring the Microsoft RPC ALG

You can configure the Microsoft RPC ALG using the following three methods:

- Configuring the MS-RPC ALG with a Predefined Microsoft Application on page 68
- Configuring the MS-RPC ALG with a Wildcard UUID on page 69
- Configuring the MS-RPC ALG with a Specific UUID on page 69

Configuring the MS-RPC ALG with a Predefined Microsoft Application

There are several predefined MS applications. To view the predefined Microsoft applications from the CLI, enter the **show configuration groups junos-defaults** command.

```
user@host> show security policies
  from-zone trust to-zone untrust {
    policy p1 {
      match {
        source-address any;
        destination-address any;
        application junos-ms-rpc-msexchange;
      }
      then {
        permit;
      }
    }
  }
```

After you commit the configuration, from the CLI, enter the **show security alg msrpc object-id-map** command to view the output.

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
1544f5e0-613c-11d1-93df-00c04fd7bd09    0x80000001
a4f1db00-ca47-1067-b31f-00dd010662da    0x80000002
f5cc5a18-4264-101a-8c59-08002b2f8426    0x80000003
```

The output shows that the UUID has been applied for the policy.

Copyright © 2016, Juniper Networks, Inc.

Configuring the MS-RPC ALG with a Wildcard UUID

To permit the configuration for any MS RPC application, add the **application junos-ms-rpc-any** statement to the Permit configuration.

```
user@host> show security policies
  from-zone trust to-zone untrust {
  policy p1 {
    match {
      source-address any;
      destination-address any;
      application junos-ms-rpc-any;
    }
    then {
      permit;
    }
  }
}
```

After you commit the configuration, from the CLI, enter the **show security alg msrpc object-id-map** command to view the output.

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
ffffffff-ffff-ffff-ffff-ffffffffffff    0x80000004
```

Configuring the MS-RPC ALG with a Specific UUID

For applications that have not been predefined, you need to manually configure a specific UUID. For example, to permit a NETLOGON application that has not been predefined, you add the **application msrpc-netlogon** statement to the Permit configuration.

```
user@host> show applications
  application msrpc-netlogon {
  term t1 protocol tcp uuid 12345678-1234-abcd-ef00-01234567cffb;
  term t2 protocol udp uuid 12345678-1234-abcd-ef00-01234567cffb;
  term t3 protocol tcp uuid 12345778-1234-abcd-ef00-0123456789ab;
}
user@host> show security policies
  from-zone trust to-zone untrust {
  match {
    source-address any;
    destination-address any;
    application msrpc-netlogon;
  }
  then {
    permit;
  }
  }
}
```

After you commit the configuration, from the CLI, enter the **show security alg msrpc object-id-map** command to verify the Microsoft universal unique identifier to Object ID

Copyright © 2016, Juniper Networks, Inc.

(UUID-to-OID) mapping table. The Microsoft RPC ALG monitors packets on TCP port 135.

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
12345778-1234-abcd-ef00-0123456789ab    0x80000006
12345678-1234-abcd-ef00-01234567cffb    0x80000005
be617c0-31a5-11cf-a7d8-00805f48a135     0x80000020
e3514235-4b06-11d1-ab04-00c04fc2dcd2    0x80000002
67df7c70-0f04-11ce-b13f-00aa003bac6c    0x80000014
```

NOTE: The **show security alg msrpc object-id-map** CLI command has a chassis cluster node option to permit the output to be restricted to a particular node or to query the entire cluster. The **show security alg msrpc object-id-map node** CLI command options are **<node-id | all | local | primary>**.

Related Documentation

- Enabling Microsoft RPC ALGs (J-Web Procedure) on page 70
- Enabling Microsoft RPC ALGs (CLI Procedure) on page 70
- Customizing Microsoft RPC Applications (CLI Procedure) on page 73

Enabling Microsoft RPC ALGs (CLI Procedure)

The MS-RPC ALG is enabled by default and requires no configuration.

To disable the Microsoft RPC ALG, enter the following command:

```
user@host# set security alg msrpc disable
```

To reenable the Microsoft RPC ALG, enter the following command:

```
user@host# delete security alg msrpc
```

Related Documentation

- Understanding Microsoft RPC ALGs on page 67
- Enabling Microsoft RPC ALGs (J-Web Procedure) on page 70
- Verifying the Microsoft RPC ALG Tables on page 71

Enabling Microsoft RPC ALGs (J-Web Procedure)

The MS-RPC ALG is enabled by default and requires no configuration.

To disable or reenable the Microsoft ALG:

1. Select **Configure>Security>ALG**.
2. Select the **Enable MSRPC** check box.
3. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options>Commit**.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Microsoft RPC ALGs on page 67
- Enabling Microsoft RPC ALGs (CLI Procedure) on page 70
- Verifying the Microsoft RPC ALG Tables on page 71

Verifying the Microsoft RPC ALG Tables

Purpose To verify the Microsoft RPC ALG, display the Microsoft universal unique identifier to Object ID (UUID-to-OID) mapping table. The Microsoft RPC ALG monitors packets on TCP port 135.

Action From the CLI, enter the **show security alg msrpc object-id-map** command.

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
1be617c0-31a5-11cf-a7d8-00805f48a135    0x80000020
e3514235-4b06-11d1-ab04-00c04fc2dcd2    0x80000002
67df7c70-0f04-11ce-b13f-00aa003bac6c    0x80000014
```

NOTE: The **show security alg msrpc object-id-map** CLI command has a chassis cluster node option to permit the output to be restricted to a particular node or to query the entire cluster. The **show security alg msrpc object-id-map node** CLI command options are **<node-id | all | local | primary>**.

Related Documentation

- Enabling Microsoft RPC ALGs (J-Web Procedure) on page 70
- Enabling Microsoft RPC ALGs (CLI Procedure) on page 70
- Customizing Microsoft RPC Applications (CLI Procedure) on page 73

Understanding Microsoft RPC Services

MS-RPC is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun RPC, the MS-RPC provides a way for a program running on one host to call procedures in a program running on another host. The MS-RPC is dynamically negotiated based on the service program's universal unique identifier (UUID). The specific UUID is mapped to a transport address.

Predefined Microsoft remote procedure call (MS-RPC) services include:

- **junos-ms-rpc-epm**
- **junos-ms-rpc-tcp**
- **junos-ms-rpc-udp**

Copyright © 2016, Juniper Networks, Inc.

MS-RPC application defaults include:

- junos-ms-rpc-iis-com-1
- junos-ms-rpc-iis-com-adminbase
- junos-ms-rpc-msexchange-directory-nsp
- junos-ms-rpc-msexchange-directory-rfr
- junos-ms-rpc-msexchange-info-store
- junos-ms-rpc-uuid-any-tcp
- junos-ms-rpc-uuid-any-udp
- junos-ms-rpc-wmic-admin
- junos-ms-rpc-wmic-admin2
- junos-ms-rpc-wmic-mgmt
- junos-ms-rpc-wmic-webm-callresult
- junos-ms-rpc-wmic-webm-classobject
- junos-ms-rpc-wmic-webm-level1login
- junos-ms-rpc-wmic-webm-login-clientid
- junos-ms-rpc-wmic-webm-login-helper
- junos-ms-rpc-wmic-webm-objectsink
- junos-ms-rpc-wmic-webm-refreshing-services
- junos-ms-rpc-wmic-webm-remote-refresher
- junos-ms-rpc-wmic-webm-services
- junos-ms-rpc-wmic-webm-shutdown

MS-RPC application-set defaults include:

- junos-ms-rpc
- junos-ms-rpc-any
- junos-ms-rpc-iis-com
- junos-ms-rpc-msexchange
- junos-ms-rpc-wmic

Table 4 lists predefined MS-RPC services, UUID values associated with each service, and a description of each service.

Copyright © 2016, Juniper Networks, Inc.

Table 4: Predefined MS-RPC services

Service	UUID	Description
EPM	e1af8308-5d1f-11c9-91a4-08002b14a0fa	MS-RPC Endpoint Mapper (EPM) protocol is a TCP/UDP port-based service that includes TCP/UDP port 135.
EXCHANGE-DATABASE	1a190310-bb9c-11cd-90f8-00aa00466520	Microsoft Exchange Database service.
EXCHANGE-DIRECTORY	f5cc5a18-4264-101a-8c59-08002b2f8426 f5cc5a7c-4264-101a-8c59-08002b2f8426 f5cc59b4-4264-101a-8c59-08002b2f8426	Microsoft Exchange Directory service.
WIN-DNS	50abc2a4-574d-40b3-9d66-ee4fd5fba076	Microsoft Windows DNS server.
WINS	5f52c28-7f9f-101a-b52b-08002b2efabe 811109bf-a4e1-11d1-ab54-00a0c91e9b45	Microsoft WINS service.
WMIC-Webm-Level1Login	f309ad18-d86a-11d0-a075-00c04fb68820	This service allows users to connect to the management services interface in a particular namespace.

Related Documentation

- Understanding Microsoft RPC ALGs on page 67
- Customizing Microsoft RPC Applications (CLI Procedure) on page 73
- Understanding Sun RPC Services on page 64

Customizing Microsoft RPC Applications (CLI Procedure)

MS-RPC applications are customized in the same way as SUN RPC applications.

MS-RPC services in security policies are:

- **0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde**
- **1453c42c-0fa6-11d2-a910-00c04f990f3b**
- **10f24e8e-0fa6-11d2-a910-00c04f990f3b**
- **1544f5e0-613c-11d1-93df-00c04fd7bd09**

The corresponding TCP/UDP ports are dynamic. To permit them, you use the following statement for each number:

set applications *application-name* term *term-name* uuid *hex-number*

Copyright © 2016, Juniper Networks, Inc.

The ALG maps the program numbers into dynamically negotiated TCP/UDP ports based on these four UUIDs and permits or denies the service based on a policy you configure.

Related Documentation

- Understanding Microsoft RPC Services on page 71
- Customizing Sun RPC Applications (CLI Procedure) on page 64
- Verifying the Microsoft RPC ALG Tables on page 71

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 8

Configuring the RSH ALG

- Understanding the RSH ALG on page 75
- Example: Configuring the RSH ALG on page 75

Understanding the RSH ALG

The Remote Shell (RSH) Application Layer Gateway (ALG) processes RSH packets that initiate requests and open two gates to allow return packets from the reverse direction to the client. One gate is used for an identification (ident) session to apply authorization and the other gate is used for a standard error (stderr) session to transfer an error message.

NOTE: The RSH ALG does not work if Port Address Translation (PAT) is configured. The RSH requires the port range to be between 512 to 1024. The source NAT module cannot match this port range.

Related Documentation

- Example: Configuring the RSH ALG on page 75

Example: Configuring the RSH ALG

This example shows how to configure the RSH ALG in route or NAT mode. The configuration allows RSH traffic to pass through a device, and it transfers remote commands and results between a client and a server located on opposite sides of a Juniper Networks device.

- Requirements on page 75
- Overview on page 76
- Configuration on page 77
- Verification on page 85

Requirements

This example uses the following hardware and software components:

- An SRX Series device

Copyright © 2016, Juniper Networks, Inc.

- Two PCs (server and client)

Before you begin:

- Understand the concepts behind ALGs. See "ALG Overview" on page 3
- Understand the basics of RSH ALG. See the "Understanding the RSH ALG" on page 75

Overview

In this example, first you configure network interfaces on the device. Create security zones and assign interfaces to the zones, and configure a policy to allow RSH traffic to go through an SRX Series device.

Then you create a static NAT rule set rs1 with a rule r1 to match with the destination address 40.0.172.10/32, and you create a static NAT prefix with address 40.0.172.45/32.

Next you create a source NAT pool src-p1 with a source rule set src-rs1 to translate packets from interface fe-3/0/0.0 to interface fe-3/0/1.0. For matching packets, the source address is translated to an IP address in the src-p1 pool.

Then you create a destination NAT pool des-p1 with a destination rule set des-rs1 to translate packets from zone trust to destination address 40.0.172.10/32. For matching packets, the destination address is translated to an IP address in the des-p1 pool. Finally, you enable RSH ALG trace options.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 4 shows the RSH ALG topology.

Figure 4: RSH ALG Topology

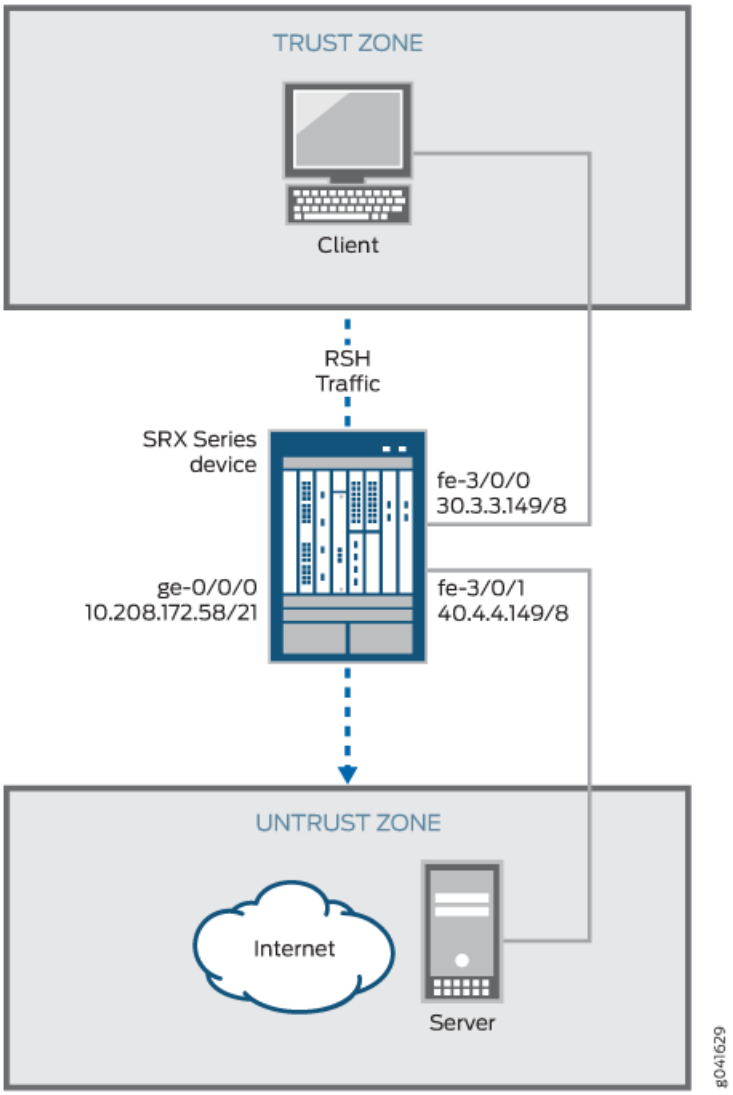

Configuration

To configure the RSH ALG, perform these tasks:

- Configuring a Route Mode on page 78
- Configuring a Static NAT Rule Set on page 80
- Configuring a Source NAT Pool and Rule Set without PAT on page 81
- Configuring a Destination NAT Pool and Rule Set on page 83
- Enabling RSH ALG Trace Options on page 84

Copyright © 2016, Juniper Networks, Inc.

Configuring a Route Mode

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.208.172.58/21
set interfaces fe-3/0/0 unit 0 family inet address 30.3.3.149/8
set interfaces fe-3/0/1 unit 0 family inet address 40.4.4.149/8
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-3/0/0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-3/0/1.0
set security policies from-zone trust to-zone untrust policy rsh match source-address any
set security policies from-zone trust to-zone untrust policy rsh match destination-address any
set security policies from-zone trust to-zone untrust policy rsh match application junos-rsh
set security policies from-zone trust to-zone untrust policy rsh then permit
```

Step-by-Step Procedure

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

To configure route mode:

1. Configure interfaces.

   ```
   [edit interfaces]
   user@host#set ge-0/0/0 unit 0 family inet address 10.208.172.58/21
   user@host#set fe-3/0/0 unit 0 family inet address 30.3.3.149/8
   user@host#set fe-3/0/1 unit 0 family inet address 40.4.4.149/8
   ```

2. Configure zones and assign interfaces to the zones.

   ```
   [edit security zones security-zone]
   user@host#set trust host-inbound-traffic system-services all
   user@host#set trust host-inbound-traffic protocols all
   user@host#set trust interfaces fe-3/0/0.0
   user@host#set untrust host-inbound-traffic system-services all
   user@host#set untrust host-inbound-traffic protocols all
   user@host#set untrust interfaces fe-3/0/0.1
   ```

3. Configure an RSH policy that allows RSH traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host#set policy rsh match source-address any
   user@host#set policy rsh match destination-address any
   user@host#set policy rsh match application junos-rsh
   user@host#set policy rsh then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show security policies** commands. If the output does not display

Copyright © 2016, Juniper Networks, Inc.

the intended configuration, repeat the configuration instructions in this example for correction.

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

```
[edit]
user@host# show interfaces
ge-0/0/0 {
   unit 0 {
      family inet {
         address 10.208.172.58/21;
      }
   }
}
fe-3/0/0 {
   unit 0 {
      family inet {
         address 30.3.3.149/8;
      }
   }
}
fe-3/0/1 {
   unit 0 {
      family inet {
         address 40.4.4.149/8;
      }
   }
}
```

```
[edit]
user@host# show security zones
..
   security-zone trust {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
      fe-3/0/0.0;
      }
}
security-zone untrust {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
```

```
        fe-3/0/1.0;
        }
}
...
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy rsh {
    match {
      source-address any;
      destination-address any;
      application junos-rsh;
    }
    then {
      permit;
    }
  }
}
```

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

Configuring a Static NAT Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone trust
set security nat static rule-set rs1 rule r1 match destination-address 40.0.172.10/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 40.0.172.45/32
```

Step-by-Step Procedure

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

To configure a static NAT rule set:

1. Create a static NAT rule set.

   ```
   [edit security nat static rule-set rs1]
   user@host#set from zone trust
   ```

2. Define the rule to match with the destination address.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 match destination-address 40.0.172.10/32
   ```

3. Define the static NAT prefix for the device.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 then static-nat prefix 40.0.172.45/32
   ```


Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
    from zone trust;
    rule r1 {
      match {
        destination-address 40.0.172.10/32;
      }
      then {
        static-nat {
          prefix {
            40.0.172.45/32;
          }
        }
      }
    }
  }
}
```

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

Configuring a Source NAT Pool and Rule Set without PAT

CLI Quick Configuration

NOTE: The RSH ALG does not support PAT configuration. The RSH ALG requires the stderr port range to be between 512 to 1024. The source NAT module cannot match this port range.

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

```
set security nat source pool src-p1 address 40.0.172.100/32 to 40.0.172.101/32
set security nat source pool src-p1 port no-translation
set security nat source rule-set src-rs1 from interface fe-3/0/0.0
set security nat source rule-set src-rs1 to interface fe-3/0/1.0
set security nat source rule-set src-rs1 rule r1 match source-address 30.0.0.0/8
set security nat source rule-set src-rs1 rule r1 match destination-address 40.0.0.0/8
set security nat source rule-set src-rs1 rule r1 then source-nat pool src-p1
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure a source NAT pool and rule set:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host#set pool src-p1 address 40.0.172.100/32 to 40.0.172.101/32
   ```

2. Create a source NAT pool with no port translation.

   ```
   [edit security nat source ]
   set pool src-p1 port no-translation
   ```

3. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 from interface fe-3/0/0.0
   user@host# set rule-set src-rs1 to interface fe-3/0/1.0
   ```

4. Configure a rule that matches packets and translates the source address to an address in the source pool.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 match source-address 30.0.0.0/8
   ```

5. Configure a rule that matches packets and translates the destination address to an address in the source pool.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 match destination-address 40.0.0.0/8
   ```

6. Configure a source NAT pool in the rule.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 then source-nat pool src-p1
   ```

Results

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

```
[edit]
user@host# show security nat
source {
   pool src-p1 {
      address {
         40.0.172.100/32 to 40.0.172.101/32;
      }
      port no-translation;
   }
   rule-set src-rs1 {
      from interface fe-3/0/0.0;
      to interface fe-3/0/1.0;
      rule r1 {
         match {
            source-address 30.0.0.0/8;
            destination-address 40.0.0.0/8;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          source-nat {
            pool {
              src-p1;
            }
          }
        }
      }
    }
  }
```

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

Configuring a Destination NAT Pool and Rule Set

CLI Quick Configuration

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

```
set security nat destination pool des-p1 address 40.0.172.45/32
set security nat destination rule-set des-rs1 from zone trust
set security nat destination rule-set des-rs1 rule des-r1 match source-address 30.0.172.12/32
set security nat destination rule-set des-rs1 rule des-r1 match destination-address 40.0.172.10/32
set security nat destination rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
```

Step-by-Step Procedure

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

To configure a destination NAT pool and rule set:

1. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host#set pool des-p1 address 40.0.172.45/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 from zone trust
   ```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match source-address 30.0.172.12/32
   ```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match destination-address 40.0.172.10/32
   ```

5. Configure a source NAT pool in the rule.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
   ```

Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show security nat
destination {
   pool des-p1 {
      address {
         40.0.172.45/32;
      }
   }
   rule-set des-rs1 {
      from zone trust;
      rule des-r1 {
         match {
            source-address 30.0.172.12/32;
            destination-address 40.0.172.10/32;
         }
         then {
            destination-nat {
               pool {
                  des--p1;
               }
            }
         }
      }
   }
}
```

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

Enabling RSH ALG Trace Options

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

```
set security alg rsh traceoptions flag all
set security alg traceoptions file trace
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
```

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

To enable RSH ALG trace options:

1. Enable RSH ALG trace options.

```
[edit security alg]
user@host#set sql traceoptions flag all
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure a filename to receive output from the tracing operation.

   ```
   [edit security alg]
   user@host#set traceoptions file trace
   ```

3. Specify the maximum trace file size.

   ```
   [edit security alg]
   user@host#set traceoptions file size 1g
   ```

4. Specify the level of tracing output.

   ```
   [edit security alg]
   user@host#set traceoptions level verbose
   ```

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

```
[edit]
user@host# show security alg
traceoptions {
   file trace size 1g;
   level verbose;
}
rsh traceoptions flag all;
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the RSH ALG Control Session on page 85
- Verifying the RSH ALG on page 86
- Verifying the RSH ALG Resource Manager Group on page 87
- Verifying the RSH ALG Resource Information on page 87

Verifying the RSH ALG Control Session

Purpose Verify that the RSH command is executed and all the RSH control and data sessions are created.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
Session ID: 2924, Policy name: rsh/6, Timeout: 2, Valid
Resource information : RSH ALG, 2, 0
  In: 30.0.172.12/1023 --> 40.0.172.45/514;tcp, If: fe-3/0/0.0, Pkts: 7, Bytes:
320
  Out: 40.0.172.45/514 --> 30.0.172.12/1023;tcp, If: fe-3/0/1.0, Pkts: 7, Bytes:
 314

Session ID: 2925, Policy name: rsh/6, Timeout: 2, Valid
Resource information : RSH ALG, 2, 24
  In: 40.0.172.45/44864 --> 30.0.172.12/113;tcp, If: fe-3/0/1.0, Pkts: 5, Bytes:
```

Copyright © 2016, Juniper Networks, Inc.

```
 278
  Out: 30.0.172.12/113 --> 40.0.172.45/44864;tcp, If: fe-3/0/0.0, Pkts: 5, Bytes:
 345

Session ID: 2926, Policy name: rsh/6, Timeout: 2, Valid
Resource information : RSH ALG, 2, 23
  In: 40.0.172.45/1023 --> 30.0.172.12/1022;tcp, If: fe-3/0/1.0, Pkts: 4, Bytes:
 216
  Out: 30.0.172.12/1022 --> 40.0.172.45/1023;tcp, If: fe-3/0/0.0, Pkts: 3, Bytes:
 164
Total sessions: 3
```

Meaning

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name, number of packets in and out.
- **Policy name**—Policy name that permitted the traffic.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and source interface for this session is fe-3/0/0.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and destination interface for this session is fe-3/0/1.0).

Verifying the RSH ALG

Purpose Verify that the RSH ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host>show security alg status
ALG Status :
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

NOTE: The RSH ALG is disabled by default. To enable the RSH ALG, enter the **set security alg rsh** command in the configuration mode.

Meaning The output shows the RSH ALG status as follows:

- Enabled—Shows the RSH ALG is enabled.
- Disabled—Shows the RSH ALG is disabled.

Copyright © 2016, Juniper Networks, Inc.

Verifying the RSH ALG Resource Manager Group

Purpose Verify the total number of resource manager groups and active groups that are used by the RSH ALG.

Action From operational mode, enter the **show security resource-manager group active** command.

```
user@host>show security resource-manager group active
Group ID 1: Application - RSH ALG
    Total groups 677, active groups 1
```

Verifying the RSH ALG Resource Information

Purpose Verify the total number of resources and active resources that are used by the RSH ALG.

Action From operational mode, enter the **show security resource-manager resource active** command.

```
user@host>show security resource-manager resource active
Resource ID 2: Group ID - 1, Application - RSH ALG

      Resource ID 1: Group ID - 1, Application - RSH ALG
     Total Resources 4044, active resources 2
```

Related Documentation

- Understanding the RSH ALG on page 75

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 9

Configuring the RTSP ALG

- Understanding the RTSP ALG on page 89
- Understanding IPv6 Support for RTSP ALG on page 91
- Understanding RTSP ALG Messages on page 91
- Understanding RTSP ALG Conversation and NAT on page 93
- Example: Configuring the RTSP ALG on page 96

Understanding the RTSP ALG

- Overview on page 89
- RTSP Modes on page 90

Overview

RTSP (Real-Time Streaming Protocol) is an Application Layer protocol for controlling the delivery of data with real-time properties. It is similar in syntax and operation to HTTP/1.1. Unlike SIP and H.323, the purpose of RTSP is to access existing media files over the network and to control the replay of the media. The typical communication is between a client (running RealPlayer for example) and a streaming media server. Commands include the ability to pause and play media files from the remote server.

RTSP is a control channel protocol between the media client and media server. The data channel uses a different protocol, usually Real-Time Transport Protocol (RTP) or RTP Control Protocol (RTCP).

In RTSP standard mode, the client sets up three network channels with the RTSP server when media data is delivered using RTP over UDP.

RTSP runs over TCP. RTP and RTCP run over UDP. The ports for RTP and RTCP packets are dynamically negotiated by the client and server using RTSP. Because RTP and RTCP ports are dynamic, these ports cannot be allowed by a static policy. The main purpose of introducing an RTSP ALG to a firewall is to create dynamic policy (pinhole) according to the result of client/server negotiation so that RTP and RTCP traffic can pass through.

When the client and server reside in different realms, they might not be able to determine how to route to the address of the RTP or RTCP offer given by the peer. In this case, ALG needs to be involved to do translation for the RTP or RTCP offer address and modify it in the payload.

Copyright © 2016, Juniper Networks, Inc.

After the connection is established, the RTSP ALG monitors the messages exchanged between the client and server, tracks the status change of the dialog, and returns all the resources it acquired to support an RTSP dialog back to the system after the dialog has completed or failed.

RTSP Modes

Standard Mode

In RTSP standard mode, the client sets up three network channels with the RTSP server when media data is delivered using RTP over UDP.

A full-duplex TCP connection is used for control and negotiation. A full-duplex UDP channel is used for media data delivery using the RTP packet format. In most cases, RTP is initiated from the server. A full-duplex UDP channel called RTCP is used to provide synchronization information to the client and packet loss information to the server.

Figure 5 shows the RTSP ALG standard mode.

Figure 5: RTSP ALG Standard Mode

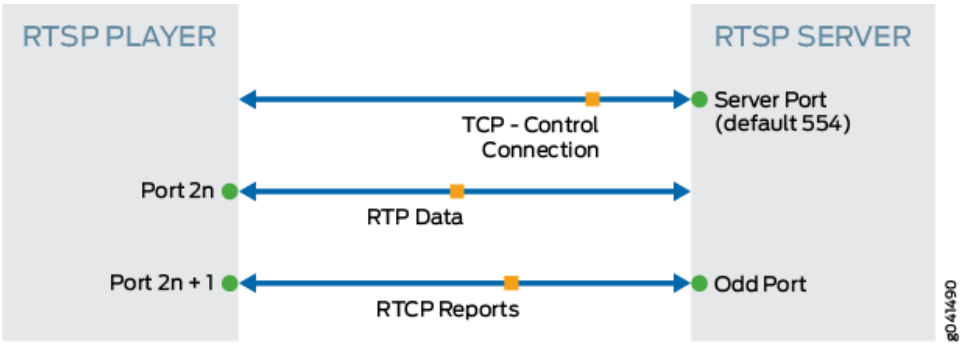

Interleave Mode

In RTSP interleave mode, media data can be made into packets using RTP or RDT over TCP. In this scenario , a single full-duplex TCP connection is used for both control and for media data delivery from the RTSP server to the client. The data stream is interleaved with the RTSP control stream.

Figure 6 shows the RTSP ALG interleave mode.

Figure 6: RTSP ALG Interleave Mode

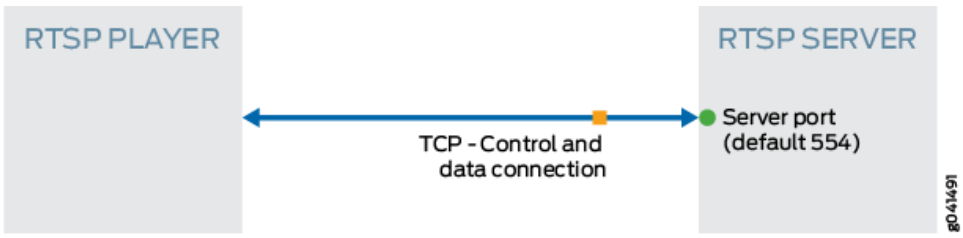

Related Documentation

- Understanding RTSP ALG Messages on page 91

Copyright © 2016, Juniper Networks, Inc.

- Understanding RTSP ALG Conversation and NAT on page 93
- Example: Configuring the RTSP ALG on page 96

Understanding IPv6 Support for RTSP ALG

IPv6 is supported on the RTSP ALG along with NAT-PT mode and NAT64 address translation.

This feature enables the RTSP ALG to parse IPv6 RTSP packets, open an IPv6 pattern pinhole, and translate the Layer 7 IPv6 address according to the NAT configuration. Also, support for IPv6 RTSP transaction pass through under permission policy and IPv6 RTSP transaction pass through under NAT-PT and NAT 64 are enabled.

The RTSP ALG with IPv6 support has the following limitations:

- Real-Time Streaming Protocol (RTSP) is an Application Layer protocol for controlling the delivery of data with real-time properties. The RTSP ALG supports a peer client, and the server transmits real-time media; it does not support third-party endpoints involved in the transaction.
- In case of destination NAT or NAT64 for IP address translation, if the RTSP message (including the Session Description Protocol [SDP] application content) length exceeds 2500 bytes, then the RTSP ALG processes only the first 2500 bytes of the message and ignores the rest of the message. In this scenario, the IP address in the RTSP message is not translated if the IP address does not appear in the first 2500 bytes.

Related Documentation

- Understanding the RTSP ALG on page 89

Understanding RTSP ALG Messages

- RTSP Messages Format on page 91
- RTSP Methods on page 92
- RTSP Status Code on page 92
- RTSP Header on page 93

RTSP Messages Format

RTSP is text based and uses the ISO 10646 character set in UTF-8 encoding. Lines are terminated by CRLF, and an empty line is the separator of the message and body.

The first line is called the start-line. For request messages from client to server, the start-line represents the RTSP method. For the response message from server to client, the start-line represents the RTSP status code as the reply of method. The status code element is a 3-digit integer result code.

Copyright © 2016, Juniper Networks, Inc.

RTSP Methods

There are nine types of methods during one transaction.

- **OPTION**—Represents a request for information about the communication options available on the request/response chain identified by the Request-URL. This method allows the client to determine the options, requirements, or both associated with a resource, or the capabilities of a server, without implying a resource action or initiating a resource retrieval.
- **DESCRIBE**—Retrieves the description of a presentation or media object identified by the request URL from a server. This method might use the Accept header to specify the description formats that the client interprets.
- **ANNOUNCE**—Request sent from client to server, this method posts the description of a presentation or media object identified by the request URL to a server. When request sent from server to client, this method updates the session description in real-time.
- **SETUP**—Requests a URI and specifies the transport mechanism to be used for the streamed media.
- **PLAY**—Informs the server to start sending data using the mechanism specified in SETUP.
- **PAUSE**—Requests the stream delivery to be interrupted temporarily.
- **TEARDOWN**—Stops the stream delivery for the given URI, freeing the resource associated with it.
- **GET_PARAMETER**—Retrieves the value of a parameter of a presentation or stream specified in the URI.
- **SET_PARAMETER**—Sets the value of a parameter for a presentation or stream specified by the URI.

RTSP Status Code

The first digit of the status code defines the class of response.

- 1**: Informational—Request received, continuing process.
- 2**: Success
- 3**: Redirection—Further action must be taken in order to complete the request.
- 4**: Client Error—The request contains bad syntax or cannot be fulfilled.
- 5**: Server Error—The server failed to fulfill an apparently valid request.

Copyright © 2016, Juniper Networks, Inc.

RTSP Header

The RTSP header consists of the following fields:

- **CSeq**—Specifies the sequence number for an RTSP request-response pair. For every RTSP request containing the given sequence number, there will be a corresponding response having the same number.
- **Content-Length**—Contains the length of the content of the method, that is, after the double CRLF following the last header.
- **TRANSPORT**—Indicates which transport protocol is to be used and configures its parameters.
- **SESSION**—Identifies an RTSP session started by the media server in a SETUP response and concluded by TEARDOWN on the presentation URL.

Related Documentation

- Understanding the RTSP ALG on page 89
- Understanding RTSP ALG Conversation and NAT on page 93
- Example: Configuring the RTSP ALG on page 96

Understanding RTSP ALG Conversation and NAT

This topic provides details on typical RTSP ALG conversation.

In general, RTP and RTCP packets are bidirectional, which means that either the client or server could initiate an RTP or an RTCP session.

Figure 7 describes an example of a sample packet capture in a standard RTSP conversation.

Copyright © 2016, Juniper Networks, Inc.

Figure 7: RTSP ALG Conversation

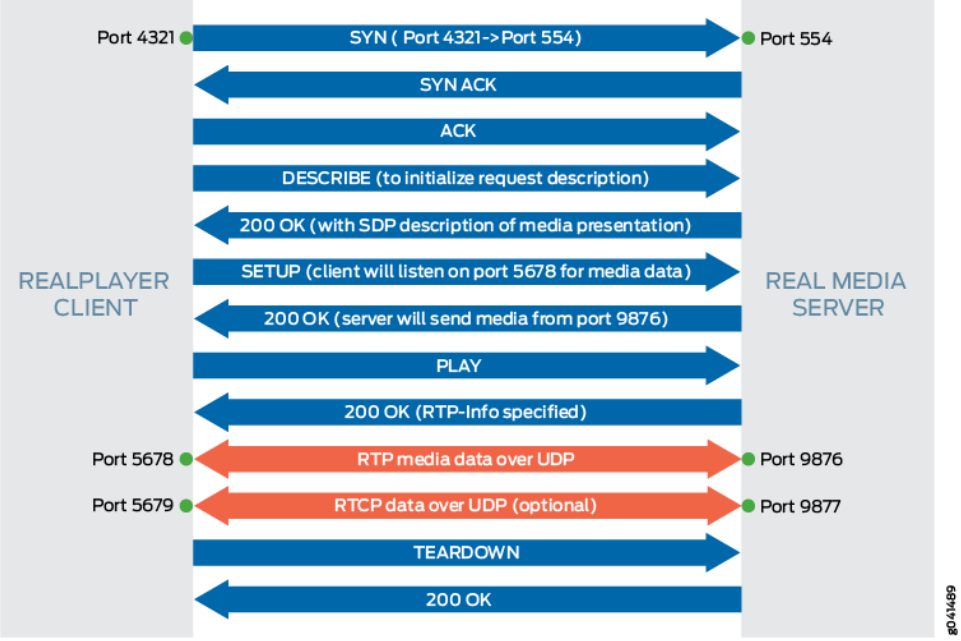

The RTSP ALG performs the following actions for a RTSP sample packet capture in a standard RTSP conversation:

1. Monitors SETUP and 200 OK messages.
2. Receives negotiated ports (6543 and 8765 in this example)
3. Opens a pinhole for UDP media data from server to client.
4. Receives the IP address in payload and translates the address if NAT is required.

Table 5 describes the RTSP payload IP NAT.

Table 5: RTSP Payload IP NAT

	Forward(C->S)	Reverse(S->C)	Pinhole	Payload IP Translate	Payload Port Translate
No NAT	A/4321->B/554	A/4321<-B/554	B/9876->A/5678 A/5678->B/9876	N/A	N/A
Source NAT (IPvx)	A/4321->B/554	A'/P'<-B/554	B/9876->A'/P'' A/5678->B/9876	N/A (*)	5678<->P''
Destination NAT (IPvx)	A/4321->B'/554	A/4321<-B/554	B/9876->A/5678 A/5678->B'/9876	B' -> B (**)	N/A

Copyright © 2016, Juniper Networks, Inc.

Table 5: RTSP Payload IP NAT *(continued)*

	Forward(C->S)	Reverse(S->C)	Pinhole	Payload IP Translate	Payload Port Translate
NAT64	A/4321->B''/554	A''/Q'<-B/554	B/9876->A''/Q'' A/5678->B''/9876	B''(IPv6)->B(IPv4)	5678<->Q''
NAT46	A/4321->B'''/554	A'''/R'<-B/554	B/9876->A'''/R'' A/5678->B'''/9876	B'''(IPv4)->B(IPv6)	5678<->R''

In Table 5, the following letters and symbols are used:

- A—RTSP client IP address
- A'—Translated IPv4 or IPv6 address of RTSP client
- A''—Translated IPv4 address
- A'''—Translated IPv6 address
- B—RTSP server IP address
- B'—RTSP server IP address before destination NAT
- B''—RTSP server IP address at IPv6 realm
- B'''— RTSP server IP address at IPv4 realm
- P'—Translated Port(translates from 4321) of RTSP client
- P''—Translated Port(translates from 5678 in message payload) of RTSP client
- Q'—Translated (IPv6 to IPv4) Port(translates from 4321) of RTSP client
- Q''—Translated (IPv6 to IPv4) Port (translates from 5678 in message payload) of RTSP client
- R'—Translated (IPv4 to IPv6) Port (translates from 4321) of RTSP client
- R''—Translated (IPv4 to IPv6) Port (translates from 5678 in message payload) of RTSP client
- (*)—RTSP server IP address B appears in payload message; it does not need to translate
- (**)—IP address B' appears in payload message from client to server; it needs to translate to B

Related Documentation

- Understanding the RTSP ALG on page 89
- Understanding RTSP ALG Messages on page 91
- Example: Configuring the RTSP ALG on page 96

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the RTSP ALG

This example shows how to configure the RTSP ALG to pass through RTSP traffic with a source NAT pool on Juniper Networks devices.

- Requirements on page 96
- Overview on page 96
- Configuration on page 96
- Verification on page 99

Requirements

- Configure proxy ARP for all IP addresses in the source NAT pool.
- Enable the RTSP ALG.
- Understand the basics concepts of the RTSP ALG. See "Understanding the RTSP ALG" on page 89.

Overview

In this example, the RTSP ALG is configured to monitor and allow RTSP traffic transferring media between client and server located on opposite sides of a Juniper Networks device.

Configuration

- Enabling RTSP ALG on page 96
- Configuring a NAT Source Pool and Rule Set and a Policy on page 96
- Configuring RTSP ALG trace options on page 98

Enabling RTSP ALG

CLI Quick Configuration

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

To configure proxy ARP for all IP addresses in the source NAT pool and to enable RTSP ALG:

```
set security nat proxy-arp interface <interface-name> address 10.10.10.1/32 to 10.10.10.10/32
set security alg rtsp
```

Enter **commit** from configuration mode.

Configuring a NAT Source Pool and Rule Set and a Policy

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

Copyright © 2016, Juniper Networks, Inc.

network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
set security zones security-zone green address-book address sa1 1.1.1.0/24
set security zones security-zone red address-book address da1 2.2.2.0/24
set security nat source rule-set rs1 from zone green
set security nat source rule-set rs1 to zone red
set security nat source rule-set rs1 rule r1 match source-address 1.1.1.0/24
set security nat source rule-set rs1 rule r1 match destination-address 2.2.2.0/24
set security nat source rule-set rs1 rule r1 then source-nat pool pool1

set security policy from-zone green to-zone red policy pol1 match destination-address da1
set security policy from-zone green to-zone red policy pol1 match source-address sa1
set security policy from-zone green to-zone red policy pol1 match application junos-rtsp
set security policy from-zone green to-zone red policy pol1 then permit
```

Enter **commit** from configuration mode.

NOTE: If you are not sure of the RTSP client and server IP address, you can replace “da1” and “sa1” with “any”.

Step-by-Step Procedure

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

To configure a source NAT pool:

1. Create a NAT source pool.

   ```
   [edit security]
   user@host# set nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
   ```

2. Configure security zone address book entries.

   ```
   [edit security zones security-zone]
   user@host# set green address-book address sa1 1.1.1.0/24
   user@host# set red address-book address da1 2.2.2.0/24
   ```

3. Create a NAT source rule set.

   ```
   [edit security nat source rule-set rs1]
   user@host# set from zone green
   user@host# set to zone red
   user@host# set rule r1 match source-address 1.1.1.0/24
   user@host# set rule r1 match destination-address 2.2.2.0/24
   user@host# set rule r1 then source-nat pool pool1
   ```

4. Configure a policy.

   ```
   [edit security policies from-zone green to-zone red policy pol1]
   user@host# set match source-address sa1
   user@host# set match destination-address da1
   user@host# set match application junos-rtsp
   user@host# set then permit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit ]
user@host# show security nat
source {
  pool pool1 {
    address {
      10.10.10.1/32 to 10.10.10.10/32;
    }
  }
  rule-set rs1 {
    from zone green;
    to zone red;
    rule r1 {
      match {
        source-address 1.1.1.0/24;
        destination-address 2.2.2.0/24;
      }
      then {
        source-nat {
          pool {
            pool1;
          }
        }
      }
    }
  }
```

```
[edit]
user@host# show security policies
from-zone green to-zone red {policy pol1 {
  policy pol1 {
    match {
      source-address sa1;
      destination-address da1;
      application [junos-rtsp];
    }
    then {
      permit;
    }
  }
}
default-policy {
  permit-all;
}
```

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

Configuring RTSP ALG trace options

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security alg rtsp traceoptions flag all
set security alg traceoptions file trace
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
```

Step-by-Step Procedure

To configure RTSP ALG trace options:

1. Enable RTSP ALG trace options.

   ```
   [edit security alg]
   user@host# set rtsp traceoptions flag all
   ```

2. Configure a filename to receive output from the tracing operation.

   ```
   [edit security alg]
   user@host# set traceoptions file trace
   ```

3. Specify the maximum trace file size.

   ```
   [edit security alg]
   user@host# set traceoptions file size 1g
   ```

4. Specify the level of tracing output.

   ```
   [edit security alg]
   user@host# set traceoptions level verbose
   ```

Results

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

```
[edit]
user@host# show security alg
traceoptions {
   file trace size 1g;
   level verbose;
   }
rtsp traceoptions flag all;
```

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

Verification

Confirm that the configuration is working properly.

- Verifying RTSP ALG on page 100
- Verifying the RTSP ALG Control Session on page 100
- Verifying the RTSP ALG Flow Gate Information on page 101
- Verifying the RTSP Resource Manager Group on page 102
- Verifying the RTSP Resource Information on page 102

Copyright © 2016, Juniper Networks, Inc.

Verifying RTSP ALG

Purpose Verify that the RTSP ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status

DNS      : Enabled
FTP      : Enabled
H323     : Enabled
RTSP     : Enabled
```

Meaning The output shows the RTSP ALG status as follows:

- Enabled—Shows the RTSP ALG is enabled.
- Disabled—Shows the RTSP ALG is disabled.

Verifying the RTSP ALG Control Session

Purpose Verify that the control session is created and all the RTSP control and data sessions are created.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
Flow Sessions on FPC5 PIC0:

Session ID: 100004087, Policy name: dns-alg/4, Timeout: 1798, Valid
Resource information : RTSP ALG, 1, 0
  In: 1.1.0.100/59889 --> 1.1.0.202/554;tcp, If: ge-0/0/1.0, Pkts: 28, Bytes:
7618
  Out: 1.1.0.202/554 --> 1.1.0.100/59889;tcp, If: ge-0/0/2.0, Pkts: 27, Bytes:
24304

Session ID: 100004088, Policy name: dns-alg/4, Timeout: 120, Valid
Resource information : RTSP ALG, 1, 1
  In: 1.1.0.202/5004 --> 1.1.0.100/62092;udp, If: ge-0/0/2.0, Pkts: 19, Bytes:
17013
  Out: 1.1.0.100/62092 --> 1.1.0.202/5004;udp, If: ge-0/0/1.0, Pkts: 0, Bytes: 0

Session ID: 100004089, Policy name: dns-alg/4, Timeout: 120, Valid
Resource information : RTSP ALG, 1, 4
  In: 1.1.0.202/5004 --> 1.1.0.100/62094;udp, If: ge-0/0/2.0, Pkts: 433, Bytes:
346183
  Out: 1.1.0.100/62094 --> 1.1.0.202/5004;udp, If: ge-0/0/1.0, Pkts: 0, Bytes: 0

Session ID: 100004090, Policy name: dns-alg/4, Timeout: 120, Valid
Resource information : RTSP ALG, 1, 3
  In: 1.1.0.100/62093 --> 1.1.0.202/5005;udp, If: ge-0/0/1.0, Pkts: 2, Bytes: 260

  Out: 1.1.0.202/5005 --> 1.1.0.100/62093;udp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0
Total sessions: 4
```


Copyright © 2016, Juniper Networks, Inc.

Meaning

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **Policy name**—Policy name that permitted the traffic.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and the source interface for this session is ge-0/0/1.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and destination interface for this session is fe-0/0/2.0).

Verifying the RTSP ALG Flow Gate Information

Purpose Verify that the flow gate is opened for TCP data channel connection.

Action From operational mode, enter the **show security flow gate** command.

```
user@host>show security flow gate

 Flow Gates on FPC5 PIC0:

Hole: 1.1.0.202-1.1.0.202/5005-5005->1.1.0.100-1.1.0.100/62093-62093
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: udp
  Application: RTSP ALG/11
  Age: 32 seconds
  Flags: 0x0080
  Zone: untrust
  Reference count: 1
  Resource: 4-1-2

Hole: 1.1.0.100-1.1.0.100/62093-62093->1.1.0.202-1.1.0.202/5005-5005
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: udp
  Application: RTSP ALG/11
  Age: 32 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 4-1-3

Hole: 1.1.0.202-1.1.0.202/5004-5004->1.1.0.100-1.1.0.100/62094-62094
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: udp
  Application: RTSP ALG/11
  Age: 32 seconds
  Flags: 0x0080
  Zone: untrust
  Reference count: 1
  Resource: 4-1-4

Hole: 1.1.0.100-1.1.0.100/62094-62094->1.1.0.202-1.1.0.202/5004-5004
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: udp
  Application: RTSP ALG/11
  Age: 32 seconds
  Flags: 0x0080
```

Copyright © 2016, Juniper Networks, Inc.

```
  Zone: trust
  Reference count: 1
  Resource: 4-1-5

Valid gates: 4
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 4
```

Meaning The sample output shows that the flow gate is opened for TCP data channel connection.

Verifying the RTSP Resource Manager Group

Purpose Verify the total number of resource manager groups and active groups that are used by the RTSP ALG.

Action From operational mode, enter the **show security resource-manager group active** command.

```
user@host>show security resource-manager group active
Group ID 1: Application - RTSP ALG
       Total groups 19763, active groups 1
```

Meaning The sample output shows the total number of resource manager groups and active groups that are used by the RTSP ALG.

Verifying the RTSP Resource Information

Purpose Verify the total number of resources and active resources that are used by the RTSP ALG.

Action From operational mode, enter the **show security resource-manager resource active** command.

```
user@host>show security resource-manager resource active
Resource ID 2: Group ID - 1, Application - RTSP ALG

       Resource ID 1: Group ID - 1, Application - RTSP ALG
       Total Resources 93286, active resources 2
```

Meaning The sample output shows the total number of resources and active resources that are used by the RTSP ALG.

Related Documentation

- ALG Overview on page 3
- Understanding the RTSP ALG on page 89
- Understanding RTSP ALG Messages on page 91
- Understanding RTSP ALG Conversation and NAT on page 93

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 10

Configuring the SQLNET ALG

- Understanding the SQLNET ALG on page 103
- Example: Configuring the SQLNET ALG on page 104

Understanding the SQLNET ALG

The SQLNET Application Layer Gateway (ALG) processes Transparent Network Substrate (TNS) REDIRECT packets for IP addresses and port information. The SQLNET ALG performs Network Address Translation (NAT) on the payload of the TNS REDIRECT packet, opens a pinhole for a new connection from a client to a server, and transfers data between a client and a server located on opposite sides of a Juniper Networks device.

SQLNET ALG supports the following types of data transfer modes:

- Redirect mode — connect-redirect type
- Interleave mode — connect-accept type
- Load balance — connect-redirect-connect-redirect type

SQLNET allows remote data access between applications and the Oracle database, or among multiple Oracle databases. SQLNET primarily establishes and maintains connection between a client application and an Oracle database server. SQLNET has several communication layers that enable clients and database servers to share, modify, and manipulate data.

Oracle SQL servers use the SQLNET protocol to execute SQL commands from clients, including load balancing and application-specific services. The SQLNET protocol uses TNS as its networking architecture, and all SQLNET traffic is encapsulated into TNS packet format.

The SQLNET ALG monitors control packets, opens pinhole for data traffic, and performs NAT and port rewrites. Support of stateful firewall and NAT services are required to configure the SQLNET ALG for TCP port 1521.

Related Documentation

- Example: Configuring the SQLNET ALG on page 104

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the SQLNET ALG

The SQLNET ALG processes TNS REDIRECT packets, performs NAT, and opens a pinhole for a new connection from a client to a server.

This example shows how to configure the SQLNET ALG in route or NAT mode, allow SQLNET traffic to pass through a device, and transfer data between a client and a server located on opposite sides of a Juniper Networks device.

- Requirements on page 104
- Overview on page 104
- Configuration on page 105
- Verification on page 113

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Two PCs (client and server)

Before you begin:

- Understand the concepts behind ALGs. See "ALG Overview" on page 3.
- Understand the basics of SQLNET ALG . See "Understanding the SQLNET ALG" on page 103.

Overview

In this example, first you configure network interfaces on the device. Create security zones and assign interfaces to the zones, and configure a policy to allow SQLNET traffic to go through an SRX Series device.

Then you create a static NAT rule set rs1 with a rule r1 to match with the destination address 40.0.172.10/32, and you create a static NAT prefix with address 40.0.172.45/32.

Next you create a source NAT pool src-p1 with a source rule set src-rs1 to translate packets from interface fe-3/0/0.0 to interface fe-3/0/1.0. For matching packets, the source address is translated to an IP address in the src-p1 pool.

Then you create a destination NAT pool des-p1 with a destination rule set des-rs1 to translate packets from zone trust to destination address 40.0.172.10/32. For matching packets, the destination address is translated to an IP address in the des-p1 pool. Finally, you enable SQLNET ALG trace options.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 8 shows the SQLNET ALG topology.

Figure 8: SQLNET ALG Topology

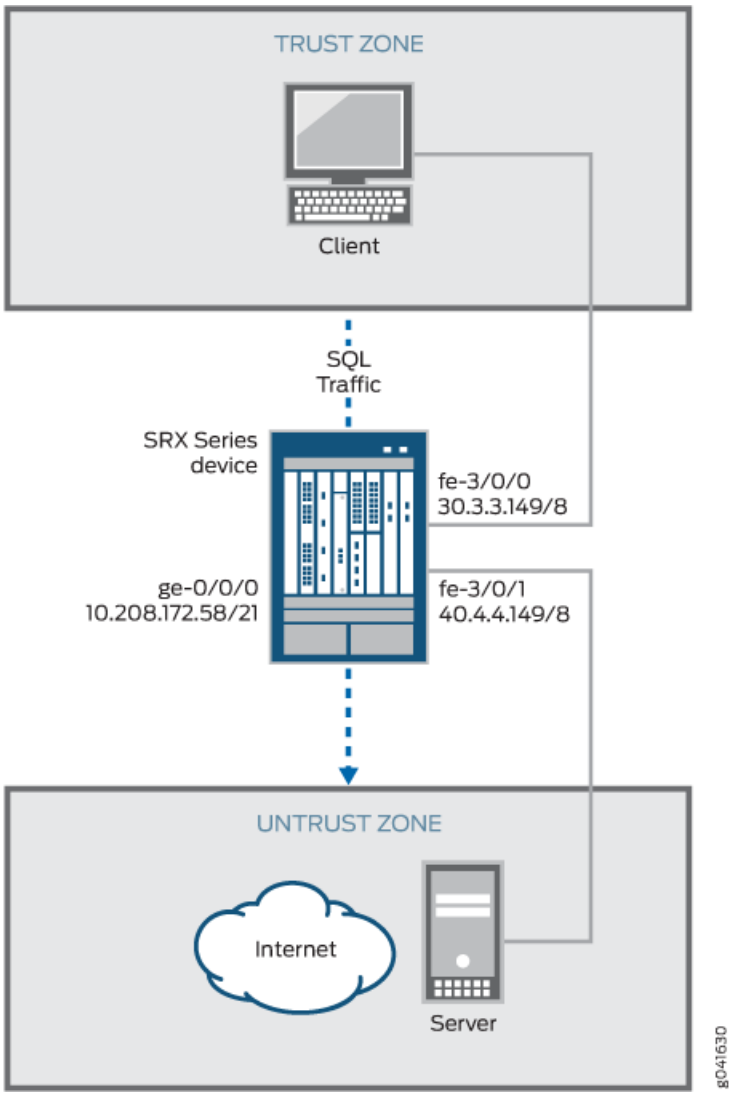

Configuration

To configure the SQLNET ALG, perform these tasks:

- Configuring a Route Mode on page 106
- Configuring a Static NAT Rule Set on page 108
- Configuring a Source NAT Pool and Rule Set on page 109
- Configuring a Destination NAT Pool and Rule Set on page 110
- Enabling SQLNET ALG Trace Options on page 112

Copyright © 2016, Juniper Networks, Inc.

Configuring a Route Mode

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.208.172.58/21
set interfaces fe-3/0/0 unit 0 family inet address 30.3.3.149/8
set interfaces fe-3/0/1 unit 0 family inet address 40.4.4.149/8
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-3/0/0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-3/0/1.0
set security policies from-zone trust to-zone untrust policy sql match source-address any
set security policies from-zone trust to-zone untrust policy sql match destination-address any
set security policies from-zone trust to-zone untrust policy sql match application junos-sqlnet-v2
set security policies from-zone trust to-zone untrust policy sql then permit
```

Step-by-Step Procedure

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

To configure route mode:

1. Configure interfaces.

   ```
   [edit interfaces]
   user@host# set ge-0/0/0 unit 0 family inet address 10.208.172.58/21
   user@host# set fe-3/0/0 unit 0 family inet address 30.3.3.149/8
   user@host# set fe-3/0/1 unit 0 family inet address 40.4.4.149/8
   ```

2. Configure zones and assign interfaces to the zones.

   ```
   [edit security zones security-zone]
   user@host# set trust host-inbound-traffic system-services all
   user@host# set trust host-inbound-traffic protocols all
   user@host# set trust interfaces fe-3/0/0.0
   user@host# set untrust host-inbound-traffic system-services all
   user@host# set untrust host-inbound-traffic protocols all
   user@host# set untrust interfaces fe-3/0/1.0
   ```

3. Configure a SQL policy that allows SQL traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust ]
   user@host# set policy sql match source-address any
   user@host# set policy sql match destination-address any
   user@host# set policy sql match application junos-sqlnet-v2
   user@host# set policy sql then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show security policies** commands. If the output does not display

Copyright © 2016, Juniper Networks, Inc.

the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.208.172.58/21;
        }
    }
}
fe-3/0/0 {
    unit 0 {
        family inet {
            address 30.3.3.149/8;
        }
    }
}
fe-3/0/1 {
    unit 0 {
        family inet {
            address 40.4.4.149/8;
        }
    }
}
```

```
[edit]
user@host# show security zones
...
    security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        fe-3/0/0.0;
        }
}
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
```

```
        fe-3/0/1.0;
        }
    }
    ...

    [edit]
    user@host# show security policies
    from-zone trust to-zone untrust {
      policy sql {
        match {
          source-address any;
          destination-address any;
          application junos-sqlnet-v2;
        }
        then {
          permit;
        }
      }
    }
```

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

Configuring a Static NAT Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone trust
set security nat static rule-set rs1 rule r1 match destination-address 40.0.172.10/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 40.0.172.45/32
```

Step-by-Step Procedure

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

To configure a static NAT rule set:

1. Create a static NAT rule set.

   ```
   [edit security nat static rule-set rs1]
   user@host#set from zone trust
   ```

2. Define a rule to match with the destination address.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 match destination-address 40.0.172.10/32
   ```

3. Define a static NAT prefix for the device.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 then static-nat prefix 40.0.172.45/32
   ```


Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
    from zone trust;
    rule r1 {
      match {
        destination-address 40.0.172.10/32;
      }
      then {
        static-nat {
          prefix {
            40.0.172.45/32;
          }
        }
      }
    }
  }
}
```

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

Configuring a Source NAT Pool and Rule Set

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

```
set security nat source pool src-p1 address 40.0.172.100/32 to 40.0.172.101/32
set security nat source rule-set src-rs1 from interface fe-3/0/0.0
set security nat source rule-set src-rs1 to interface fe-3/0/1.0
set security nat source rule-set src-rs1 rule r1 match source-address 30.0.0.0/8
set security nat source rule-set src-rs1 rule r1 match destination-address 40.0.0.0/8
set security nat source rule-set src-rs1 rule r1 then source-nat pool src-p1
```

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

To configure a source NAT pool and rule set:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host#set pool src-p1 address 40.0.172.100/32 to 40.0.172.101/32
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 from interface fe-3/0/0.0
   user@host# set rule-set src-rs1 to interface fe-3/0/1.0
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Configure a rule that matches packets and translates the source address to an address in the source pool.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 match source-address 30.0.0.0/8
   ```

4. Configure a rule that matches packets and translates the destination address to an address in the source pool.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 match destination-address 40.0.0.0/8
   ```

5. Configure a source NAT pool in the rule.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule r1 then source-nat pool src-p1
   ```

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

```
[edit]
user@host# show security nat
source {
  pool src-p1 {
    address {
      40.0.172.100/32 to 40.0.172.101/32;
    }
  }
  rule-set src-rs1 {
    from interface fe-3/0/0.0;
    to interface fe-3/0/1.0;
    rule r1 {
      match {
        source-address 30.0.0.0/8;
        destination-address 40.0.0.0/8;
      }
      then {
        source-nat {
          pool {
            src-p1;
          }
        }
      }
    }
  }
}
```

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

Configuring a Destination NAT Pool and Rule Set

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat destination pool des-p1 address 40.0.172.45/32
set security nat destination rule-set des-rs1 from zone trust
set security nat destination rule-set des-rs1 rule des-r1 match source-address 30.0.172.12/32
set security nat destination rule-set des-rs1 rule des-r1 match destination-address 40.0.172.10/32
set security nat destination rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
```

Step-by-Step Procedure

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

To configure a destination NAT pool and rule set:

1. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host#set pool des-p1 address 40.0.172.45/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 from zone trust
   ```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match source-address 30.0.172.12/32
   ```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match destination-address 40.0.172.10/32
   ```

5. Configure a source NAT pool in the rule.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
   ```

Results

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

```
[edit]
user@host# show security nat
destination {
   pool des-p1 {
      address {
         40.0.172.45/32;
      }
   }
   rule-set des-rs1 {
      from zone trust;
      rule des-r1 {
         match {
```

```
        source-address 30.0.172.12/32;
        destination-address 40.0.172.10/32;
      }
      then {
        destination-nat {
          pool {
            des-p1;
          }
        }
      }
    }
  }
```

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

Enabling SQLNET ALG Trace Options

CLI Quick Configuration

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

```
set security alg sql traceoptions flag all
set security alg traceoptions file trace
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
```

Step-by-Step Procedure

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

To enable SQLNET ALG trace options:

1. Enable SQLNET ALG trace options.

   ```
   [edit security alg]
   user@host#set sql traceoptions flag all
   ```

2. Configure a filename to receive output from the tracing operation.

   ```
   [edit security alg]
   user@host#set traceoptions file trace
   ```

3. Specify the maximum trace file size.

   ```
   [edit security alg]
   user@host#set traceoptions file size 1g
   ```

4. Specify the level of tracing output.

   ```
   [edit security alg]
   user@host#set traceoptions level verbose
   ```

Results

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

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security alg
  traceoptions {
    file trace size 1g;
    level verbose;
  }
sql traceoptions flag all;
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the SQLNET ALG Control Session on page 113
- Verifying the SQLNET ALG on page 114
- Verifying the SQLNET ALG Resource Manager Group on page 114
- Verifying the SQLNET ALG Resource Information on page 114

Verifying the SQLNET ALG Control Session

Purpose Verify that the SQL command is executed and all the SQL control and data sessions are created.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
Session ID: 10880, Policy name: sql, Timeout: 2, Valid
  In: 30.0.172.12/52315 --> 40.0.172.35/1521;tcp, If: fe-3/0/0.0, Pkts: 6, Bytes:
 492
  Out: 40.0.172.35/1521 --> 30.0.172.12/52315;tcp, If: fe-3/0/1.0, Pkts: 4, Bytes:
 227

Session ID: 10881, Policy name: sql, Timeout: 1800, Valid
Resource information : SQLV2 ALG, 5, 18
  In: 30.0.172.12/45944 --> 40.0.172.35/1114;tcp, If: fe-3/0/0.0, Pkts: 18, Bytes:
 4240
  Out: 40.0.172.35/1114 --> 30.0.172.12/45944;tcp, If: fe-3/0/1.0, Pkts: 15,
Bytes: 3989
Total sessions: 2
```

Meaning

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name, number of packets in and out.
- **Policy name**—Policy name that permitted the traffic.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and source interface for this session is fe-3/0/0.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and destination interface for this session is fe-3/0/1.0).

Copyright © 2016, Juniper Networks, Inc.

Verifying the SQLNET ALG

Purpose Verify that the SQLNET ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host>show security alg status
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Disabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the SQLNET ALG status as follows:

- Enabled—Shows the SQLNET ALG is enabled
- Disabled—Shows the SQLNET ALG is disabled.

Verifying the SQLNET ALG Resource Manager Group

Purpose Verify the total number of resource manager groups and active groups that are used by the SQLNET ALG.

Action From operational mode, enter the **show security resource-manager group active** command.

```
user@host>show security resource-manager group active
Group ID 1: Application - SQL ALG
    Total groups 677, active groups 1
```

Verifying the SQLNET ALG Resource Information

Purpose Verify the total number of resources and active resources that are used by the SQLNET ALG.

Action From operational mode, enter the **show security resource-manager resource active** command.

```
user@host>show security resource-manager resource active
Resource ID 2: Group ID - 1, Application - SQL ALG

      Resource ID 1: Group ID - 1, Application - SQL ALG
     Total Resources 4044, active resources 2
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding the SQLNET ALG on page 103

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 11

Configuring the TALK ALG

- Understanding the TALK ALG on page 117
- Example: Configuring the TALK ALG on page 117

Understanding the TALK ALG

The TALK ALG is a visual communication program used for interactive communication between two users. The TALK ALG processes TALK packets, performs Network Address Translation (NAT), and opens two gates (TCP and UDP) on the receiver side. One gate is used for the next LOOKUP packet. The other gate is used for make a connection from a client to a server and to initiate communication between a client and a server located on opposite sides of a Juniper Networks device.

There are two types of TALK servers: ntalkd and talkd.

The TALK ALG processes both ntalk and talkd packets. The TALK ALG uses port UDP517 and port UDP518 to establish a connection between a client and a server.

Related Documentation

- Example: Configuring the TALK ALG on page 117

Example: Configuring the TALK ALG

This example show how to configure the TALK ALG in route or NAT mode, allow the TALK traffic to pass through a device, and initiate communication between a client and a server located on opposite sides of a Juniper Networks device.

- Requirements on page 117
- Overview on page 118
- Configuration on page 119
- Verification on page 127

Requirements

This example uses the following hardware and software components:

- An SRX Series device

Copyright © 2016, Juniper Networks, Inc.

- Two PCs (client and server)

Before you begin:

- Understand the concepts behind ALGs. See "ALG Overview" on page 3.
- Understand the basics of TALK ALG. See "Understanding the TALK ALG" on page 117.

Overview

In this example, first you configure network interfaces on the device, create security zones and assign interfaces to the zones, and configure a policy to allow TALK traffic to go through an SRX Series device.

Then you create a static NAT rule set rs1 with a rule r1 to match the destination address 40.5.2.120/32, and you create a static NAT prefix with address 20.5.2.120/32.

Next you create a source NAT pool src-p1 with a source rule set src-rs1 to translate packets from zone trust to zone untrust. For matching packets, the source address is translated to an IP address in the src-p1 pool.

Then you create a destination NAT pool des-p1 with a destination rule set des-rs1 to translate packets from zone trust to destination address 40.5.2.121/32. For matching packets, the destination address is translated to an IP address in the des-p1 pool. Finally, you configure TALK ALG trace options.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 9 shows the TALK ALG topology.

Figure 9: TALK ALG Topology

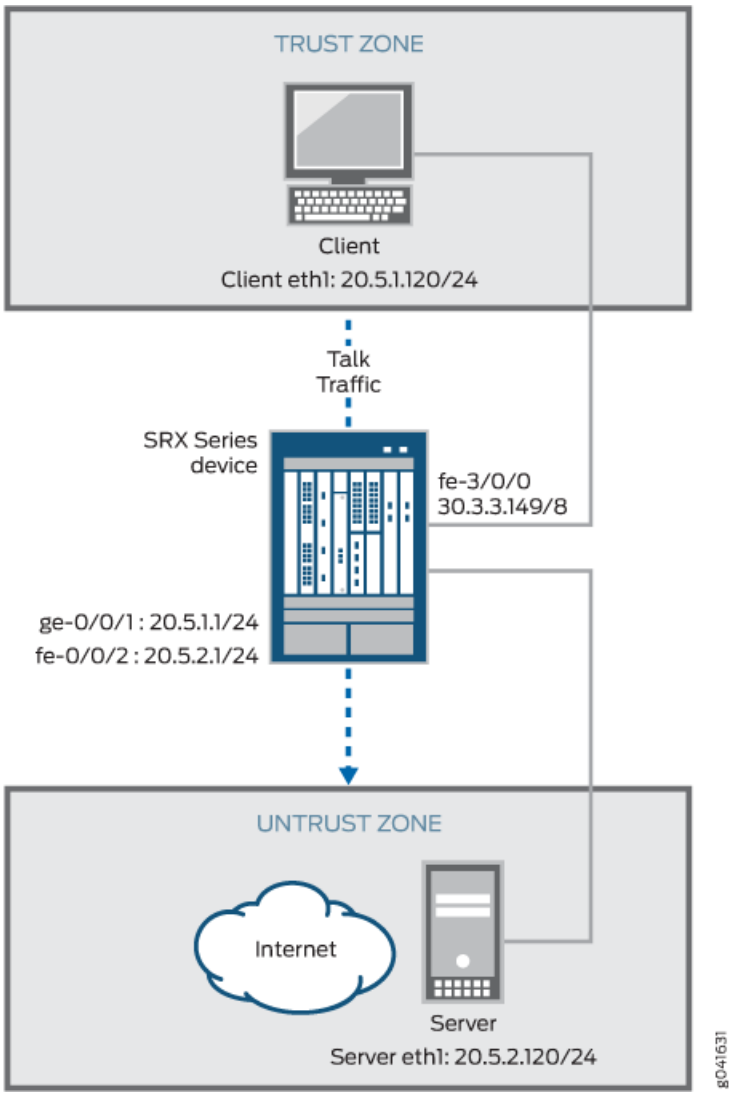

Configuration

To configure the TALK ALG, perform these tasks:

- Configuring a Route Mode on page 120
- Configuring a Static NAT Rule Set on page 122
- Configuring a Source NAT Pool and Rule Set on page 123
- Configuring a Destination NAT Pool and Rule Set on page 125
- Configuring TALK ALG trace options on page 126

Copyright © 2016, Juniper Networks, Inc.

Configuring a Route Mode

CLI Quick Configuration

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

```
set interfaces ge-0/0/1 unit 0 family inet address 20.5.1.1/24
set interfaces fe-0/0/2 unit 0 family inet address 20.5.2.1/24
set security zones security-zone trust interfaces ge-0/0/1 host-inbound-traffic system-services
 all
set security zones security-zone trust interfaces ge-0/0/1 host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/2 host-inbound-traffic system-services
 all
set security zones security-zone untrust interfaces fe-0/0/2 host-inbound-traffic protocols all
set security policies from-zone trust to-zone untrust policy talk match source-address any
set security policies from-zone trust to-zone untrust policy talk match destination-address any
set security policies from-zone trust to-zone untrust policy talk match application junos-ntalk
set security policies from-zone trust to-zone untrust policy talk then permit
```

Step-by-Step Procedure

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

To configure route mode:

1. Configure interfaces.

   ```
   [edit interfaces]
   user@host#set ge-0/0/1 unit 0 family inet address 20.5.1.1/24
   user@host#set fe-0/0/2 unit 0 family inet address 20.5.2.1/24
   ```

2. Configure zones and assign interfaces to the zones.

   ```
   [edit security zones security-zone trust]
   user@host#set interfaces ge-0/0/1 host-inbound-traffic system-services all
   user@host#set interfaces ge-0/0/1 host-inbound-traffic protocols all
   [edit security zones security-zone untrust]
   user@host#set interfaces fe-0/0/2 host-inbound-traffic system-services all
   user@host#set interfaces fe-0/0/2 host-inbound-traffic protocols all
   ```

3. Configure a TALK policy that allows TALK traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone untrust to-zone trust]
   user@host#set policy talk match source-address any
   user@host#set policy talk match destination-address any
   user@host#set policy talk match application junos-ntalk
   user@host#set policy talk then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show interfaces
...
   ge-0/0/1 {
   unit 0 {
      family inet {
         address 20.5.1.1/24;
      }
   }
}
...
fe-0/0/2 {
   unit 0 {
      family inet {
         address 20.5.2.1/24;
      }
   }
}
```

```
[edit]
user@host# show security zones
security-zone trust {
   ....
      interfaces {
      ge-0/0/1.0 {
         host-inbound-traffic {
            system-services {
               all;
            }
            protocols {
               all;
            }
         }
      }
   }
}
...
security-zone untrust {
   interfaces {
      fe-0/0/2.0 {
         host-inbound-traffic {
            system-services {
               all;
            }
            protocols {
               all;
            }
         }
      }
   }
}
```

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security policies
from-zone trust to-zone untrust {
  policy talk {
    match {
      source-address any;
      destination-address any;
      application junos-ntalk;
    }
    then {
      permit;
    }
  }
}
```

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

Configuring a Static NAT Rule Set

CLI Quick Configuration

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

```
set security nat static rule-set rs1 from zone trust
set security nat static rule-set rs1 rule r1 match destination-address 40.5.2.120/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 20.5.2.120/32
```

Step-by-Step Procedure

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

To configure a static NAT rule set:

1. Create a static NAT rule set.

   ```
   [edit security nat static rule-set rs1]
   user@host#set from zone trust
   ```

2. Define the rule to match with the destination address.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 match destination-address 40.5.2.120/32
   ```

3. Define the static NAT prefix for the device.

   ```
   [edit security nat static rule-set rs1]
   user@host# set rule r1 then static-nat prefix 20.5.2.120/32
   ```

Results

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

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
```


Copyright © 2016, Juniper Networks, Inc.

```
        from zone trust;
        rule r1 {
          match {
            destination-address 40.5.2.120/32
          }
          then {
            static-nat {
              prefix {
                20.5.2.120/32;
              }
            }
          }
        }
      }
    }
```

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

Configuring a Source NAT Pool and Rule Set

CLI Quick Configuration

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

```
set security nat source pool src-p1 address 40.5.1.120/32
set security nat source rule-set src-rs1 from zone trust
set security nat source rule-set src-rs1 to zone untrust
set security nat source rule-set src-rs1 rule src-r1 match source-address 20.5.1.120/32
set security nat source rule-set src-rs1 rule src-r1 match destination-address 20.5.2.120/32
set security nat source rule-set src-rs1 rule src-r1 then source-nat pool src-p1
```

Step-by-Step Procedure

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

To configure a source NAT pool and rule set:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host#set pool src-p1 address 40.5.1.120/32
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 from zone trust
   user@host# set rule-set src-rs1 to zone untrust
   ```

3. Configure a rule that matches packets and translates the source address to an address in the source pool.

   ```
   [edit security nat source]
   user@host# set rule-set src-rs1 rule src-r1 match source-address 20.5.1.120/32
   ```

4. Configure a rule that matches packets and translates the destination address to an address in the source pool.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat source]
user@host# set rule-set src-rs1 rule src-r1 match destination-address 20.5.2.120/32
```

5. Configure a source NAT pool in the rule.

```
[edit security nat source]
user@host# set rule-set src-rs1 rule src-r1 then source-nat pool src-p1
```

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

```
[edit]
user@host# show security nat
source {
  pool src-p1 {
    address {
      40.5.1.120/32;
    }
  }
  rule-set src-rs1 {
    from zone trust;
    to zone untrust;
    rule src-r1 {
      match {
        source-address 20.5.1.120/32;
        destination-address 20.5.2.120/32;
      }
      then {
        source-nat {
          pool {
            src-p1;
          }
        }
      }
    }
  }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Configuring a Destination NAT Pool and Rule Set

CLI Quick Configuration

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

```
set security nat destination pool des-p1 address 20.5.2.120/32
set security nat destination rule-set des-rs1 from zone trust
set security nat destination rule-set des-rs1 rule des-r1 match source-address 20.5.1.120/32
set security nat destination rule-set des-rs1 rule des-r1 match destination-address 40.5.2.120/32
set security nat destination rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
```

Step-by-Step Procedure

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

To configure a destination NAT pool and rule set:

1. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host#set pool des-p1 address 20.5.2.120/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 from zone trust
   ```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match source-address 20.5.1.120/32
   ```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 match destination-address 40.5.2.120/32
   ```

5. Configure a source NAT pool in the rule.

   ```
   [edit security nat destination]
   user@host# set rule-set des-rs1 rule des-r1 then destination-nat pool des-p1
   ```

Results

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

```
[edit]
user@host# show security nat
destination {
   pool des-p1 {
      address {
         20.5.2.120/32;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
rule-set des-rs1 {
   from zone trust;
   rule des-r1 {
      match {
         source-address 20.5.1.120/32;
         destination-address 40.5.2.120/32;
      }
      then {
         destination-nat {
            pool {
               des-p1;
            }
         }
      }
   }
}
```

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

Configuring TALK ALG trace options

CLI Quick Configuration

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

```
set security alg talk traceoptions flag all
set security alg traceoptions file trace
set security alg traceoptions file size 1g
set security alg traceoptions level verbose
```

Step-by-Step Procedure

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

To configure TALK ALG trace options:

1. Enable TALK ALG trace otpions.

   ```
   [edit security alg]
   user@host#set talk traceoptions flag all
   ```

2. Configure the filename to receive output from the tracing operation.

   ```
   [edit security alg]
   user@host#set traceoptions file trace
   ```

3. Specify the maximum trace file size.

   ```
   [edit security alg]
   user@host#set traceoptions file size 1g
   ```

4. Specify the level of tracing output.

   ```
   [edit security alg]
   user@host#set traceoptions level verbose
   ```


Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show security alg
traceoptions {
   file trace size 1g;
   level verbose;
   }
talk traceoptions flag all;
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the TALK ALG Control Session on page 127
- Verifying the TALK Flow Gate Information on page 128
- Verifying TALK ALG on page 129
- Verifying the TALK Resource Manager Group on page 129
- Verifying the TALK Resource Information on page 129

Verifying the TALK ALG Control Session

Purpose Verify that the TALK control session is created and all the TALK control and data sessions are created.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session

 ession ID: 128570, Policy name: p11/4, Timeout: 56, Valid
Resource information : TALK ALG, 2, 0
  In: 5.1.1.200/1105 --> 6.1.1.200/518;udp, If: ge-0/0/1.0, Pkts: 3, Bytes: 336
  Out: 6.1.1.200/518 --> 5.1.1.200/1105;udp, If: ge-0/0/2.0, Pkts: 3, Bytes: 156

Session ID: 128617, Policy name: p11/4, Timeout: 1796, Valid
Resource information : TALK ALG, 2, 2
  In: 6.1.1.200/42224 --> 5.1.1.200/518;udp, If: ge-0/0/2.0, Pkts: 1, Bytes: 112

  Out: 5.1.1.200/518 --> 6.1.1.200/42224;udp, If: ge-0/0/1.0, Pkts: 1, Bytes: 52

Session ID: 128618, Policy name: p11/4, Timeout: 1796, Valid
Resource information : TALK ALG, 2, 3
  In: 6.1.1.200/51430 --> 5.1.1.200/32905;tcp, If: ge-0/0/2.0, Pkts: 4, Bytes:
219
  Out: 5.1.1.200/32905 --> 6.1.1.200/51430;tcp, If: ge-0/0/1.0, Pkts: 3, Bytes:
167
```

Copyright © 2016, Juniper Networks, Inc.

Meaning

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **Policy name**—Policy name that permitted the traffic.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and the source interface for this session is ge-0/0/1.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, session is TCP, and destination interface for this session is fe-0/0/2.0).

Verifying the TALK Flow Gate Information

Purpose Verify that the gates are opened for TCP data channel and reverse UDP reply.

Action From operational mode, enter the **show security flow gate** command.

```
user@host>show security flow gate
Hole: 6.1.1.200-6.1.1.200/0-0->5.1.1.200-5.1.1.200/518-518
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: udp
  Application: TALK ALG/65
  Age: 110 seconds
  Flags: 0x0080
  Zone: untrust
  Reference count: 1
  Resource: 11-2-2

Hole: 6.1.1.200-6.1.1.200/0-0->5.1.1.200-5.1.1.200/32905-32905
  Translated: 0.0.0.0/0->0.0.0.0/0
  Protocol: tcp
  Application: TALK ALG/65
  Age: 110 seconds
  Flags: 0x0080
  Zone: untrust
  Reference count: 1
  Resource: 11-2-3
```

Meaning

- **Hole**—Range of flows permitted by the pinhole.
- **Translated**—Tuples used to create the session if it matches the pinhole (source and destination IP addresses with their respective source and destination port numbers).
- **Protocol**—Application protocol, such as UDP or TCP.
- **Application**—Name of the application.
- **Age**—Idle timeout for the pinhole.
- **Flags**— Internal debug flags for the pinhole.
- **Zone**—Security zone such as from zone and to zone.
- **Reference count**—Number of resource manager references to the pinhole.
- **Resource**—Resource manager information about the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Verifying TALK ALG

Purpose Verify that the TALK ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host>show security alg status
ALG Status :
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the TALK ALG status as follows:

- Enabled—Shows the TALK ALG is enabled.
- Disabled—Shows the TALK ALG is disabled.

Verifying the TALK Resource Manager Group

Purpose Verify the total number of resource manager groups and active groups that are used by the TALK ALG.

Action From operational mode, enter the **show security resource-manager group active** command.

```
user@host>show security resource-manager group active
Group ID 2: Application - TALK ALG
       Total groups 3276, active groups 1
```

Verifying the TALK Resource Information

Purpose Verify the total number of resources and active resources that are used by the TALK ALG.

Action From operational mode, enter the **show security resource-manager resource active** command.

```
user@host>show security resource-manager resource active
Resource ID 3: Group ID - 2, Application - TALK ALG

Resource ID 2: Group ID - 2, Application - TALK ALG
Total Resources 6015, active resources 2
```

Related Documentation

- Understanding the TALK ALG on page 117

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 12

Configuring the TFTP ALG

- Understanding the TFTP ALG on page 131
- Understanding TFTP ALG Conversation on page 132
- Understanding IPv6 Support for the TFTP ALG on page 133
- Example: Configuring the TFTP ALG on page 134

Understanding the TFTP ALG

- Overview on page 131
- TFTP Packets on page 131
- TFTP Session on page 132

Overview

Trivial File Transfer Protocol (TFTP) is a simple protocol used for files transfer *(RFC 1350)*. TFTP is implemented on top of UDP, with destination port 69 as the well-known port. The TFTP Application Layer Gateway (ALG) processes TFTP packets that initiate the request and creates pinholes to allow return packets from the reverse direction.

In flow processing there are two sessions for one TFTP conversation, one is the TFTP control session created by a read request (RRQ) or write request (WRQ) packet; the other one is the TFTP data session created by a DATA packet (for RRQ) or acknowledgment (ACK) packet (for WRQ).

In a Junos OS firewall, the TFTP control session is permitted through the junos-tftp application policy. The data session is permitted through the TFTP ALG open pinhole from any port of the server to the TID (port) of the client when the control session packet is received. No NAT translation is required, because the NAT translation has already been performed and the information is available from the session data structure.

On SRX320 and SRX340 devices, broadcast TFTP is not supported when flow is enabled on the device.

TFTP Packets

Any transfer begins with a request to read or write a file. A data packet of less than 512 bytes signals termination of a transfer.

Copyright © 2016, Juniper Networks, Inc.

TFTP supports five types of packets:

- Read request (RRQ)
- Write request (WRQ)
- Data (DATA)
- Acknowledgment (ACK)
- Error (ERROR)

TFTP Session

The TFTP ALG is based on UDP, which is a stateless transport protocol. In a firewall, the TFTP ALG acts as a UDP session with timeout. If there is no packet refresh session, the session is terminated after timeout. Although the TFTP client and server determine the termination of a TFTP conversation, they are sometimes unaware of the session in Fireware. Therefore, the client and server could request a new TFTP conversation in this scenario.

The TFTP ALG session can proceed in any of the following ways:

- When the TFTP control session reaches timeout, the session is not terminated if the data session is still alive.
- A TFTP session might terminate or get corrupted by the **clear security flow session all** or the **clear specific session** CLI commands regardless of whether the data session is ongoing or not.
- If a new TFTP session request arrives and reaches the existing session, the TFTP ALG will open the pinhole again for the new request.
- If the pinhole already exists, the TFTP ALG will not open the pinhole again and there will be no packet drop.
- The TFTP ALG will not drop any packet.

Related Documentation

- ALG Overview on page 3
- Understanding TFTP ALG Conversation on page 132
- Understanding IPv6 Support for the TFTP ALG on page 133
- Example: Configuring the TFTP ALG on page 134

Understanding TFTP ALG Conversation

By default TFTP servers listen for incoming requests from TFTP clients on port 69. A TFTP client chooses its source tunnel identifier (TID) port and sends its initial request to the server. In response, the server uses the TID chosen as the source port and sends a response to the client's TID as the destination port. The two TIDs ports are then used for the rest of the data transfer.

Copyright © 2016, Juniper Networks, Inc.

Read file conversation steps:

1. Host A (client) sends an RRQ packet to host B (server) with A's TID as source and port 69 as destination.
2. Host B (server) sends a DATA packet to host A (client) with B's TID as source and A's TID as destination.
3. Host A (client) sends an ACK packet to host B (server) with A's TID as source and B's TID as destination.
4. DATA and ACK packets conversation continues until file data transferring is complete.

Write file conversation steps:

1. Host A (client) sends a WRQ packet to host B (server) with A's TID as source and port 69 as destination.
2. Host B (server) sends an ACK packet to host A (client) with B's TID as source and A's TID as destination.
3. Host A (client) sends a DATA packet to host B (server) with A's TID as source and B's TID as destination.
4. Host B (server) sends an ACK packet to host A (client) with B's TID as source and A's TID as destination.

Related Documentation

- ALG Overview on page 3
- Understanding IPv6 Support for the TFTP ALG on page 133
- Example: Configuring the TFTP ALG on page 134

Understanding IPv6 Support for the TFTP ALG

Trivial File Transfer Protocol (TFTP) Application Layer Gateway (ALG) has been enhanced to support IPv6 and IPv4 TFTP conversation, which has IPv6 and IPv4 addresses for both the source IP address and destination IP address.

TFTP ALG processes packets that initiate the routing request and create pinholes to allow return packets from the reverse direction to the port that sent the request.

The data session is set up by the first packet from the client to the server. TFTP ALG monitors the first packet and opens a pinhole from any port on the server to the client. This process helps the return packets from the server and subsequent data packets to pass through.

Related Documentation

- ALG Overview on page 3
- Understanding TFTP ALG Conversation on page 132
- Example: Configuring the TFTP ALG on page 134

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the TFTP ALG

The TFTP ALG processes TFTP packets that initiate the request and opens a gate to allow return packets from the reverse direction to the port that sends the request.

This example shows how to configure the TFTP ALG to pass through TFTP traffic with a source NAT pool on Juniper Networks devices.

- Requirements on page 134
- Overview on page 134
- Configuration on page 134
- Verification on page 136

Requirements

- Configure proxy ARP for all IP addresses in the source NAT pool.
- Understand the basic concepts of TFTP ALG. See "Understanding the TFTP ALG" on page 131.

Overview

In this example, the TFTP ALG is configured to monitor and allow TFTP traffic, transferring files between the client and server located on opposite sides of a Juniper Networks device.

Configuration

Configuring a NAT Source Pool, Rule Set, and a Policy

CLI Quick Configuration

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

```
set security nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
set security zones security-zone green address-book address sa1 1.1.1.0/24
set security zones security-zone red address-book address da1 2.2.2.0/24
set security nat source rule-set rs1 from zone green
set security nat source rule-set rs1 to zone red
set security nat source rule-set rs1 rule r1 match source-address 1.1.1.0/24
set security nat source rule-set rs1 rule r1 match destination-address 2.2.2.0/24
set security nat source rule-set rs1 rule r1 then source-nat pool pool1

set security policy from-zone green to-zone red policy pol1 match destination-address da1
set security policy from-zone green to-zone red policy pol1 match source-address sa1
set security policy from-zone green to-zone red policy pol1 match application junos-tftp
set security policy from-zone green to-zone red policy pol1 then permit
```

NOTE: If you are not sure of the TFTP client and server IP address, you can replace "da1" and "sa1" with "any".

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure a source NAT pool:

1. Create a NAT source pool.

```
[edit security]
user@host# set nat source pool pool1 address 10.10.10.1/32 to 10.10.10.10/32
```

2. Configure security zone address book entries.

```
[edit security zones security-zone]
user@host# set green address-book address sa1 1.1.1.0/24
user@host# set red address-book address da1 2.2.2.0/24
```

3. Create a NAT source rule set.

```
[edit security nat source rule-set rs1]
user@host# set from zone green
user@host# set to zone red
user@host# set rule r1 match source-address 1.1.1.0/24
user@host# set rule r1 match destination-address 2.2.2.0/24
user@host# set rule r1 then source-nat pool pool1
```

4. Configure a policy

```
[edit security policies from-zone green to-zone red policy pol1]
user@host# set match source-address sa1
user@host# set match destination-address da1
user@host# set match application junos-tftp
user@host# set then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool pool1 {
    address {
      10.10.10.1/32 to 10.10.10.10/32;
    }
  }
  rule-set rs1 {
    from zone green;
    to zone red;
    rule r1 {
      match {
        source-address 1.1.1.0/24;
        destination-address 2.2.2.0/24;
      }
      then {
        source-nat {
          pool {
```

Copyright © 2016, Juniper Networks, Inc.

```
            pool1;
          }
        }
      }
    }
  }
[edit]
user@host# show security policies
from-zone green to-zone red {policy pol1 {
  policy pol1 {
    match {
      source-address sa1;
      destination-address da1;
      application [junos-tftp];
    }
    then {
      permit;
    }
  }
}
default-policy {
  permit-all;
}
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the NAT Source Pool and Rule Set on page 136
- Verifying TFTP ALG on page 136

Verifying the NAT Source Pool and Rule Set

Purpose Verify that the NAT source pool and rule set used to support the TFTP ALG are working properly.

Action From operational mode, enter the **show security nat static rule r1** command.

Verifying TFTP ALG

Purpose Verify that the TFTP ALG is enabled.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status

DNS      : Enabled
FTP      : Enabled
H323     : Enabled
TFTP     : Enabled
```

Meaning The output shows the TFTP ALG status as follows:

Copyright © 2016, Juniper Networks, Inc.

- Enabled—Shows the TFTP ALG is enabled.
- Disabled—Shows the TFTP ALG is disabled.

Related Documentation

- ALG Overview on page 3
- Understanding TFTP ALG Conversation on page 132
- Understanding IPv6 Support for the TFTP ALG on page 133

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 3

Configuring VoIP ALGs

- VoIP ALG Types on page 141
- Configuring VoIP Rewrite Rules on page 143
- Configuring the H.323 ALG on page 145
- Configuring the MGCP ALG on page 183
- Configuring the SCCP ALG on page 217
- Configuring the SIP ALG on page 241

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 13

VoIP ALG Types

- Understanding VoIP ALG Types on page 141

Understanding VoIP ALG Types

Junos OS supports voice-over-IP Application Layer Gateways (VoIP ALGs) and basic data ALGs. (Note that supported ALG types vary depending on which hardware device you are using.)

VoIP ALGs provide stateful Application Layer inspection and Network Address Translation (NAT) capabilities to VoIP signaling and media traffic. The ALG inspects the state of transactions, or calls, and forwards or drops packets based on those states.

Junos OS supports the following VoIP ALGs:

- H.323—The H.323 ALG provides support for the H.323 legacy VoIP protocol. The ALG lets you secure VoIP communication between terminal hosts, such as IP phones and multimedia devices. In such a telephony system, the gatekeeper device manages call registration, admission, and call status for VoIP calls. Gatekeepers can reside in the two different zones or in the same zone.
- SIP—The SIP ALG provides support for the Session Initiation Protocol (SIP). SIP is an Internet Engineering Task Force (IETF)-standard protocol for initiating, modifying, and terminating multimedia sessions over the Internet. Such sessions might include conferencing, telephony, or multimedia, with features such as instant messaging and application-level mobility in network environments.
- SCCP—The SCCP ALG provides support for Skinny Client Control Protocol (SCCP). SCCP is a Cisco proprietary protocol for call signaling. Skinny is based on a call-agent-based call-control architecture. The control protocol uses binary-coded frames encoded on TCP frames sent to well-known TCP port number destinations to set up and tear down RTP media sessions.
- MGCP—The MGCP ALG provides support for Media Gateway Control Protocol (MGCP). MGCP is a text-based Application Layer protocol used for call setup and call control between the media gateway and the media gateway controller (MGC). .

For information about enabling and configuring each of these ALGs through J-Web, select the **Configure>Security>ALG** page in the J-Web user interface and click **Help**.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- ALG Overview on page 3
- Understanding the ALG for IKE and ESP on page 33
- Understanding H.323 ALG on page 145
- Understanding the SIP ALG on page 241
- Understanding SCCP ALGs on page 217
- Understanding the MGCP ALG on page 183
- Understanding RPC ALGs on page 61

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 14

Configuring VoIP Rewrite Rules

- Understanding VoIP DSCP Rewrite Rules on page 143
- Example: Configuring VoIP DSCP Rewrite Rules on page 144

Understanding VoIP DSCP Rewrite Rules

This topic describes the voice over IP Application Layer Gateway (VoIP ALG) mechanism for modifying the Differentiated Services Code Point (DSCP) field of Real-Time Transport Protocol (RTP) packets. The VoIP ALG mechanism is applicable for the RTP session, which is recognized by the ALG.

DSCP is a modification of the type of service byte for class of service (CoS). Six bits of this byte are reallocated for use as the DSCP field, where each DSCP specifies a particular per-hop behavior that is applied to a packet.

To avoid VoIP quality degradation caused by network congestion, the RTP packets are required to mark the DSCP bit to ensure they get higher routing priority. A downstream router can put those packets in a higher priority queue for faster forwarding. To provide this functionality, there needs to be a per-VoIP mechanism for modifying the DSCP field of RTP packets according to the specific configuration. This will ensure that all RTP packets based on User Datagram Protocol/Transport Control Protocol (UDP/TCP) that encounter the ALG will be assigned a specific DSCP bit.

A rewrite rule modifies the appropriate CoS bits in an outgoing packet to meet the requirements of the targeted peer. Each rewrite rule reads the current CoS value that is configured at the VoIP ALG level. Every packet that hits the VoIP ALG is marked by this CoS value.

This feature supports ALG DSCP marking for H323, Session Initiation Protocol (SIP), Media Gateway Control Protocol (MGCP), and Skinny Client Control Protocol (SCCP). It provides a 6-bit DSCP value configuration for each of these. When the first RTP packet hits the ALG, this feature receives the 6-bit DSCP value form the configuration and sets it to the RTP session that the packet has created. This first RTP packet and the following RTP packets passing through the RTP session are marked according to the 6-bit DSCP value in the session.

Related Documentation

- Example: Configuring VoIP DSCP Rewrite Rules on page 144

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring VoIP DSCP Rewrite Rules

This example shows how to configure VoIP DSCP.

- Requirements on page 144
- Overview on page 144
- Configuration on page 144
- Verification on page 144

Requirements

This example uses an SRX320 device. The example assumes that the ALG has been enabled.

Overview

This example shows how to configure four ALG DSCP markings; SIP, H323, MGCP, and SCCP. You set the 6-bit DSCP value configuration for each ALG DSCP.

Configuration

Step-by-Step Procedure

To configure VoIP DSCP rewrite rules:

1. Set the DSCP for each VoIP ALG.

```
[edit]
user@host# set security alg sip dscp-rewrite code-point 101010
user@host# set security alg h323 dscp-rewrite code-point 010101
user@host# set security alg mgcp dscp-rewrite code-point 111000
user@host# set security alg sccp dscp-rewrite code-point 000111
```

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

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

Verification

To verify that the configuration is working properly, enter the **show security alg** command.

Related Documentation

- Understanding VoIP DSCP Rewrite Rules on page 143

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 15

Configuring the H.323 ALG

- Understanding H.323 ALG on page 145
- Understanding the Avaya H.323 ALG on page 147
- H.323 ALG Configuration Overview on page 149
- Example: Passing H.323 ALG Traffic to a Gatekeeper in the Private Zone on page 149
- Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 154
- Example: Using NAT with the H.323 ALG to Enable Incoming Calls on page 160
- Example: Using NAT with the H.323 ALG to Enable Outgoing Calls on page 168
- Understanding H.323 ALG Endpoint Registration Timeouts on page 174
- Example: Setting H.323 ALG Endpoint Registration Timeouts on page 175
- Understanding H.323 ALG Media Source Port Ranges on page 176
- Example: Setting H.323 ALG Media Source Port Ranges on page 176
- Understanding H.323 ALG DoS Attack Protection on page 177
- Example: Configuring H.323 ALG DoS Attack Protection on page 178
- Understanding H.323 ALG Unknown Message Types on page 179
- Example: Allowing Unknown H.323 ALG Message Types on page 179

Understanding H.323 ALG

The H.323 standard is a legacy voice-over-IP (VoIP) protocol defined by the International Telecommunication Union (ITU-T). H.323 consists of a suite of protocols (such as H.225.0 and H.245) that are used for call signaling and call control for VoIP.

H.323 uses the ASN.1 coding format. It sets up the dynamic links for data, video, and audio streams, following the protocols Q.931 (with port number 1720) and H.245. There are three major processes in H.323:

- Gatekeeper Discovery—An endpoint finds its gatekeeper through the gatekeeper discovery process, through broadcast or unicast (to a known IP and the well-known UDP port 1719). (Junos OS supports unicast only.)
- Endpoint Registration, Admission, and Status—An endpoint registers to a gatekeeper and asks for its management. Before making a call, an endpoint asks its gatekeeper for permission to place the call. In both registration and admission phases, the

Copyright © 2016, Juniper Networks, Inc.

Registration, Admission, and Status (RAS) channel is used. The Transport Service Access Point (TSAP) can be either the well-known UDP port (1719) or a dynamically assigned port from the discovery or registration phase.

- Call Control and Call Setup—Calls can be established within a zone or across two zones, or even across multiple zones (multipoint conference). The call setup and tear down is performed through the call signaling channel whose TSAP is the well-known TCP port (1720). The call control, including opening/closing media channels between two endpoints, is performed through the call control channel whose TSAP is dynamically assigned from the previous call signaling process. H.245 messages are used in the call control channel, and are encoded using ASN.1.

NOTE: Detailed information on H.323 can be found in ITU-T Recommendation H.323.

The H.323 Application Layer Gateway (ALG) lets you secure VoIP communication between terminal hosts, such as IP phones and multimedia devices. In such a telephony system, the gatekeeper device manages call registration, admission, and call status for VoIP calls. Gatekeepers can reside in the two different zones or in the same zone. (See Figure 10.)

Figure 10: H.323 ALG for VoIP Calls

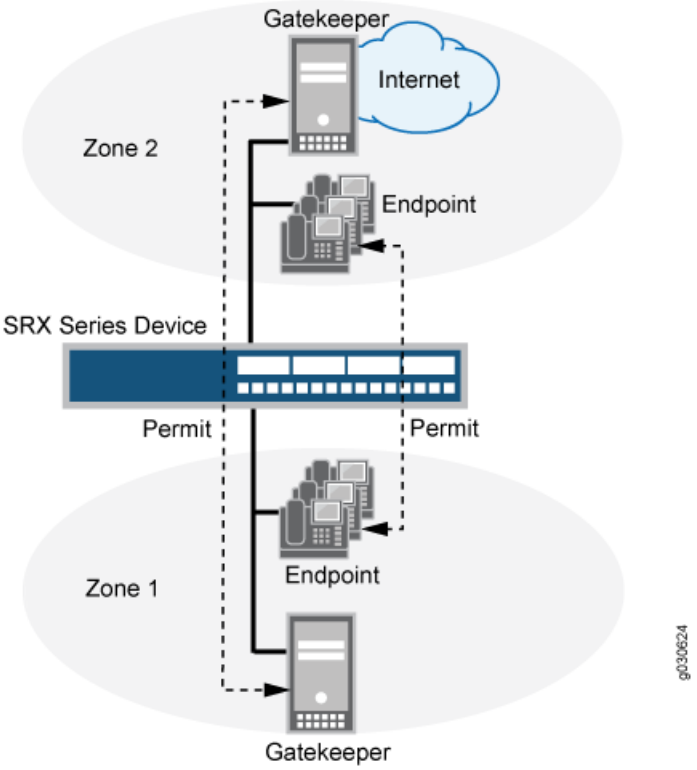

Copyright © 2016, Juniper Networks, Inc.

NOTE: The illustration uses IP phones for illustrative purposes, although it is possible to make configurations for other hosts that use VoIP, such as Microsoft NetMeeting multimedia devices.

Related Documentation

- ALG Overview on page 3
- Understanding the Avaya H.323 ALG on page 147
- H.323 ALG Configuration Overview on page 149

Understanding the Avaya H.323 ALG

The H.323 standard is a legacy voice-over-IP (VoIP) protocol defined by the International Telecommunication Union (ITU-T). H.323 consists of a suite of protocols (such as H.225.0 and H.245) that are used for call signaling and call control for VoIP. The processes for configuring the H.323 standard Application Layer Gateway (ALG) and the proprietary Avaya H.323 ALG are the same.

However, Avaya H.323 ALG has some special features. To understand and configure the Avaya H.323-specific features listed here, see the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

This topic contains the following sections:

- Avaya H.323 ALG-Specific Features on page 147
- Call Flow Details in the Avaya H.323 ALG on page 148

Avaya H.323 ALG-Specific Features

Avaya H.323-specific features are as follows:

- H.323 Fast Connect
- H.323 asymmetric media
- Call waiting
- Call forwarding
- Voice mail
- Call identification
- Conference calling

Copyright © 2016, Juniper Networks, Inc.

Call Flow Details in the Avaya H.323 ALG

- Connecting the Phone into the Network—Avaya performs the Q.931 Setup/Connect negotiation when the phone is wired into the network rather than when a call is being initiated.
- Making a call—When a call is made, because the PBX has already stored the capabilities for each phone when the phone is connected to the network, no further Q.931 and PBX negotiations are required to set up the call. It no longer exchanges Q.931 Setup and Connect messages with the PBX. The phone and the PBX exchange H.323 Facility messages to set up the call.
- Registering with a CM—When a call has been made, Avaya H.323 registers with the Avaya Communication Manager (CM). The registration process is similar to a generic H.323 standard registration process.

  NOTE: The direct mode and tunnel mode are not defined by Avaya H.323 ALG.

  For a call to work, the CM must be deployed with Avaya Endpoints. During the call, RAS and Q.931 messages are exchanged between the CM and the Avaya Endpoints.

  NOTE: For Avaya H.323 with a source Network Address Translation (NAT) pool, the registration process allows only one IP address in the pool.

- Setting up Real-Time Transport Protocol (RTP)/Real-Time Control Protocol (RTCP) ports—The Q.931 Setup, Facility and Information messages are used to set up RTP/RTCP ports. The hierarchy for an Avaya H.323 session is Q.931, RTP/RTCP, Parent, and then Child.

  NOTE: H.245 ports are not used in an Avaya call flow process.

- Using Avaya H.323 counters—The counters for calls and active calls are not applicable to the Avaya H.323 ALG. The call creation and tearing down is done by Facility messages afterward. When resources are allocated for a call, all counters for calls and active calls increment. If resources are allocated for a call multiple times, messages belonging to the same call that pass the firewall multiple times will trigger multiple increments of the counters. In other words, messages that belong to the same call and pass the firewall multiple times might trigger multiple increments of the counters if the resource for a call needs to be allocated multiple times.

  For example, in the two-zone case, the setup and connect message pair allocates one call resource. The active call counter is increased once. Each time the setup and connect message pair passes the firewall, a different call resource with unique interfaces and NAT is allocated. Therefore, the counter increments twice in a three-zone scenario.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- ALG Overview on page 3
- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149

H.323 ALG Configuration Overview

The H.323 Application Layer Gateway (ALG) is enabled by default on the device—no action is required to enable it. However, you might choose to fine-tune H.323 ALG operations by using the following instructions:

1. Specify how long an endpoint registration entry remains in the Network Address Translation (NAT) table. For instructions, see "Example: Setting H.323 ALG Endpoint Registration Timeouts" on page 175.
2. Enable media traffic on a narrow or wide range of ports. For instructions, see "Example: Setting H.323 ALG Media Source Port Ranges" on page 176.
3. Protect the H.323 gatekeeper from denial-of-service (DoS) flood attacks. For instructions, see "Example: Configuring H.323 ALG DoS Attack Protection" on page 178.
4. Enable unknown messages to pass when the session is in NAT mode and route mode. For instructions, see "Example: Allowing Unknown H.323 ALG Message Types" on page 179.

Related Documentation
- Understanding H.323 ALG on page 145
- Example: Passing H.323 ALG Traffic to a Gatekeeper in the Private Zone on page 149
- Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 154
- Example: Using NAT with the H.323 ALG to Enable Incoming Calls on page 160
- Example: Using NAT with the H.323 ALG to Enable Outgoing Calls on page 168

Example: Passing H.323 ALG Traffic to a Gatekeeper in the Private Zone

This example shows how to set up two policies that allow H.323 traffic to pass between IP phone hosts and a gatekeeper in the private zone, and an IP phone host (2.2.2.5/32) in the public zone.

- Requirements on page 149
- Overview on page 150
- Configuration on page 150
- Verification on page 153

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.
- Configure security zones. See "Understanding Security Zones" on page 1050.

Overview

This example shows how to set up two policies that allow H.323 traffic to pass between IP phone hosts and a gatekeeper in the private zone, and an IP phone host (2.2.2.5/32) in the public zone. The connection to the device can either be with or without NAT. See Figure 11.

Figure 11: H.323 Gatekeeper in the Private Zone

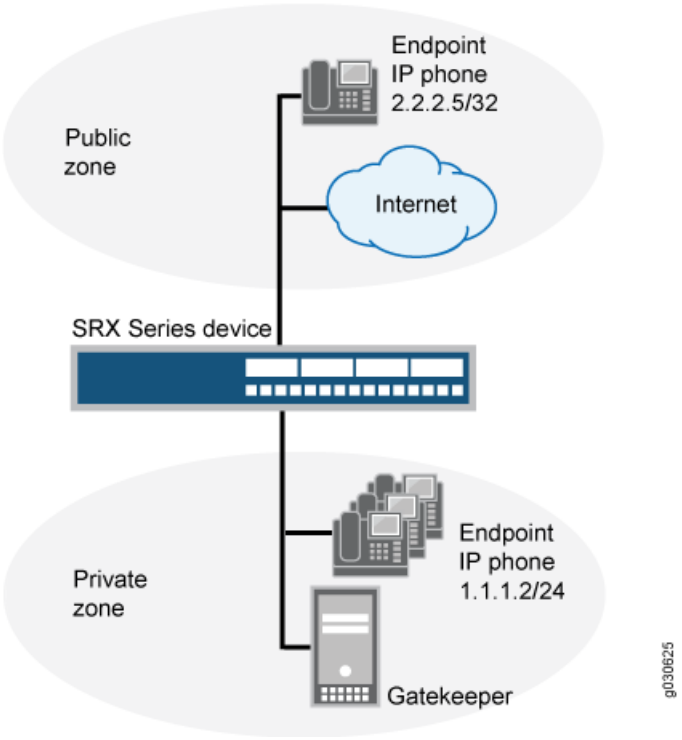

Configuration

CLI Quick Configuration

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

```
set security zones security-zone public address-book address ip_phone 2.2.2.5/32
set security zones security-zone private address-book address gateway 2.2.2.5/32
set security policies from-zone private to-zone public policy P1 match source-address
    any
set security policies from-zone private to-zone public policy P1 match destination-address
    IP_Phone
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone private to-zone public policy P1 match application
  junos-h323
set security policies from-zone private to-zone public policy P1 then permit
set security policies from-zone public to-zone private policy P2 match source-address
  any
set security policies from-zone public to-zone private policy P2 match destination-address
  gateway
set security policies from-zone public to-zone private policy P2 match application
  junos-h323
set security policies from-zone public to-zone private policy P2 then permit
```

Step-by-Step Procedure

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

To configure the device to pass H.323 ALG traffic to a gatekeeper in the private zone:

1. Configure two address books.

   ```
   [edit]

   user@host# set security zones security-zone public address-book address ip_phone
     2.2.2.5/32

   set security zones security-zone private address-book address gateway 2.2.2.5/32
   ```

2. Configure policy P1 from the private zone to the public zone.

   ```
   [edit]

   user@host# set security policies from-zone private to-zone public policy P1 match
     source-address any

   user@host# set security policies from-zone private to-zone public policy P1 match
     destination-address IP_Phone

   user@host# set security policies from-zone private to-zone public policy P1 match
     application junos-h323

   user@host# set security policies from-zone private to-zone public policy P1 then
     permit
   ```

3. Configure policy P2 from the public zone to the private zone.

   ```
   [edit]

   user@host# set security policies from-zone public to-zone private policy P2 match
     source-address any

   user@host# set security policies from-zone public to-zone private policy P2 match
     destination-address gateway

   user@host# set security policies from-zone public to-zone private policy P2 match
     application junos-h323
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security policies from-zone public to-zone private policy P2 then
  permit
```

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

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

```
[edit]
user@host# show security policies
...
from-zone trust to-zone trust {
    policy default-permit {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone trust to-zone untrust {
    policy default-permit {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone untrust to-zone trust {
    policy default-deny {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            deny;
        }
    }
}
from-zone private to-zone public {
    policy P1 {
        match {
            source-address any;
            destination-address IP_Phone;
            application junos-h323;
        }
        then {
            permit;
        }
    }
```

 Copyright © 2016, Juniper Networks, Inc.

```
}
from-zone public to-zone private {
    policy P2 {
        match {
            source-address any;
            destination-address gateway;
            application junos-h323;
        }
        then {
            permit;
        }
    }
}
...
```

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

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying H.323 ALG Configurations on page 153

Verifying H.323 ALG Configurations

Purpose Display information about active calls.

NOTE: H.323 counters for calls and active calls in the output to this **show security** command do not apply to the proprietary Avaya implementation of H.323. This is because Q.931 setup and connect messages are exchanged right after the phone is powered up and call creation and tear down is done by Facility messages.

Counters for calls and active calls are increased when the resources allocated for calls are increased—that is, messages belonging to the same call and that pass the firewall multiple times increment the counters. This applies when resources for a call need to be allocated multiple times. For example, in a two-zone scenario the setup and connect message pair allocates one call resource, and the active call counter is increased by one. But in a three-zone scenario the setup and connect message pair passes the firewall twice, each time allocating different call resources. In this case, the counter is incremented.

Action From the J-Web interface, select **Monitor>ALGs>H323**. Alternatively, from the CLI, enter the **show security alg h323 counters** command.

Counters for H.245 messages received also will not be accurate in the case of H.245 tunneling. Because H.245 messages are encapsulated in Q.931 packets, the counter for

Copyright © 2016, Juniper Networks, Inc.

H.245 messages received will remain zero even when there are H.245 messages. The **Other H245** counter will, however, reflect these packet transmissions.

```
[edit]
user@host> show security alg h323 counters
H.323 counters summary:
  Packets received         : 0
  Packets dropped          : 0
  RAS message received     : 0
  Q.931 message received   : 0
  H.245 message received   : 0
  Number of calls          : 0
  Number of active calls   : 0
H.323 error counters:
  Decoding errors          : 0
  Message flood dropped    : 0
  NAT errors               : 0
  Resource manager errors  : 0
H.323 message counters:
  RRQ       : 0
  RCF       : 0
  ARQ       : 0
  ACF       : 0
  URQ       : 0
  UCF       : 0
  DRQ       : 0
  DCF       : 0
  Oth RAS   : 0
  Setup     : 0
  Alert     : 0
  Connect   : 0
  CallProd  : 0
  Info      : 0
  RelCmpl   : 0
  Facility  : 0
  Empty     : 0
  OLC       : 0
  OLC-ACK   : 0
  Oth H245  : 0
```

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149

Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone

This example shows how to set up two policies to allow H.323 traffic to pass between IP phone hosts in the internal zone, and the IP phone at IP address 2.2.2.5/32 (and the gatekeeper) in the external zone.

- Requirements on page 155
- Overview on page 155
- Configuration on page 155
- Verification on page 158

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.
- Configure security zones. See "Understanding Security Zones" on page 1050.

Overview

Because route mode does not require address mapping of any kind, a device configuration for a gatekeeper in the external, or public, zone is usually identical to the configuration for a gatekeeper in an internal, or private, zone. This example shows how to set up two policies to allow H.323 traffic to pass between IP phone hosts in the internal zone, and the IP phone at IP address 2.2.2.5/32 (and the gatekeeper) in the external zone. The device can be in transparent or route mode. See Figure 12.

Figure 12: H.323 Gatekeeper in the External Zone

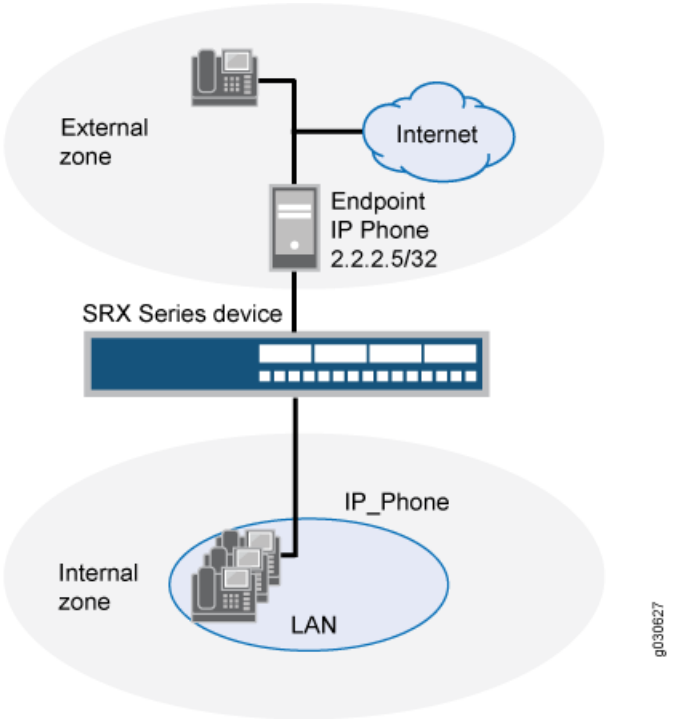

Configuration

CLI Quick Configuration

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

Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone external address-book address IP_Phone 2.2.2.5/32
set security zones security-zone internal address-book address gatekeeper 2.2.2.10/32
set security policies from-zone internal to-zone external policy P1 match source-address
  any
set security policies from-zone internal to-zone external policy P1 match
  destination-address IP_Phone
set security policies from-zone internal to-zone external policy P1 match application
  junos-h323
set security policies from-zone internal to-zone external policy P1 then permit
set security policies from-zone internal to-zone external policy P2 match source-address
  any
set security policies from-zone internal to-zone external policy P2 match
  destination-address gatekeeper
set security policies from-zone internal to-zone external policy P2 match application
  junos-h323
set security policies from-zone internal to-zone external policy P2 then permit
set security policies from-zone external to-zone internal policy P3 match source-address
  IP_Phone
set security policies from-zone external to-zone internal policy P3 match
  destination-address any
set security policies from-zone external to-zone internal policy P3 match application
  junos-h323
set security policies from-zone external to-zone internal policy P3 then permit
set security policies from-zone external to-zone internal policy P4 match source-address
  gatekeeper
set security policies from-zone external to-zone internal policy P4 match
  destination-address any
set security policies from-zone external to-zone internal policy P4 match application
  junos-h323
set security policies from-zone external to-zone internal policy P4 then permit
```

Step-by-Step Procedure

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

To configure the device to pass H.323 ALG traffic to a gatekeeper in the external zone:

1. Configure two address books.

```
[edit]
user@host# set security zones security-zone external address-book address
  IP_Phone 2.2.2.5/32
user@host# set security zones security-zone internal address-book address
  gatekeeper 2.2.2.10/32
```

2. Configure policy P1 from the internal zone to the external zone.

```
[edit]
user@host# set security policies from-zone internal to-zone external policy P1 match
  source-address any
user@host# set security policies from-zone internal to-zone external policy P1 match
  destination-address IP_Phone
user@host# set security policies from-zone internal to-zone external policy P1 match
  application junos-h323
user@host# set security policies from-zone internal to-zone external policy P1 then
  permit
```


Copyright © 2016, Juniper Networks, Inc.

3. Configure policy P2 to allow traffic between the internal zone and the gatekeeper in the external zone.

   ```
   [edit]
   user@host# set security policies from-zone internal to-zone external policy P2 match
     source-address any
   user@host# set security policies from-zone internal to-zone external policy P2 match
     destination-address gatekeeper
   user@host# set security policies from-zone internal to-zone external policy P2 match
     application junos-h323
   user@host# set security policies from-zone internal to-zone external policy P2 then
     permit
   ```

4. Configure policy P3 to allow traffic between phones in the internal zone and the external zone.

   ```
   [edit]
   user@host# set security policies from-zone external to-zone internal policy P3 match
     source-address IP_Phone
   user@host# set security policies from-zone external to-zone internal policy P3 match
     destination-address any
   user@host# set security policies from-zone external to-zone internal policy P3 match
     application junos-h323
   user@host# set security policies from-zone external to-zone internal policy P3 then
     permit
   ```

5. Configure policy P4 to allow traffic between phones in the internal zone and the gatekeeper in the external zone.

   ```
   [edit]
   user@host# set security policies from-zone external to-zone internal policy P4
     match source-address gatekeeper
   user@host# set security policies from-zone external to-zone internal policy P4
     match destination-address any
   user@host# set security policies from-zone external to-zone internal policy P4
     match application junos-h323
   user@host# set security policies from-zone external to-zone internal policy P4 then
     permit
   ```

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

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

```
[edit]
user@host# show security policies
...
from-zone internal to-zone external {
    policy P1 {
        match {
            source-address any;
            destination-address IP_Phone;
            application junos-h323;
        }
        then {
            permit;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    policy P2 {
        match {
            source-address any;
            destination-address gatekeeper;
            application junos-h323;
        }
        then {
            permit;
        }
    }
}
from-zone external to-zone internal {
    policy P3 {
        match {
            source-address IP_Phone;
            destination-address any;
            application junos-h323;
        }
        then {
            permit;
        }
    }
    policy P4 {
        match {
            source-address gatekeeper;
            destination-address any;
            application junos-h323;
        }
        then {
            permit;
        }
    }
}
...
```

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

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying H.323 ALG Configurations on page 158

Verifying H.323 ALG Configurations

Purpose Display information about active calls.

Copyright © 2016, Juniper Networks, Inc.

NOTE: H.323 counters for calls and active calls in the output to this **show security** command do not apply to the proprietary Avaya implementation of H.323. This is because Q.931 setup and connect messages are exchanged right after the phone is powered up and call creation and tear down is done by Facility messages.

Counters for calls and active calls are increased when the resources allocated for calls are increased—that is, messages belonging to the same call and that pass the firewall multiple times increment the counters. This applies when resources for a call need to be allocated multiple times. For example, in a two-zone scenario the setup and connect message pair allocates one call resource, and the active call counter is increased by one. But in a three-zone scenario the setup and connect message pair passes the firewall twice, each time allocating different call resources. In this case, the counter is incremented.

Action From the J-Web interface, select **Monitor>ALGs>H323**. Alternatively, from the CLI, enter the **show security alg h323 counters** command.

Counters for H.245 messages received also will not be accurate in the case of H.245 tunneling. Because H.245 messages are encapsulated in Q.931 packets, the counter for H.245 messages received will remain zero even when there are H.245 messages. The **Other H245** counter will, however, reflect these packet transmissions.

```
[edit]
user@host> show security alg h323 counters
H.323 counters summary:
  Packets received         : 0
  Packets dropped          : 0
  RAS message received     : 0
  Q.931 message received   : 0
  H.245 message received   : 0
  Number of calls          : 0
  Number of active calls   : 0
H.323 error counters:
  Decoding errors          : 0
  Message flood dropped    : 0
  NAT errors               : 0
  Resource manager errors  : 0
H.323 message counters:
  RRQ       : 0
  RCF       : 0
  ARQ       : 0
  ACF       : 0
  URQ       : 0
  UCF       : 0
  DRQ       : 0
  DCF       : 0
  Oth RAS   : 0
  Setup     : 0
  Alert     : 0
  Connect   : 0
  CallProd  : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Info       : 0
RelCmpl    : 0
Facility   : 0
Empty      : 0
OLC        : 0
OLC-ACK    : 0
Oth H245   : 0
```

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149

Example: Using NAT with the H.323 ALG to Enable Incoming Calls

This example shows how to configure NAT with the H.323 ALG to enable calls from a public to a private network.

- Requirements on page 160
- Overview on page 160
- Configuration on page 161
- Verification on page 166

Requirements

Before you begin, understand H.323 ALGs. See "Understanding H.323 ALG" on page 145.

Overview

In a two-zone scenario with a server in the private zone, you can use NAT for incoming calls by configuring a NAT pool on the interface to the public zone.

In this example (see Figure 13), IP-Phone1 and a server called gatekeeper are in the private zone, and IP-Phone2 is in the public zone. You configure a static nat rule set and a source NAT pool to do NAT. You also create two policies, private-to-public and public-to-private, to permit ALG H.323 traffic from and to the private and public zones.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 13 shows NAT with the H.323 ALG incoming calls.

Figure 13: NAT with the H.323 ALG—Incoming Calls

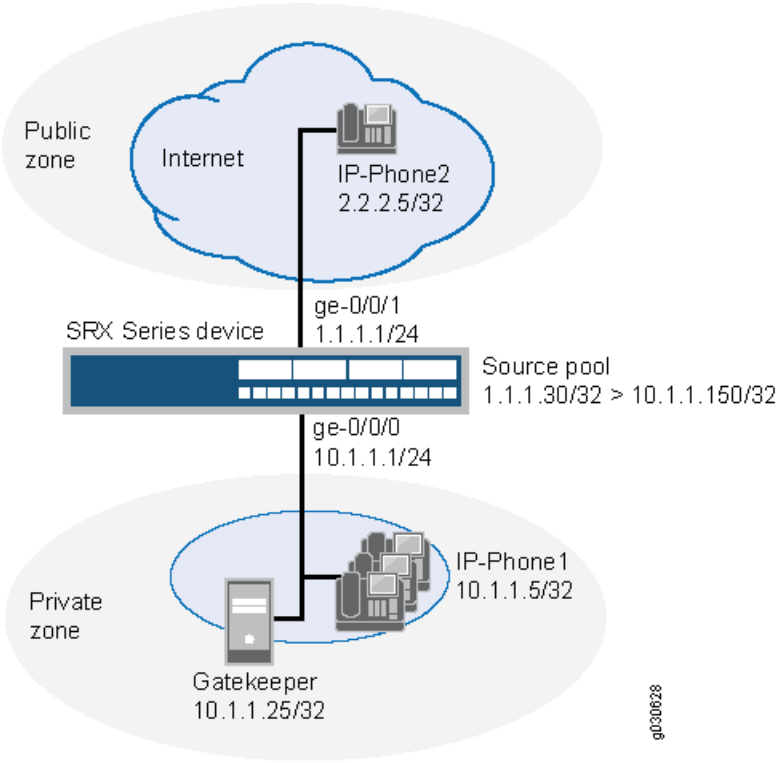

In this example, you configure source NAT as follows:

- Create a static NAT rule set called gatekeeper with a rule called gatekeeper to match packets from the public zone with the destination address 1.1.1.25/32. For matching packets, the destination IP address is translated to the private address 10.1.1.25/32.
- Define a source NAT pool called h323-nat-pool to contain the IP address range from 1.1.1.30/32 through 1.1.1.150/32.
- Create a source NAT rule set called h323-nat with rule h323-r1 to match packets from the private zone to the public zone with the source IP address 10.1.1.0/24. For matching packets, the source address is translated to the IP address in h323-nat-pool.
- Configure proxy ARP for the addresses 1.1.1.30/32 through 1.1.1.150/32 on interface ge-0/0/1.0. This allows the system to respond to ARP requests received on the interface for these addresses.

Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

Copyright © 2016, Juniper Networks, Inc.

network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address IP-Phone1 10.1.1.5/32
set security zones security-zone private address-book address gatekeeper 10.1.1.25/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address IP-Phone2 2.2.2.5/32
set security zones security-zone public interfaces ge-0/0/1.0
set security nat source pool h323-nat-pool address 1.1.1.30/32 to 1.1.1.150/32
set security nat source address-persistent
set security nat source rule-set h323-nat from zone private
set security nat source rule-set h323-nat to zone public
set security nat source rule-set h323-nat rule h323-r1 match source-address 10.1.1.0/24
set security nat source rule-set h323-nat rule h323-r1 then source-nat pool h323-nat-pool
set security nat proxy-arp interface ge-0/0/1.0 address 1.1.1.30/32 to 1.1.1.150/32
set security policies from-zone private to-zone public policy private-to-public match
  source-address IP-Phone1
set security policies from-zone private to-zone public policy private-to-public match
  source-address gatekeeper
set security policies from-zone private to-zone public policy private-to-public match
  destination-address IP-Phone2
set security policies from-zone private to-zone public policy private-to-public match
  application junos-h323
set security policies from-zone private to-zone public policy private-to-public then permit
set security policies from-zone public to-zone private policy public-to-private match
  source-address IP-Phone2
set security policies from-zone public to-zone private policy public-to-private match
  destination-address IP-Phone1
set security policies from-zone public to-zone private policy public-to-private match
  destination-address gatekeeper
set security policies from-zone public to-zone private policy public-to-private match
  application junos-h323
set security policies from-zone public to-zone private policy public-to-private then permit
```

Step-by-Step Procedure

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

To configure NAT with H.323 ALG to enable calls from a public to a private network:

1. Configure interfaces.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
   user@host# set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
   ```

2. Configure zones and assign addresses to them.

   ```
   [edit security zones]
   user@host# set security-zone private interfaces ge-0/0/0.0
   user@host# set security-zone private address-book address IP-Phone1 10.1.1.5/32
   user@host# set security-zone private address-book address gatekeeper 10.1.1.25/32
   user@host# set security-zone public interfaces ge-0/0/1.0
   user@host# set security-zone public address-book address IP-Phone2 2.2.2.5/32
   ```


Copyright © 2016, Juniper Networks, Inc.

3. Create a static NAT rule set.

```
[edit security nat static rule-set ip-phones]
user@host# set from zone public
user@host# set match destination-address 1.1.1.25/32
user@host# set then static-nat prefix 10.1.1.25/32
```

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/1.0 address 1.1.1.25/32
```

5. Configure a source NAT rule set.

```
[edit security nat]
set source pool h323-nat-pool address 1.1.1.30/32 to 1.1.1.150/32
set source address-persistent
set source rule-set h323-nat from zone private
set source rule-set h323-nat to zone public
set source rule-set h323-nat rule h323-r1 match source-address 10.1.1.0/24
set source rule-set h323-nat rule h323-r1 then source-nat pool h323-nat-pool
set proxy-arp interface ge-0/0/1.0 address 1.1.1.30/32 to 1.1.1.150/32
```

6. Configure policies for outgoing traffic.

```
[edit security policies from-zone private to-zone public policy private-to-public]
user@host# set match source-address IP-Phone1
user@host# set match source-address gatekeeper
user@host# set match destination-address IP-Phone2
user@host# set match application junos-h323
user@host# set then permit
```

7. Configure policies for incoming traffic.

```
[edit security policies from-zone public to-zone private policy public-to-private]
user@host# set match source-address IP-Phone2
user@host# set match destination-address IP-Phone1
user@host# set match destination-address gatekeeper
user@host# set match application junos-h323
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
    ge-0/0/1 {
      unit 0 {
        family inet {
```

Copyright © 2016, Juniper Networks, Inc.

```
            address 1.1.1.1/24;
          }
        }
      }
[edit]
user@host# show security zones
security-zone private {
  address-book {
    address IP-Phone1 10.1.1.5/32;
    address gatekeeper 10.1.1.25/32;
  }
  interfaces {
    ge-0/0/0.0;
  }
}
  security-zone public {
    address-book {
      address IP-Phone2 2.2.2.5/32;
    }
    interfaces {
      ge-0/0/1.0;
    }
  }
[edit]
user@host# show security nat
source {
  pool h323-nat-pool {
    address {
      1.1.1.30/32 to 1.1.1.150/32;
    }
  }
  address-persistent;
  rule-set h323-nat {
    from zone private;
    to zone public;
    rule h323-r1 {
      match {
        source-address 10.1.1.0/24;
      }
      then {
        source-nat {
          pool {
            h323-nat-pool;
          }
        }
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/1.0 {
      address {
        1.1.1.30/32 to 1.1.1.150/32;
      }
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    static {
      rule-set ip-phones {
        from zone public;
        rule gatekeeper {
          match {
            destination-address 1.1.1.25/32;
          }
          then {
            static-nat prefix 10.1.1.25/32;
          }
        }
      }
    }
      proxy-arp {
        interface ge-0/0/1.0 {
          address {
            1.1.1.25/32;
          }
        }
      }
[edit]
user@host# show security policies
from-zone private to-zone public {
  policy private-to-public {
    match {
      source-address [IP-Phone1 gatekeeper];
      destination-address IP-Phone2;
      application junos-h323;
    }
    then {
      permit;
    }
  }
}
  from-zone public to-zone private {
    policy public-to-private {
      match {
        source-address IP-Phone2;
        destination-address [IP-Phone1 gatekeeper];
        application junos-h323;
      }
      then {
        permit;
      }
    }
  }
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying H.323 ALG Status on page 166
- Verifying Security ALG H.323 Counters on page 166
- Verifying Source NAT Rule Usage on page 167

Verifying H.323 ALG Status

Purpose Verify that H.323 ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Enabled
  MGCP     : Enabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the H323 ALG status as follows:

- Enabled—Shows the H323 ALG is enabled.
- Disabled—Shows the H323 ALG is disabled.

Verifying Security ALG H.323 Counters

Purpose Verify that there is a security counters for ALG H323.

Action From operational mode, enter the **show security alg h323 counters** command.

```
user@host> show security alg h323 counters
```

```
H.323 counters summary:
Packets received :4060
Packets dropped :24
RAS message received :3690
Q.931 message received :202
H.245 message received :145
Number of calls :25
Number of active calls :0

H.323 Error Counters:
```


Copyright © 2016, Juniper Networks, Inc.

```
Decoding errors :24
Message flood dropped :0
NAT errors :0
Resource manager errors :0

H.323 Message Counters:
RRQ : 431
RCF : 49
ARQ : 60
ACF : 33
URQ : 34
UCF : 25
DRQ : 55
DCF : 44
oth RAS : 2942
Setup : 28
Alert : 9
Connect : 25
CallPrcd : 18
Info : 0
RelCmpl : 39
Facility : 14
Progress : 0
Empty : 65
OLC : 20
OLC-ACK : 20
```

Meaning The sample output gives the rundown of security ALG H.323 counters expressing that, there are security counters for ALG H323.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat source rule all
source NAT rule: h323-r1      Rule-set: h323-nat
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : private
  To zone                   : public
  Match
    Source addresses        : 0.0.0.0         - 255.255.255.255
    Destination port        : 0               - 0
  Action                      : interface
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
    Max session number        : 0
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 154
- H.323 ALG Configuration Overview on page 149

Example: Using NAT with the H.323 ALG to Enable Outgoing Calls

This example shows how to configure static NAT with H.323 ALG to enable calls from a private to a public network.

- Requirements on page 168
- Overview on page 168
- Configuration on page 169
- Verification on page 173

Requirements

Before you begin, understand the H.323 ALG and its processes. See "Understanding H.323 ALG" on page 145.

Overview

In this example (see Figure 14), IP-Phone 1 and a server called gatekeeper are in the private zone and IP-Phone2 is in the public zone. You configure static NAT to enable IP-Phone1 and gatekeeper to call IP-Phone2 in the public zone. You then create a policy called public-to-private to allow ALG H.323 traffic from the public zone to the private zone and a policy called private-to-public to allow ALG H.323 traffic from the private zone to the public zone.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 14 shows NAT with the H.323 ALG outgoing calls.

Figure 14: NAT with the H.323 ALG—Outgoing Calls

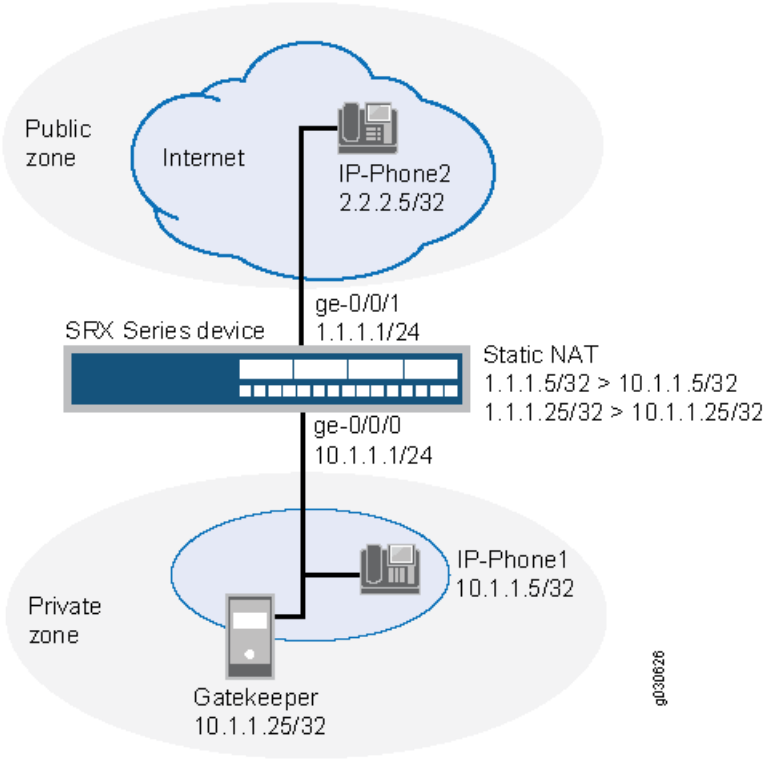

In this example, you configure static NAT as follows:

- Create a static NAT rule set called ip-phones with a rule called phone1 to match packets from the public zone with the destination address 1.1.1.5/32. For matching packets, the destination IP address is translated to the private address 10.1.1.5/32.
- Define a second rule called gatekeeper to match packets from the public zone with the destination address 1.1.1.25/32. For matching packets, the destination IP address is translated to the private address 10.1.1.25/32.
- Create proxy ARP for the addresses 1.1.1.5/32 and 1.1.1.25/32 on interface ge-0/0/1. This allows the system to respond to ARP requests received on the specified interface for these addresses.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
```

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address IP-Phone1 10.1.1.5/32
set security zones security-zone private address-book address gatekeeper 10.1.1.25/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address IP-Phone2 2.2.2.5/32
set security zones security-zone public interfaces ge-0/0/1.0
set security nat static rule-set ip-phones from zone public
set security nat static rule-set ip-phones rule phone1 match destination-address 1.1.1.5/32
set security nat static rule-set ip-phones rule phone1 then static-nat prefix 10.1.1.5/32
set security nat static rule-set ip-phones rule gatekeeper match destination-address
  1.1.1.25/32
set security nat static rule-set ip-phones rule gatekeeper then static-nat prefix 10.1.1.25/32
set security nat proxy-arp interface ge-0/0/1.0 address 1.1.1.5/32
set security nat proxy-arp interface ge-0/0/1.0 address 1.1.1.25/32
set security policies from-zone public to-zone private policy public-to-private match
  source-address IP-Phone2
set security policies from-zone public to-zone private policy public-to-private match
  destination-address gatekeeper
set security policies from-zone public to-zone private policy public-to-private match
  application junos-h323
set security policies from-zone public to-zone private policy public-to-private then permit
set security policies from-zone private to-zone public policy private-to-public match
  source-address IP-Phone1
set security policies from-zone private to-zone public policy private-to-public match
  source-address gatekeeper
set security policies from-zone private to-zone public policy private-to-public match
  destination-address IP-Phone2
set security policies from-zone private to-zone public policy private-to-public match
  application junos-h323
set security policies from-zone private to-zone public policy private-to-public then permit
```

Step-by-Step Procedure

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

To configure static NAT with the H.323 ALG to enable calls from a private to a public network:

1. Configure interfaces.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
```

2. Create zones and assign addresses to them.

```
[edit security zones]
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone public interfaces ge-0/0/1.0
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone private address-book address IP-Phone1 10.1.1.5/32
user@host# set security-zone private address-book address gatekeeper 10.1.1.25/32
user@host# set security-zone public interfaces ge-0/0/1.0
user@host# set security-zone public address-book address IP-Phone2 2.2.2.5/32
```

3. Configure static NAT rule set with rules.

```
[edit security nat static rule-set ip-phones]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from zone public
user@host# set rule phone1 match destination-address 1.1.1.5/32
user@host# set rule phone1 then static-nat prefix 10.1.1.5/32
user@host# set rule gatekeeper match destination-address 1.1.1.25/32
user@host# set rule gatekeeper then static-nat prefix 10.1.1.25/32
```

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/1 address 1.1.1.5/32
user@host# set proxy-arp interface ge-0/0/1 address 1.1.1.25/32
```

5. Configure a security policy for incoming traffic.

```
[edit security policies from-zone public to-zone private policy public-to-private]
user@host# set match source-address IP-Phone2
user@host# set match destination-address gatekeeper
user@host# set match application junos-h323
user@host# set then permit
```

6. Configure a security policy for outgoing traffic.

```
[edit security policies from-zone private to-zone public policy private-to-public]
user@host# set match source-address IP-Phone1
user@host# set match source-address gatekeeper
user@host# set match destination-address IP-Phone2
user@host# set match application junos-h323
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
    ge-0/0/1 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }
[edit]
user@host# show security zones
  security-zone private {
    address-book {
      address IP-Phone1 10.1.1.5/32;
      address gatekeeper 10.1.1.25/32;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
        interfaces {
            ge-0/0/0.0;
        }
    }
        security-zone public {
            address-book {
                address IP-Phone2 2.2.2.5/32;
            }
            interfaces {
                ge-0/0/1.0;
            }
        }
[edit]
user@host# show security nat
    static {
        rule-set ip-phones {
            from zone public;
            rule phone1 {
                match {
                    destination-address 1.1.1.5/32;
                }
                then {
                    static-nat prefix 10.1.1.5/32;
                }
            }
            rule gatekeeper {
                match {
                    destination-address 1.1.1.25/32;
                }
                then {
                    static-nat prefix 10.1.1.25/32;
                }
            }
        }
    }
        proxy-arp {
            interface ge-0/0/1.0 {
                address {
                    1.1.1.5/32;
                    1.1.1.25/32;
                }
            }
        }
[edit]
user@host# show security policies
    from-zone public to-zone private {
        policy public-to-private {
            match {
                source-address IP-Phone2;
                destination-address gatekeeper;
                application junos-h323;
            }
            then {
                permit;
            }
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
  from-zone private to-zone public {
    policy private-to-public {
      match {
        source-address [ IP-Phone1 gatekeeper ];
        destination-address IP-Phone2;
        application junos-h323;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying H.323 ALG Status on page 173
- Verifying Security ALG H.323 Counters on page 173

Verifying H.323 ALG Status

Purpose Verify that H.323 ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Enabled
  MGCP     : Enabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Enabled
  RTSP     : Enabled
  SCCP     : Enabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the H323 ALG status as follows:

- Enabled—Shows the H323 ALG is enabled.
- Disabled—Shows the H323 ALG is disabled.

Verifying Security ALG H.323 Counters

Purpose Verify that there is a security counters for ALG H323.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security alg h323 counters** command.

```
user@host> show security alg h323 counters

H.323 counters summary:
Packets received :4060
Packets dropped :24
RAS message received :3690Q.931 message received :202
H.245 message received :145
Number of calls :25
Number of active calls :0

H.323 Error Counters:
Decoding errors :24
Message flood dropped :0
NAT errors :0
Resource manager errors :0

H.323 Message Counters:
RRQ : 431
RCF : 49
ARQ : 60
ACF : 33
URQ : 34
UCF : 25
DRQ : 55
DCF : 44
oth RAS : 2942
Setup : 28
Alert : 9
Connect : 25
CallPrcd : 18
Info : 0
RelCmpl : 39
Facility : 14
Progress : 0
Empty : 65
OLC : 20
OLC-ACK : 20
```

Meaning The sample output gives the synopsis of security ALG H.323 counters expressing that there are security counters for ALG H.323.

Related Documentation

- Example: Passing H.323 ALG Traffic to a Gatekeeper in the External Zone on page 154
- H.323 ALG Configuration Overview on page 149

Understanding H.323 ALG Endpoint Registration Timeouts

In Network Address Translation (NAT) mode, when endpoints in the protected network behind the Juniper Networks device register with the H.323 gatekeeper, the device adds an entry to the NAT table containing a mapping of the public-to-private address for each endpoint. These entries make it possible for endpoints in the protected network to receive incoming calls.

Copyright © 2016, Juniper Networks, Inc.

You set an endpoint registration timeout to specify how long an endpoint registration entry remains in the NAT table. To ensure uninterrupted incoming call service, set the endpoint registration timeout to a value equal to or greater than the keepalive value the administrator configures on the gatekeeper. The range is 10 to 50,000 seconds, the default value is 3600 seconds.

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149
- Example: Setting H.323 ALG Endpoint Registration Timeouts on page 175

Example: Setting H.323 ALG Endpoint Registration Timeouts

This example shows how to specify the endpoint registration timeout.

- Requirements on page 175
- Overview on page 175
- Configuration on page 175
- Verification on page 176

Requirements

Before you begin, understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

Overview

You set an endpoint registration timeout range to specify how long an endpoint registration entry remains in the NAT table. The range is 10 to 50,000 seconds, and the default value is 3600 seconds.

Configuration

GUI Step-by-Step Procedure

To specify the H.323 ALG endpoint registration timeout:

1. Select **Configure>Security>ALG**.
2. Select the **H323** tab.
3. In the Timeout for endpoints box, type **5000**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

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

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

Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security alg h323 counters** command.

Related Documentation

- Understanding H.323 ALG Endpoint Registration Timeouts on page 174
- H.323 ALG Configuration Overview on page 149

Understanding H.323 ALG Media Source Port Ranges

The media source port feature enables you to configure the device to allow media traffic on a narrow or wide range of ports. By default, the device listens for H.323 traffic on a wide range of ports. If your endpoint equipment allows you to specify a sending port and a listening port, you might want to narrow the range of ports the device allows media traffic on. This enhances security by opening a smaller pinhole for H.323 traffic.

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149
- Example: Setting H.323 ALG Media Source Port Ranges on page 176

Example: Setting H.323 ALG Media Source Port Ranges

This example shows how to enable the H.323 ALG media source port feature.

- Requirements on page 176
- Overview on page 176
- Configuration on page 177
- Verification on page 177

Requirements

Before you begin, understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

Overview

The media source port feature enables you to configure the device to allow media traffic on a narrow or wide range of ports. By default, the device listens for H.323 traffic on a narrow range of ports. This example shows how to configure the device to open a wide gate for media traffic by enabling the media source port feature.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To enable the H.323 ALG media source port feature:

1. Select **Configure>Security>ALG**.
2. Select the **H323** tab.
3. Select the **Enable Permit media from any source port** check box.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To enable the H.323 ALG media source port feature:

1. Set a narrow gate for media traffic by disabling the media source port for the H.323 ALG.

   ```
   [edit]
   user@host# delete security alg h323 media-source-port-any
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg h323 counters** command.

Related Documentation

- Understanding H.323 ALG Media Source Port Ranges on page 176
- H.323 ALG Configuration Overview on page 149

Understanding H.323 ALG DoS Attack Protection

You can protect the H.323 gatekeeper from denial-of-service (DoS) flood attacks by limiting the number of Registration, Admission, and Status (RAS) messages per second it will attempt to process. Incoming RAS request messages exceeding the threshold you specify are dropped by H.323 Application Layer Gateway (ALG). The range is 2 to 50,000 messages per second, the default value is 1000.

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149
- Example: Configuring H.323 ALG DoS Attack Protection on page 178

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring H.323 ALG DoS Attack Protection

This example shows how to configure the H.323 ALG DoS attack protection feature.

- Requirements on page 178
- Overview on page 178
- Configuration on page 178
- Verification on page 179

Requirements

Before you begin, understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

Overview

You can protect the H.323 gatekeeper from DoS flood attacks by limiting the range of Registration, Admission, and Status (RAS) messages per second it will attempt to process. The range is 2 to 50,000 messages per second, and the default value is 1000. This example limits the number of incoming RAS request messages to 5000 messages per second.

Configuration

GUI Step-by-Step Procedure

To configure the H.323 ALG DoS attack protection feature:

1. Select **Configure>Security>ALG**.
2. Select the **H323** tab.
3. In the Message flood gatekeeper threshold box, type **5000**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure the H.323 ALG DoS attack protection feature:

1. Configure the gatekeeper for the H.323 ALG and set the threshold.

```
[edit]
user@host# set security alg h323 application-screen message-flood gatekeeper
  threshold 5000
```

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

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


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security alg h323 counters** command.

Related Documentation

- Understanding H.323 ALG DoS Attack Protection on page 177
- H.323 ALG Configuration Overview on page 149

Understanding H.323 ALG Unknown Message Types

This feature enables you to specify how unidentified H.323 messages are handled by the device. The default is to drop unknown (unsupported) messages.

You can protect the H.323 gatekeeper from denial-of-service (DoS) flood attacks by limiting the number of Registration, Admission, and Status (RAS) messages per second it will attempt to process. Incoming RAS request messages exceeding the threshold you specify are dropped by the H.323 Application Layer Gateway (ALG). The range is 2 to 50,000 messages per second, the default value is 1000.

We do not recommend permitting unknown messages because they can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown H.323 messages can help you get your network operational, so that you can analyze your voice-over-IP (VoIP) traffic to determine why some messages were being dropped. The unknown H.323 message type feature enables you to configure the device to accept H.323 traffic containing unknown message types in both Network Address Translation (NAT) mode and route mode.

NOTE: **This option applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.**

Related Documentation

- Understanding H.323 ALG on page 145
- H.323 ALG Configuration Overview on page 149
- Example: Allowing Unknown H.323 ALG Message Types on page 179

Example: Allowing Unknown H.323 ALG Message Types

This example shows how to configure the device to allow unknown H.323 message types in both route and NAT modes.

- Requirements on page 180
- Overview on page 180

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 180
- Verification on page 181

Requirements

Before you begin, understand and configure any Avaya H.323-specific features. See the *Administrator Guide for Avaya Communication Manager*, *Avaya IP Telephony Implementation Guide*, and *Avaya Application Solutions IP Telephony Deployment Guide* at http://support.avaya.com.

Overview

This feature enables you to specify how unidentified H.323 messages are handled by the device. The default is to drop unknown (unsupported) messages. The Enable Permit NAT applied option and the **permit-nat-applied** configuration statement specify that unknown messages be allowed to pass if the session is in NAT mode. The Enable Permit routed option and the **permit-routed** configuration statement specify that unknown messages be allowed to pass if the session is in route mode. (Sessions in transparent mode are treated as route mode.)

Configuration

GUI Step-by-Step Procedure

To configure the device to allow unknown H.323 message types in both route and NAT modes:

1. Select **Configure>Security>ALG**.
2. Select the **H323** tab.
3. Select the **Enable Permit NAT applied** check box.
4. Select the **Enable Permit routed** check box.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure the device to allow unknown H.323 message types in both route and NAT modes:

1. Specify that unknown messages be allowed to pass if the session is in NAT mode.

   ```
   [edit]
   user@host# set security alg h323 application-screen unknown-message
     permit-nat-applied
   ```

2. Specify that unknown messages be allowed to pass if the session is in route mode.

   ```
   [edit]
   user@host# set security alg h323 application-screen unknown-message
     permit-routed
   ```

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

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


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security alg h323** command and the **show security alg h323 counters** command.

Related Documentation

- Understanding H.323 ALG Unknown Message Types on page 179
- H.323 ALG Configuration Overview on page 149

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 16

Configuring the MGCP ALG

- Understanding the MGCP ALG on page 183
- MGCP ALG Configuration Overview on page 189
- Example: Configuring Media Gateways in Subscriber Homes Using MGCP ALGs on page 189
- Example: Configuring Three-Zone ISP-Hosted Service Using MGCP ALG and NAT on page 196
- Understanding MGCP ALG Call Duration and Timeouts on page 208
- Example: Setting MGCP ALG Call Duration on page 209
- Example: Setting MGCP ALG Inactive Media Timeout on page 210
- Example: Setting MGCP ALG Transaction Timeout on page 211
- Understanding MGCP ALG DoS Attack Protection on page 212
- Example: Configuring MGCP ALG DoS Attack Protection on page 213
- Understanding MGCP ALG Unknown Message Types on page 214
- Example: Allowing Unknown MGCP ALG Message Types on page 214

Understanding the MGCP ALG

The Media Gateway Control Protocol (MGCP) is a text-based Application Layer protocol used for call setup and call control between the media gateway and the media gateway controller (MGC).

The protocol is based on a master/slave call control architecture: the MGC (call agent) maintains call control intelligence, and media gateways carry out the instructions from the call agent. Both signaling packets and media packets are transmitted over UDP. Junos OS supports MGCP in route mode and Network Address Translation (NAT) mode.

The MGCP Application Layer Gateway (ALG) performs the following procedures:

- Conducts voice-over-IP (VoIP) signaling payload inspection. The payload of the incoming VoIP signaling packet is fully inspected based on related RFCs and proprietary standards. Any malformed packet attack is blocked by the ALG.
- Conducts MGCP signaling payload inspection. The payload of the incoming MGCP signaling packet is fully inspected in accordance with RFC 3435. Any malformed-packet attack is blocked by the ALG.

Copyright © 2016, Juniper Networks, Inc.

- Provides stateful processing. The corresponding VoIP-based state machines are invoked to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- Performs NAT. Any embedded IP address and port information in the payload is properly translated based on the existing routing information and network topology, and is then replaced with the translated IP address and port number, if necessary.
- Manages pinholes for VoIP traffic. To keep the VoIP network secure, the IP address and port information used for media or signaling is identified by the ALG, and any needed pinhole is dynamically created and closed during call setup.

This topic contains the following sections:

- MGCP Security on page 184
- Entities in MGCP on page 184
- Commands on page 186
- Response Codes on page 188

MGCP Security

The MGCP ALG includes the following security features:

- Denial-of-service (DoS) attack protection. The ALG performs stateful inspection at the UDP packet level, the transaction level, and the call level. MGCP packets matching the RFC 3435 message format, transaction state, and call state, are processed. All other messages are dropped.
- Security policy enforcement between gateway and gateway controller (signaling policy).
- Security policy enforcement between gateways (media policy).
- Per-gateway MGCP message flooding control. Any malfunctioning or hacked gateway will not disrupt the whole VoIP network. Combined with per-gateway flooding control, damage is contained within the impacted gateway.
- Per-gateway MGCP connection flooding control.
- Seamless switchover/failover if calls, including calls in progress, are switched to the standby firewall in case of system failure.

Entities in MGCP

There are four basic entities in MGCP:

- Endpoint on page 185
- Connection on page 185
- Call on page 185
- Call Agent on page 185

Copyright © 2016, Juniper Networks, Inc.

Endpoint

A media gateway is a collection of endpoints. An endpoint can be an analog line, trunk, or any other access point. An endpoint contains the following elements:

```
local-endpoint-name@domain-name
```

The following examples are some valid endpoint IDs:

```
group1/Trk8@example.net
group2/Trk1/*@[192.168.10.8] (wild-carding)
$@example.net (any endpoint within the media gateway)
*@example.net (all endpoints within the media gateway)
```

Connection

Connections are created on each endpoint by an MG during call setup. A typical VoIP call involves two connections. A complex call, for example a three-party call or conference call, might require more connections. The MGC can instruct media gateways to create, modify, delete, and audit a connection.

A connection is identified by its connection ID, which is created by the MG when it is requested to create a connection. Connection ID is presented as a hexadecimal string, and its maximum length is 32 characters.

Call

A call is identified by its call ID, which is created by the MGC when establishing a new call. Call ID is a hexadecimal string with a maximum length of 32 characters. Call ID is unique within the MGC. Two or more connections can have the same call ID if they belong to the same call.

Call Agent

One or more call agents (also called media gateway controllers) are supported in MGCP to enhance reliability in the VoIP network. The following two examples are of call agent names:

```
CallAgent@voipCA.example.com
voipCA.example.com
```

Several network addresses can be associated under one domain name in the Domain Name System (DNS). By keeping track of the time to live (TTL) of DNS query/response data and implementing retransmission using other alternative network addresses, switchover and failover is achieved in MGCP.

The concept of a *notified entity* is essential in MGCP. The notified entity for an endpoint is the call agent currently controlling that endpoint. An endpoint should send any MGCP command to its notified entity. However, different call agents might send MGCP commands to this endpoint.

The notified entity is set to a provisioned value upon startup, but can be changed by a call agent through the use of the **NotifiedEntity** parameter contained in an MGCP message. If the notified entity for an endpoint is empty or has not been set explicitly, its value

Copyright © 2016, Juniper Networks, Inc.

defaults to the source address of the last successful non-audit MGCP command received for that endpoint.

Commands

The MGCP protocol defines nine commands for controlling endpoints and connections. All commands are composed of a command header, optionally followed by Session Description Protocol (SDP) information. A command header has the following elements:

- A command line: command verb + transaction ID + endpointId + MGCP version.
- Zero or more parameter lines, composed of a parameter name followed by a parameter value.

Table 6 lists supported MGCP commands and includes a description of each, the command syntax, and examples. Refer to RFC 2234 for a complete explanation of command syntax.

Table 6: MGCP Commands

Command	Description	Command Syntax	Example
EPCF	EndpointConfiguration—Used by a call agent to inform a gateway of coding characteristics (a-law or mu-law) expected by the line side of the endpoint.	**ReturnCode [PackageList] EndpointConfiguration (EndpointId,[BearerInformation])**	**EPCF 2012 wxx/T2@example.com MGCP 1.0B: e:mu**
CRCX	CreateConnection—Used by a call agent to instruct the gateway to create a connection with, and endpoint inside, the gateway.	**ReturnCode, [ConnectionId,] [SpecificEndPointId,] [LocalConnectionDescriptor,] [SecondEndPointId,] [SecondConnectionId,] [PackageList] CreateConnection (CallId, EndpointId, [NotifiedEntity,] [LocalConnectionOption,] Mode, [{RemoteConnectionDescriptor | SecondEndpoindId},] [encapsulated RQNT,] [encapsulated EPCF])**	**CRCX 1205 aaln/1@gw-25.example.net MGCP 1.0C: A3C47F21456789F0L: p:10, a:PCMUM: sendrecvX: 0123456789ADR: L/hdS: L/rgv=0o=- 25678 753849 IN IP4 128.96.41.1s=-c=IN IP4 128.96.41.1t=0 0m=audio 3456 RTP/AVP 0**
MDCX	ModifyConnection—Used by a call agent to instruct a gateway to change the parameters for an existing connection.	**ReturnCode, [LocalConnectionDescriptor,] [PackageList] ModifyConnection (CallId, EndpointId, ConnectionId, [NotifiedEntity,] [LocalConnectionOption,] [Mode,]** **[RemoteConnectionDescriptor,] [encapsulated RQNT,] [encapsulated EPCF])**	**MDCX 1210 aaln/1@rgw-25.example.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8M: recvonlyX: 0123456789AER: L/huS: G/rtv=0o=- 4723891 7428910 IN IP4 128.96.63.25s=-c=IN IP4 128.96.63.25t=0 0m=audio 3456 RTP/AVP 0**

Copyright © 2016, Juniper Networks, Inc.

Table 6: MGCP Commands *(continued)*

Command	Description	Command Syntax	Example
DLCX	DeleteConnection—Used by a call agent to instruct a gateway to delete an existing connection. DeleteConnection can also be used by a gateway to release a connection that can no longer be sustained.	**ReturnCode, ConnectionParameters, [PackageList] DeleteConnection (CallId, EndpointId, ConnectionId, [NotifiedEntity,] [encapsulated RQNT,] [encapsulated EPCF])**	Example 1: MGC -> MG **DLCX 9210 aaln/1@rgw-25.example.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8** Example 2: MG -> MGC **DLCX 9310 aaln/1@rgw-25.example.net MGCP 1.0C: A3C47F21456789F0I: FDE234C8E: 900 - Hardware errorP: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48**
RQNT	NotificationRequest command—Used by a call agent to instruct an MG to monitor for certain event(s) or signal(s) for a specific endpoint.	**ReturnCode, [PackageList] NotificationRequest[(EndpointId, [NotifiedEntity,] [RequestedEvents,] RequestIdentifier, [DigitMap,] [SignalRequests,] [QuarantineHandling,] [DetectEvents,] [encapsulated EPCF])**	**RQNT 1205 aaln/1@rgw-25.example.net MGCP 1.0N: ca-new@callagent-ca.example.netX: 0123456789AAR: L/hd(A, E(S(L/dl),R(L/oc,L/hu,D/[0-9#*T](D))))D: (0T|00T|xx|91xxxxxxxxxx|9011x.T)S:T: G/ft**
NTFY	Notify—Used by a gateway to inform the call agent when requested event(s) or signal(s) occur.	**ReturnCode, [PackageList] Notify (EndpointID, [NotifiedEntity,] RequestIdentifier, ObservedEvents)**	**NTFY 2002 aaln/1@rgw-25.example.net MGCP 1.0N: ca@ca1.example.net:5678X: 0123456789ACO: L/hd,D/9,D/1,D/2,D/0,D/1,D/8,D/2,D/9,D/4, D/2,D/6,D/6**
AUEP	AuditEndpoint—Used by a call agent to audit the status of the endpoint.	**ReturnCode, EndPointIdList, | { [RequestedEvents,] [QuarantineHandling,] [DigitMap,] [SignalRequests,] [RequestedIdentifier,] [NotifiedEntity,] [ConnectionIdentifier,] [DetectEvents,] [ObservedEvents,] [EventStats,] [BearerInformation,] [BearerMethod,] [RestartDelay,] [ReasonCode,] [MaxMGCPDatagram,] [Capabilities]} [PackageList] AuditEndpoint (EndpointId, [RequestedInfo])**	Example 1: **AUEP 1201 aaln/1@rgw-25.example.netMGCP 1.0F: A, R,D,S,X,N,I,T,O** Example 2: **AUEP 1200 *@rgw-25.example.net MGCP 1.0**
AUCX	AuditConnection—Used by a call agent to collect the parameters applied to a connection.	**ReturnCode, [CallId,] [NotifiedEntity,] [LocalConnectionOptions,] [Mode,] [RemoteConnectionDescriptor,] [LocalConnectionDescriptor,] [ConnectionParameters,] [PackageList] AuditConnection (EndpointId, ConnectionId, RequestedInfo)**	**AUCX 3003 aaln/1@rgw-25.example.net MGCP 1.0I: 32F345E2F: C,N,L,M,LC,P**

Copyright © 2016, Juniper Networks, Inc.

Table 6: MGCP Commands *(continued)*

Command	Description	Command Syntax	Example
RSIP	RestareInProgress—Used by a gateway to notify a call agent that one or more endpoints are being taken out of service or placed back in service.	**ReturnCode, [NotifiedEntity,] [PackageList] RestartInProgress (EndpointId, RestartMethod, [RestartDelay,] [ReasonCode])**	**RSIP 5200 aaln/1@rg2-25.example.net MGCP 1.0RM: gracefulRD: 300**

Response Codes

Every command sent by the calling agent or gateway, whether successful or not, requires a response code. The response code is in the header of the response message, and optionally is followed by session description information.

The response header is composed of a response line, followed by zero or more parameter lines, each containing a parameter name letter followed by its value. The response header is composed of a three-digit response code, transaction ID, and optionally followed by commentary. The response header in the following response message shows response code 200 (successful completion), followed by ID 1204 and the comment:OK.

```
200 1204 OK
I: FDE234C8
v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=audio 3456 RTP/AVP 96
a=rtpmap:96 G726-32/8000
```

The ranges of response codes are defined as follows:

- 000 — 099 indicate a response acknowledgement.
- 100 — 199—indicate a provisional response.
- 200 — 299 indicate a successful completion (final response).
- 400 — 499 indicate a transient error (final response).
- 500 — 599 indicate a permanent error (final response).

Refer to RFC 3661 for detailed information about response codes.

A response to a command is sent to the source address of the command, not to the current notified entity. A media gateway can receive MGCP commands from various network addresses simultaneously, and send back responses to corresponding network addresses. However, it sends all MGCP commands to its current notified entity.

Related Documentation

- ALG Overview on page 3
- MGCP ALG Configuration Overview on page 189

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Media Gateways in Subscriber Homes Using MGCP ALGs on page 189
- Example: Configuring Three-Zone ISP-Hosted Service Using MGCP ALG and NAT on page 196

MGCP ALG Configuration Overview

The Media Gateway Control Protocol (MGCP ALG) is enabled by default on the device—no action is required to enable it. However, you might choose to fine-tune MGCP ALG operations by using the following instructions:

1. Free up bandwidth when calls fail to properly terminate. See "Example: Setting MGCP ALG Call Duration" on page 209.
2. Control how long a call can remain active without any media traffic. See "Example: Setting MGCP ALG Inactive Media Timeout" on page 210.
3. Track and clear signaling traffic when it times out. See "Example: Setting MGCP ALG Transaction Timeout" on page 211.
4. Protect the media gateway from denial-of-service (DoS) flood attacks. See "Example: Configuring MGCP ALG DoS Attack Protection" on page 213.
5. Enable unknown messages to pass when the session is in Network Address Translation (NAT) mode and route mode. See "Example: Allowing Unknown MGCP ALG Message Types" on page 214.

Related Documentation

- Understanding the MGCP ALG on page 183
- Example: Configuring Media Gateways in Subscriber Homes Using MGCP ALGs on page 189
- Example: Configuring Three-Zone ISP-Hosted Service Using MGCP ALG and NAT on page 196

Example: Configuring Media Gateways in Subscriber Homes Using MGCP ALGs

This example shows how to configure media gateways in subscriber homes using MGCP ALGs.

- Requirements on page 189
- Overview on page 190
- Configuration on page 191
- Verification on page 194

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Configure zones. See "Example: Creating Security Zones" on page 1051.
- Configure addresses and interfaces. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Configure security policies. See "Security Policies Configuration Overview" on page 1093.

Overview

When a cable service provider offers MGCP services to residential subscribers, they locate the Juniper Networks device and call agent on their premises and install a set-top box, in each subscriber's home. The set-top boxes act as gateways for the residences.

After creating zones—external_subscriber for the customer and internal_ca for the service provider—you configure addresses, then interfaces, and finally policies to allow signaling between endpoints. Note that although gateways frequently reside in different zones, requiring policies for media traffic, in this example both gateways are in the same subnet. Note also that because RTP traffic between the gateways never passes through the device, no policy is needed for the media. See Figure 15.

Figure 15: Media Gateway in Subscriber Homes

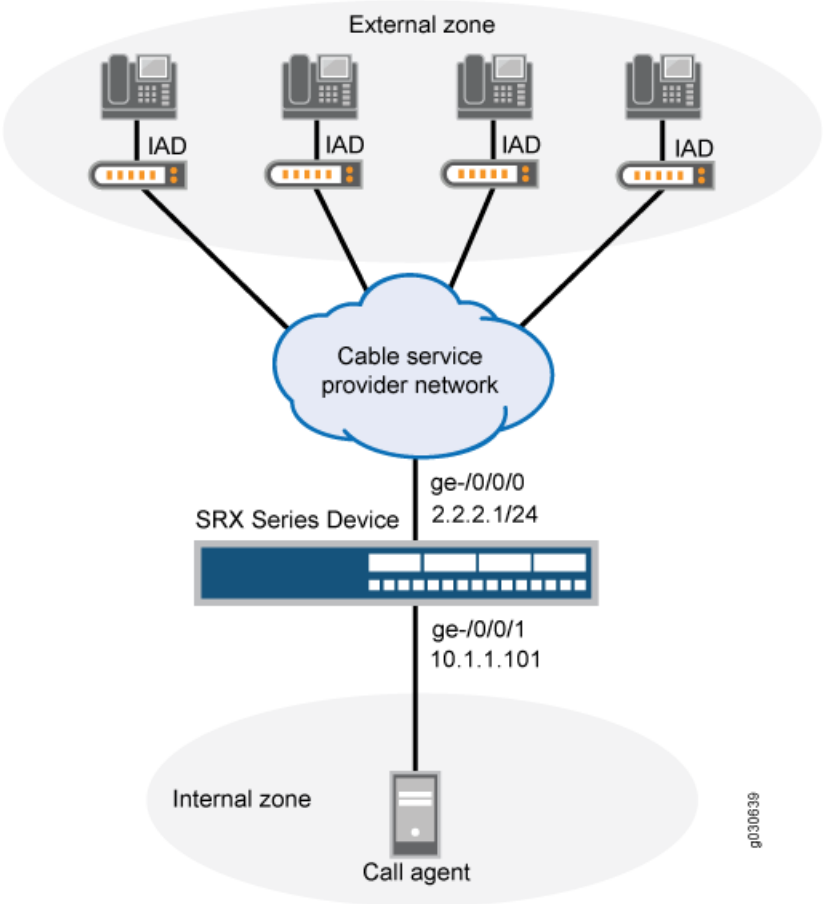

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

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

```
set security zones security-zone external-subscriber host-inbound-traffic system-services
  all
set security zones security-zone external-subscriber host-inbound-traffic protocols all
set security zones security-zone internal-ca host-inbound-traffic system-services all
set security zones security-zone internal-ca host-inbound-traffic protocols all
set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
set interfaces ge-0/0/0 unit 0 family inet
set security zones security-zone external-subscriber interfaces ge-0/0/0
set security zones security-zone internal-ca interfaces ge-0/0/1
set security address-book book1 address ca-agent 110.1.1.101/32
set security address-book book1 attach zone internal-ca
set security address-book book2 address subscriber-subnet 2.2.2.1/24
set security address-book book2 attach zone external-subscriber
set security policies from-zone internal-ca to-zone external-subscriber policy
  ca-to-subscribers match source-address ca-agent-1
set security policies from-zone internal-ca to-zone external-subscriber policy
  ca-to-subscribers match destination-address subscriber-subnet
set security policies from-zone internal-ca to-zone external-subscriber policy
  ca-to-subscribers match application junos-mgcp
set security policies from-zone internal-ca to-zone external-subscriber policy
  ca-to-subscribers then permit
set security policies from-zone external-subscriber to-zone internal-ca policy
  subscriber-to-ca match source-address subscriber-subnet
set security policies from-zone external-subscriber to-zone internal-ca policy
  subscriber-to-ca match destination-address ca-agent-1
set security policies from-zone external-subscriber to-zone internal-ca policy
  subscriber-to-ca match application junos-mgcp
set security policies from-zone external-subscriber to-zone internal-ca policy
  subscriber-to-ca then permit
set security policies from-zone internal-ca to-zone internal-ca policy intra-ca match
  source-address any
set security policies from-zone internal-ca to-zone internal-ca policy intra-ca match
  destination-address any
set security policies from-zone internal-ca to-zone internal-ca policy intra-ca match
  application any
set security policies from-zone internal-ca to-zone internal-ca policy intra-ca then permit
set security policies from-zone external-subscriber to-zone external-subscriber policy
  intra-subscriber match source-address any
set security policies from-zone external-subscriber to-zone external-subscriber policy
  intra-subscriber match destination-address any
set security policies from-zone external-subscriber to-zone external-subscriber policy
  intra-subscriber match application any
set security policies from-zone external-subscriber to-zone external-subscriber policy
  intra-subscriber then permit
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure media gateways in subscriber homes using MGCP ALGs:

1. Create security zones for the customer and the service provider.

```
[edit security zones security-zone external-subscriber]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all

[edit security zones security-zone internal-ca]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
```

2. Configure interfaces for the zones.

```
[edit]
user@host# edit security zones security-zone external-subscriber interfaces
  ge-0/0/0
user@host# set interfaces ge-0/0/0 unit 0 family inet
user@host# set security zones security-zone internal-ca interfaces ge-0/0/1
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
```

3. Configure address books and attach zones to them.

```
[edit security address-book book1]
user@host# set address ca-agent 110.1.1.101/32
user@host# set attach zone internal-ca

[edit security address-book book2]
user@host# set address subscriber-subnet 2.2.2.1/24
user@host# set attach zone external-subscriber
```

4. Configure policies for traffic from the internal to the external zone.

```
[edit security policies from-zone internal-ca to-zone external-subscriber policy
  ca-to-subscribers]
user@host# edit match source-address ca-agent-1
user@host# set match destination-address subscriber-subnet
user@host# set match application junos-mgcp
user@host# set then permit
```

5. Configure policies for traffic from the external to the internal zone.

```
[edit security policies from-zone external-subscriber to-zone internal-ca policy
  subscriber-to-ca]
user@host# edit match source-address subscriber-subnet
user@host# set match destination-address ca-agent-1
user@host# set match application junos-mgcp
user@host# set then permit
```

6. Configure policies for traffic between two internal zones.

```
[edit security policies from-zone internal-ca to-zone internal-ca policy intra-ca]
user@host# edit match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set then permit
```


Copyright © 2016, Juniper Networks, Inc.

7. Configure policies for traffic between two external zones.

```
[edit security policies from-zone external-subscriber to-zone external-subscriber
  policy intra-subscriber]
user@host# edit match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set then permit
```

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

```
from-zone internal-ca to-zone external-subscriber {
  policy ca-to-subscribers {
    match {
      source-address ca-agent-1;
      destination-address subscriber-subnet;
      application junos-mgcp;
    }
    then {
      permit;
    }
  }
}
from-zone external-subscriber to-zone internal-ca {
  policy subscriber-to-ca {
    match {
      source-address subscriber-subnet;
      destination-address ca-agent-1;
      application junos-mgcp;
    }
    then {
      permit;
    }
  }
}
from-zone internal-ca to-zone internal-ca {
  policy intra-ca {
    match {
      ssource-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone external-subscriber to-zone external-subscriber {
  policy intra-subscriber {
    match {
      source-address any;
      destination-address any;
      application any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          permit;
        }
      }
    }
```

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

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying MGCP ALGs on page 194
- Verifying MGCP ALG Calls on page 194
- Verifying MGCP ALG Endpoints on page 195
- Verifying MGCP ALG Counters on page 195

Verifying MGCP ALGs

Purpose Verify the MGCP ALG verification options.

Action From operational mode, enter the **show security alg mgcp ?** command.

```
user@host> show security alg mgcp ?
Possible completions:
  calls                Show MGCP calls
  counters             Show MGCP counters
  endpoints            Show MGCP endpoints
```

Meaning The output shows a list of all MGCP verification parameters. Verify the following information:

- All MGCP calls
- Counters for all MGCP calls
- Information about all MGCP endpoints

Verifying MGCP ALG Calls

Purpose Verify information about active MGCP calls.

Action From operational mode, enter the **show security alg mgcp calls** command.

```
user@host> show security alg mgcp calls
Endpoint@GW                    Zone         Call ID                   RM Group
d001@101.50.10.1               Trust        10d55b81140e0f76          512
   Connection Id> 0
   Local SDP>  o: 101.50.10.1                x_o: 101.50.10.1
               c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
   Remote SDP> c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
Endpoint@GW                    Zone         Call ID                   RM Group
d001@3.3.3.5                   Untrust      3a104e9b41a7c4c9          511
   Connection Id> 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Local SDP>  o: 3.3.3.5                    x_o: 3.3.3.5
            c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
Remote SDP> c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
```

Meaning The output displays information about all MGCP calls. Verify the following information:

- Endpoint
- Zone
- Call identifier
- Resource Manager group

Verifying MGCP ALG Endpoints

Purpose Verify information about MGCP endpoints.

Action From operational mode, enter the **show security alg mgcp endpoints** command.

```
user@host> show security alg mgcp endpoints
Gateway: 101.50.10.1 Zone: Trust IP: 101.50.10.1 -> 101.50.10.1
   Endpoint             Trans #  Call #   Notified Entity
   d001                 1        1        0.0.0.0/0->0.0.0.0/0
Gateway: 3.3.3.5 Zone: Untrust IP: 3.3.3.5 -> 3.3.3.5
   Endpoint             Trans #  Call #   Notified Entity
   d001                 1        1        0.0.0.0/0->0.0.0.0/0
```

Meaning The output displays information about all MGCP endpoints. Verify the following information:

- Gateway IP address and zone of both endpoints
- Endpoint identifier, transaction number, call number, and notified entity for each gateway

Verifying MGCP ALG Counters

Purpose Verify information about MGCP counters.

Action From operational mode, enter the **show security alg mgcp counters** command.

```
user@host> show security alg mgcp counters
MGCP counters summary:
Packets received             :284
Packets dropped              :0
Message received             :284
Number of connections        :4
Number of active connections :3
Number of calls              :4
Number of active calls       :3
Number of transactions       :121
Number of active transactions:52
Number of re-transmission    :68
MGCP Error Counters:
Unknown-method               :0
Decoding error               :0
```

Copyright © 2016, Juniper Networks, Inc.

```
Transaction error            :0
Call error                   :0
Connection error             :0
Connection flood drop        :0
Message flood drop           :0
IP resolve error             :0
NAT error                    :0
Resource manager error       :0
MGCP Packet Counters:
CRCX     :4       MDCX     :9       DLCX     :2
AUEP     :1       AUCX     :0       NTFY     :43
RSIP     :79      EPCF     :0       RQNT     :51
000-199  :0       200-299  :95      300-999  :0
```

Meaning The output displays information about all MGCP counters. Verify the following information:

- Summary of MGCP counters
- MGCP error counters
- MGCP packet counters

Related Documentation

- Understanding the MGCP ALG on page 183
- MGCP ALG Configuration Overview on page 189

Example: Configuring Three-Zone ISP-Hosted Service Using MGCP ALG and NAT

This example shows how to configure a three-zone configuration using MGCP ALG and NAT.

- Requirements on page 196
- Overview on page 196
- Configuration on page 198
- Verification on page 206

Requirements

Before you begin, understand NAT support with MGCP ALG. See "Understanding the MGCP ALG" on page 183.

Overview

Typically, a three-zone configuration is used when an ISP in one geographical location provides service to two networks in different geographical locations.

In this example (see Figure 16), an ISP located on the USA West Coast provides MGCP service to customers in separate networks in Asia and San Francisco. Asia customers are in the asia-3 zone and are supported by the asia-gw gateway; San Francisco customers are in the sf-2 zone and are supported by the sf-gw gateway. A call agent, west-ca, is in the DMZ. The gateways and the call agent are listed in Table 7, showing the corresponding IP address, interface, and zone.

Copyright © 2016, Juniper Networks, Inc.

In this example, after creating zones and setting addresses for the gateways and the call agent, you associate the zones to interfaces, and then configure static NAT to the call agent and source NAT for communication from an IP phone in the sf-2 zone to phones in the asia-3 zone. You also configure a policy between the zones to allow the communication.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 16 shows a three-zone ISP-hosted service.

Figure 16: Three-Zone ISP-Hosted Service

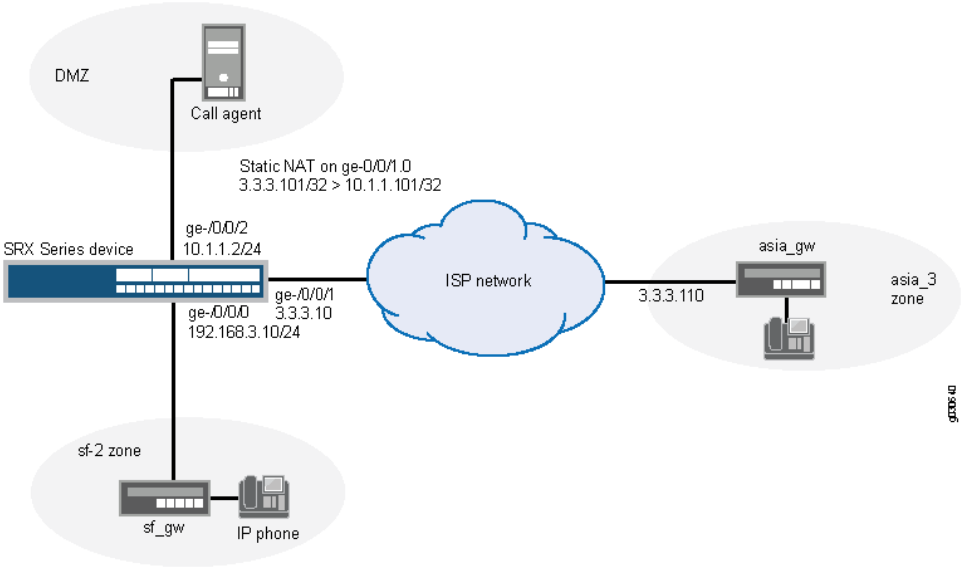

Table 7: Three-Zone ISP-Host Service

Gateway	IP Address	Interface	Zone
sf-gw	192.168.3.201	ge-0/0/0	sf-2
asia-gw	3.3.3.101	ge-0/0/1	asia-3
west-ca	10.1.1.101	ge-0/0/2	DMZ

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 192.168.3.10/24
set interfaces ge-0/0/1 unit 0 family inet address 3.3.3.10/24
set interfaces ge-0/0/2 unit 0 family inet address 10.1.1.2/24
set security zones security-zone sf-2 interfaces ge-0/0/0.0
set security zones security-zone asia-3 interfaces ge-0/0/1.0
set security zones security-zone dmz interfaces ge-0/0/2.0
set security address-book book1 address sf-gw 192.168.3.201/32
set security address-book book1 attach zone sf-2
set security address-book book2 address asia-gw 3.3.3.101/32
set security address-book book2 attach zone asia-3
set security address-book book3 address west-ca 10.1.1.101/32
```


Copyright © 2016, Juniper Networks, Inc.

```
set security address-book book3 attach zone dmz
set security nat source pool ip-phone-pool address 3.3.3.20/32
set security nat source rule-set phones from zone sf-2
set security nat source rule-set phones to zone asia-3
set security nat source rule-set phones rule phone1 match source-address 192.168.3.10/32
set security nat source rule-set phones rule phone1 match destination 3.3.3.101/32
set security nat source rule-set phones rule phone1 then source-nat pool ip-phone-pool
set security nat static rule-set to-callagent from zone asia-3
set security nat static rule-set to-callagent rule phone1 match destination-address
  3.3.3.101/32
set security nat static rule-set to-callagent rule phone1 then static-nat prefix 10.1.1.101/32
set security nat proxy-arp interface ge-0/0/1.0 address 3.3.3.101/32
set security nat proxy-arp interface ge-0/0/1.0 address 3.3.3.20/32
set security policies from-zone dmz to-zone asia-3 policy pol-dmz-to-asia-3 match
  source-address west-ca
set security policies from-zone dmz to-zone asia-3 policy pol-dmz-to-asia-3 match
  destination-address asia-gw
set security policies from-zone dmz to-zone asia-3 policy pol-dmz-to-asia-3 match
  application junos-mgcp
set security policies from-zone dmz to-zone asia-3 policy pol-dmz-to-asia-3 then permit
set security policies from-zone asia-3 to-zone dmz policy pol-asia-3-to-dmz match
  source-address asia-gw
set security policies from-zone asia-3 to-zone dmz policy pol-asia-3-to-dmz match
  destination-address west-ca
set security policies from-zone asia-3 to-zone dmz policy pol-asia-3-to-dmz match
  application junos-mgcp
set security policies from-zone asia-3 to-zone dmz policy pol-asia-3-to-dmz then permit
set security policies from-zone sf-2 to-zone dmz policy pol-sf-2-to-dmz match
  source-address sf-gw
set security policies from-zone sf-2 to-zone dmz policy pol-sf-2-to-dmz match
  destination-address west-ca
set security policies from-zone sf-2 to-zone dmz policy pol-sf-2-to-dmz match application
  junos-mgcp
set security policies from-zone sf-2 to-zone dmz policy pol-sf-2-to-dmz then permit
set security policies from-zone dmz to-zone sf-2 policy pol-dmz-to-sf-2 match
  source-address west-ca
set security policies from-zone dmz to-zone sf-2 policy pol-dmz-to-sf-2 match
  destination-address sf-gw
set security policies from-zone dmz to-zone sf-2 policy pol-dmz-to-sf-2 match application
  junos-mgcp
set security policies from-zone dmz to-zone sf-2 policy pol-dmz-to-sf-2 then permit
set security policies from-zone sf-2 to-zone asia-3 policy pol-sf-2-to-asia-3 match
  source-address sf-gw
set security policies from-zone sf-2 to-zone asia-3 policy pol-sf-2-to-asia-3 match
  destination-address asia-gw
set security policies from-zone sf-2 to-zone asia-3 policy pol-sf-2-to-asia-3 match
  application junos-mgcp
set security policies from-zone sf-2 to-zone asia-3 policy pol-sf-2-to-asia-3 then permit
set security policies from-zone asia-3 to-zone sf-2 policy pol-asia-3-to-sf-2 match
  source-address asia-gw
set security policies from-zone asia-3 to-zone sf-2 policy pol-asia-3-to-sf-2 match
  destination sf-gw
set security policies from-zone asia-3 to-zone sf-2 policy pol-asia-3-to-sf-2 match
  application junos-mgcp
set security policies from-zone asia-3 to-zone sf-2 policy pol-asia-3-to-sf-2 then permit
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone sf-2 to-zone sf-2 policy pol-intra-sf-2 match source-address
  any
set security policies from-zone sf-2 to-zone sf-2 policy pol-intra-sf-2 match
  destination-address any
set security policies from-zone sf-2 to-zone sf-2 policy pol-intra-sf-2 match application
  any
set security policies from-zone sf-2 to-zone sf-2 policy pol-intra-sf-2 then permit
set security policies from-zone asia-3 to-zone asia-3 policy pol-intra-asia-3 match
  source-address any
set security policies from-zone asia-3 to-zone asia-3 policy pol-intra-asia-3 match
  destination-address any
set security policies from-zone asia-3 to-zone asia-3 policy pol-intra-asia-3 match
  application any
set security policies from-zone asia-3 to-zone asia-3 policy pol-intra-asia-3 then permit
```

Step-by-Step Procedure

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

To configure a three-zone configuration using MGCP ALG and NAT:

1. Configure interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 192.168.3.10/24
user@host# set interfaces ge-0/0/1 unit 0 family inet address 3.3.3.10/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 10.1.1.2/24
```

2. Create security zones.

```
[edit security zones]
user@host# set security-zone sf-2 interfaces ge-0/0/0
user@host# set security-zone asia-3 interfaces ge-0/0/1
user@host# set security-zone dmz interfaces ge-0/0/2
```

3. Create address books and assign zones to them.

```
[edit security address-book book1]
user@host# set address sf-gw 192.168.3.201/32
user@host# set attach zone sf-2

[edit security address-book book2]
user@host# set address asia-gw 3.3.3.101/32
user@host# set attach zone asia-3

[edit security address-book book3]
user@host# set address west-ca 10.1.1.101/32
user@host# set attach zone dmz
```

4. Create a static NAT rule set and set the match conditions and actions for it.

```
[edit security nat static rule-set to-callagent]
user@host# set from zone asia-3
user@host# set rule phone1 match destination-address 3.3.3.101/32
user@host# set rule phone1 then static-nat prefix 10.1.1.101/32
```

5. Configure proxy ARP for address **3.3.3.101/32** on interface **ge-0/0/1.0**.

```
[edit security nat ]
user@host# set proxy-arp interface ge-0/0/1.0 address 3.3.3.101/32
```


Copyright © 2016, Juniper Networks, Inc.

6. Create a source NAT pool.

   ```
   [edit security nat]
   user@host# set source pool ip-phone-pool address 3.3.3.20/32
   ```

7. Create a source NAT rule set and set the match conditions and actions for it.

   ```
   [edit security nat source rule-set phones]
   user@host# set from zone sf-2
   user@host# set to zone asia-3
   user@host# set rule phone1 match source-address 192.168.3.10/32
   user@host# set rule phone1 match destination-address 3.3.3.101/32
   user@host# set rule phone1 then source-nat pool ip-phone-pool
   ```

8. Configure proxy ARP for address **3.3.3.20/32** on interface **ge-0/0/1.0**.

   ```
   [edit security nat ]
   user@host# set proxy-arp interface ge-0/0/1.0 address 3.3.3.20/32
   ```

9. Configure a policy to allow traffic from DMZ to Asia.

   ```
   [edit security policies from-zone dmz to-zone asia-3 policy pol-dmz-to-asia-3]
   user@host# set match source-address west-ca
   user@host# set match destination-address asia-gw
   user@host# set match application junos-mgcp
   user@host# set then permit
   ```

10. Configure a policy to allow traffic from Asia to DMZ.

    ```
    [edit security policies from-zone asia-3 to-zone dmz policy pol-asia-3-to-dmz]
    user@host# set match source-address asia-gw
    user@host# set match destination-address west-ca
    user@host# set match application junos-mgcp
    user@host# set then permit
    ```

11. Configure a policy to allow traffic from San Francisco to DMZ.

    ```
    [edit security policies from-zone sf-2 to-zone dmz policy pol-sf-2-to-dmz]
    user@host# set match source-address sf-gw
    user@host# set match destination-address west-ca
    user@host# set match application junos-mgcp
    user@host# set then permit
    ```

12. Configure a policy to allow traffic from DMZ to San Francisco.

    ```
    [edit security policies from-zone dmz to-zone sf-2 policy pol-dmz-to-sf-2]
    user@host# set match source-address west-ca
    user@host# set match destination-address sf-gw
    user@host# set match application junos-mgcp
    user@host# set then permit
    ```

13. Configure a policy to allow traffic from San Francisco to Asia.

    ```
    [edit security policies from-zone sf-2 to-zone asia-3 policy pol-sf-2-to-asia-3]
    user@host# set match source-address sf-gw
    user@host# set match destination-address asia-gw
    user@host# set match application junos-mgcp
    user@host# set then permit
    ```

14. Configure a policy to allow traffic from Asia to San Francisco.

    ```
    [edit security policies from-zone asia-3 to-zone sf-2 policy pol-asia-3-to-sf-2]
    ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set match source-address asia-gw
user@host# set match destination-address sf-gw
user@host# set match application junos-mgcp
user@host# set then permit
```

15. Configure a policy to allow traffic on devices within San Francisco.

```
[edit security policies from-zone sf-2 to-zone sf-2 policy pol-intra-sf-2]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set then permit
```

16. Configure a policy to allow traffic on devices within Asia.

```
[edit security policies from-zone asia-3 to-zone asia-3 policy pol-intra-asia-3]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security address-book**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.3.10/24;
      }
    }
  }
    ge-0/0/1 {
      unit 0 {
        family inet {
          address 3.3.3.10/24;
        }
      }
    }
    ge-0/0/2 {
      unit 0 {
        family inet {
          address 10.1.1.2/24;
        }
      }
    }
[edit]
user@host# show security zones
security-zone sf-2 {
  interfaces {
    ge-0/0/0.0;
  }
}
```

 Copyright © 2016, Juniper Networks, Inc.

```
    security-zone asia-3 {
      interfaces {
        ge-0/0/1.0;
      }
    }
    security-zone dmz {
      interfaces {
        ge-0/0/2.0;
      }
    }
[edit]
user@host# show security address-book
book1 {
  address sf-gw 192.168.3.201/32;
  attach {
    zone sf-2;
  }
}
  book2 {
    address asia-gw 3.3.3.101/32;
    attach {
      zone asia-3;
    }
  }
  book3 {
    address west-ca 10.1.1.101/32;
    attach {
      zone dmz;
    }
  }
[edit]
user@host# show security nat
source {
  pool ip-phone-pool {
    address {
      3.3.3.20/32;
    }
  }
  rule-set phones {
    from zone sf-2;
    to zone asia-3;
    rule phone1 {
      match {
        source-address 192.168.3.10/32;
        destination-address 3.3.3.101/32;
      }
      then {
        source-nat {
          pool {
            ip-phone-pool;
          }
        }
      }
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
    static {
      rule-set to-callagent {
        from zone asia-3;
        rule phone1 {
          match {
            destination-address 3.3.3.101/32;
          }
          then {
            static-nat prefix 10.1.1.101/32;
          }
        }
      }
    }
    proxy-arp {
      interface ge-0/0/1.0 {
        address {
          3.3.3.101/32;
          3.3.3.20/32;
        }
      }
    }
[edit]
user@host# show security policies
from-zone dmz to-zone asia-3 {
    policy pol-dmz-to-asia-3 {
      match {
        source-address west-ca;
        destination-address asia-gw;
        application junos-mgcp;
      }
      then {
        permit;
      }
    }
    }
    from-zone asia-3 to-zone dmz {
      policy pol-asia-3-to-dmz {
        match {
          source-address asia-gw;
          destination-address west-ca;
          application junos-mgcp;
        }
        then {
          permit;
        }
      }
    }
    from-zone sf-2 to-zone dmz {
      policy pol-sf-2-to-dmz {
        match {
          source-address sf-gw;
          destination-address west-ca;
          application junos-mgcp;
        }
        then {
          permit;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
from-zone dmz to-zone sf-2 {
  policy pol-dmz-to-sf-2 {
    match {
      source-address west-ca;
      destination-address sf-gw;
      application junos-mgcp;
    }
    then {
      permit;
    }
  }
}
from-zone sf-2 to-zone asia-3 {
  policy pol-sf-2-to-asia-3 {
    match {
      source-address sf-gw;
      destination-address asia-gw;
      application junos-mgcp;
    }
    then {
      permit;
    }
  }
}
from-zone asia-3 to-zone sf-2 {
  policy pol-asia-3-to-sf-2 {
    match {
      source-address asia-gw;
      destination-address sf-gw;
      application junos-mgcp;
    }
    then {
      permit;
    }
  }
}
from-zone sf-2 to-zone sf-2 {
  policy pol-intra-sf-2 {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone asia-3 to-zone asia-3 {
  policy pol-intra-asia-3 {
    match {
      source-address any;
      destination-address any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        application any;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying MGCP ALG on page 206
- Verifying MGCP Calls on page 206
- Verifying MGCP ALG Statistics on page 207
- Verifying MGCP Endpoints on page 207

Verifying MGCP ALG

Purpose Verify if the MGCP ALG is enabled.

Action From operational mode, enter the **show security alg status | match mgcp** command.

```
user@host> show security alg status | match mgcp
MGCP      : Enabled
```

Meaning The output shows the MGCPALG status as follows:

- Enabled—Shows the MGCP ALG is enabled.
- Disabled—Shows the MGCP ALG is disabled.

Verifying MGCP Calls

Purpose Verify the MGCP calls that are currently active.

Action From operational mode, enter the **show security alg mgcp calls** command.

```
user@host> show security alg mgcp calls
Endpoint@GW                    Zone         Call ID                   RM Group
d001@101.50.10.1               Trust        10d55b81140e0f76          512
   Connection Id> 0
   Local SDP>  o: 101.50.10.1                x_o: 101.50.10.1
               c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
   Remote SDP> c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
Endpoint@GW                    Zone         Call ID                   RM Group
d001@3.3.3.5                   Untrust      3a104e9b41a7c4c9          511
   Connection Id> 0
   Local SDP>  o: 3.3.3.5                    x_o: 3.3.3.5
               c: 3.3.3.5/16928              x_c: 3.3.3.5/16928
   Remote SDP> c: 101.50.10.1/32206          x_c: 101.50.10.1/32206
```


Copyright © 2016, Juniper Networks, Inc.

Meaning The output displays information about all MGCP calls. Verify the following information:

- Endpoint
- Zone
- Call identifier
- Resource Manager group

Verifying MGCP ALG Statistics

Purpose Verify the MGCP ALG statistics.

Action From operational mode, enter the **show security alg mgcp counters** command.

```
user@host> show security alg mgcp counters
```

```
MGCP counters summary:
Packets received             :284
Packets dropped              :0
Message received             :284
Number of connections        :4
Number of active connections :3
Number of calls              :4
Number of active calls       :3
Number of transactions       :121
Number of active transactions:52
Number of re-transmission    :68
MGCP Error Counters:
Unknown-method               :0
Decoding error               :0
Transaction error            :0
Call error                   :0
Connection error             :0
Connection flood drop        :0
Message flood drop           :0
IP resolve error             :0
NAT error                    :0
Resource manager error       :0
MGCP Packet Counters:
CRCX     :4       MDCX     :9       DLCX     :2
AUEP     :1       AUCX     :0       NTFY     :43
RSIP     :79      EPCF     :0       RQNT     :51
000-199  :0       200-299  :95      300-999  :0
```

Meaning The output displays information about all MGCP counters. Verify the following information:

- Summary of MGCP counters
- MGCP error counters
- MGCP packet counters

Verifying MGCP Endpoints

Purpose Verify the MGCP endpoints.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security alg mgcp endpoints** command.

```
user@host> show security alg mgcp endpoints

Gateway: 101.50.10.1 Zone: Trust IP: 101.50.10.1 -> 101.50.10.1
   Endpoint             Trans #  Call #   Notified Entity
   d001                 1        1        0.0.0.0/0->0.0.0.0/0
Gateway: 3.3.3.5 Zone: Untrust IP: 3.3.3.5 -> 3.3.3.5
   Endpoint             Trans #  Call #   Notified Entity
   d001                 1        1        0.0.0.0/0->0.0.0.0/0
```

Meaning The output displays information about all MGCP endpoints. Verify the following information:

- Gateway IP address and zone of both endpoints
- Endpoint identifier, transaction number, call number, and notified entity for each gateway

Related Documentation

- Static NAT Configuration Overview on page 4949
- Understanding Source NAT on page 4863

Understanding MGCP ALG Call Duration and Timeouts

The call duration feature gives you control over Media Gateway Control Protocol (MGCP) call activity and helps you to manage network resources.

Typically a Delete Connection (DLCX) message will be sent out to delete a connection. The MCGP Application Layer Gateway (ALG) intercepts it and removes all media sessions for that connection.

A call can have one or more voice channels. Each voice channel has two sessions (or two media streams), one for Real-Time Transport Protocol (RTP) traffic and one for Real-Time Control Protocol (RTCP) signaling. When managing the sessions, the device considers the sessions in each voice channel as one group. Timeouts and call duration settings apply to a group as opposed to each session.

The following parameters govern MGCP call activity:

- **maximum-call-duration**—This parameter sets the absolute maximum length of a call. When a call exceeds this parameter setting, the MGCP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, and the range is 3 through 720 minutes. This setting also frees up bandwidth in cases where calls fail to properly terminate.
- **inactive-media-timeout**—This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the MGCP ALG gates opened for media are closed. The default setting is 120 seconds, and the range is 10 through 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

 Copyright © 2016, Juniper Networks, Inc.

NOTE: The **inactive-media-timeout** value should be less than the **maximum-call-duration** value.

- **transaction-timeout**—A transaction is a command and its mandatory response. For example, an NTFY from the gateway to the call agent or a 200 OK from the call agent to the gateway. The Juniper Networks device tracks these transactions and clears them when they time out. The timeout range for MGCP transactions is 3 through 50 seconds and the default is 30 seconds.

Related Documentation

- Understanding the MGCP ALG on page 183
- MGCP ALG Configuration Overview on page 189
- Example: Setting MGCP ALG Call Duration on page 209
- Example: Setting MGCP ALG Inactive Media Timeout on page 210
- Example: Setting MGCP ALG Transaction Timeout on page 211

Example: Setting MGCP ALG Call Duration

This example shows how to set call duration for the MGCP ALG.

- Requirements on page 209
- Overview on page 209
- Configuration on page 209
- Verification on page 210

Requirements

Before you begin, determine the type of parameter used to control the MGCP call activity and manage its network resources. See "Understanding MGCP ALG Call Duration and Timeouts" on page 208.

Overview

The **maximum-call-duration** parameter governs MGCP call activity and sets the absolute maximum length of a call. When a call exceeds this parameter setting, the MGCP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, and the range is 3 through 720 minutes. This setting also frees up bandwidth in cases where calls fail to properly terminate. In this example, the call duration is set to 600 minutes.

Configuration

GUI Step-by-Step Procedure

To set call duration for the MGCP ALG:

1. Select **Configure** >**Security** >**ALG**.
2. Select the **MGCP** tab.

Copyright © 2016, Juniper Networks, Inc.

3. In the Maximum call duration box, enter **600**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options** >**Commit**.

Step-by-Step Procedure

To set call duration for the MGCP ALG:

1. Configure the MGCP ALG call duration.

   ```
   [edit]
   user@host# set security alg mgcp maximum-call-duration 600
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg mgcp** command.

Related Documentation

- Understanding MGCP ALG Call Duration and Timeouts on page 208
- MGCP ALG Configuration Overview on page 189

Example: Setting MGCP ALG Inactive Media Timeout

This example shows how to set the inactive media timeout value for the MGCP ALG.

- Requirements on page 210
- Overview on page 210
- Configuration on page 211
- Verification on page 211

Requirements

Before you begin, determine the type of parameter used to control the MGCP call activity and manage its network resources. See "Understanding MGCP ALG Call Duration and Timeouts" on page 208.

Overview

The **inactive-media-timeout** parameter governs MGCP call activity and indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the MGCP ALG gates opened for media are closed. The default setting is 120 seconds, and the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated. In this example, the inactive media timeout is set to 90 seconds.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To set the inactive media timeout for the MGCP ALG:

1. Select **Configure>Security>ALG**.
2. Select the **MGCP** tab.
3. In the Inactive Media Timeout box, enter **90**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To set the inactive media timeout for the MGCP ALG:

1. Configure the MGCP ALG inactive media timeout value.

   ```
   [edit]
   user@host# set security alg mgcp inactive-media-timeout 90
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg mgcp** command.

Related Documentation

- Understanding MGCP ALG Call Duration and Timeouts on page 208
- MGCP ALG Configuration Overview on page 189

Example: Setting MGCP ALG Transaction Timeout

This example shows how to set the transaction timeout for the MGCP ALG.

- Requirements on page 211
- Overview on page 211
- Configuration on page 212
- Verification on page 212

Requirements

Before you begin, determine the type of parameter used to control the MGCP call activity and manage its network resources. See "Understanding MGCP ALG Call Duration and Timeouts" on page 208.

Overview

The **transaction-timeout** parameter governs MGCP call activity and is a signaling message; for example, a NTFY from the gateway to the call agent or a 200 OK from the call agent

Copyright © 2016, Juniper Networks, Inc.

to the gateway. The Juniper Networks device tracks these transactions, and clears them when they time out. The timeout range for MGCP transactions is from 3 to 50 seconds, and the default is 30 seconds. In this example, the transaction timeout is set to 20 seconds.

Configuration

GUI Step-by-Step Procedure

To set the transaction timeout for the MGCP ALG:

1. Select **Configure>Security>ALG**.
2. Select the **MGCP** tab.
3. In the Transaction Timeout box, enter **20**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To set the transaction timeout for the MGCP ALG:

1. Configure the MGCP ALG transaction timeout value.

   ```
   [edit]
   user@host# set security alg mgcp transaction-timeout 20
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg mgcp** command.

Related Documentation

- Understanding MGCP ALG Call Duration and Timeouts on page 208
- MGCP ALG Configuration Overview on page 189

Understanding MGCP ALG DoS Attack Protection

You can protect the Media Gateway Control Protocol (MGCP) media gateway from denial-of-service (DoS) flood attacks by limiting the number of remote access service (RAS) messages and connections per second it will attempt to process.

When you configure MGCP message flood protection, the MGCP Application Layer Gateway (ALG) drops any messages exceeding the threshold you set. The range is 2 to 50,000 messages per second per media gateway, and the default is 1000 messages per second per media gateway.

When you configure MGCP connection flood protection, the MGCP ALG drops any connection request exceeding the threshold you set. This limits the rate of processing of **CreateConnection** (**CRCX**) commands, thereby indirectly limiting pinhole creation. The

Copyright © 2016, Juniper Networks, Inc.

range is 2 to 10,000 connection requests per second per media gateway, the default is 200.

Related Documentation

- Understanding the MGCP ALG on page 183
- MGCP ALG Configuration Overview on page 189
- Example: Configuring MGCP ALG DoS Attack Protection on page 213

Example: Configuring MGCP ALG DoS Attack Protection

This example shows how to configure connection flood protection for the MGCP ALG.

- Requirements on page 213
- Overview on page 213
- Configuration on page 213
- Verification on page 214

Requirements

Before you begin, determine whether to protect the MGCP media gateway from DoS flood attacks. See "Understanding MGCP ALG DoS Attack Protection" on page 212.

Overview

In this example, you configure the MGCP ALG to drop any message requests exceeding 10,000 requests per second and to drop any connection requests exceeding 4000 per second.

Configuration

GUI Step-by-Step Procedure

To configure connection flood protection for the MGCP ALG:

1. Select **Configure>Security>ALG**.
2. Select the **MGCP** tab.
3. In the Message flood gatekeeper threshold box, type **10000**.
4. In the Connection flood threshold box, type **4000**.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure connection flood protection for the MGCP ALG:

1. Configure the connection flood threshold value.

```
[edit]
user@host# set security alg mgcp application-screen message-flood threshold
  10000
user@host# set security alg mgcp application-screen connection-flood threshold
  4000
```

Copyright © 2016, Juniper Networks, Inc.

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

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

Verification

To verify the configuration is working properly, enter the **show security alg mgcp** command.

Related Documentation

- Understanding MGCP ALG DoS Attack Protection on page 212
- MGCP ALG Configuration Overview on page 189

Understanding MGCP ALG Unknown Message Types

To accommodate on-going development of the Media Gateway Control Protocol (MGCP), you might want to allow traffic containing new MGCP message types. The unknown MGCP message type feature enables you to configure the Juniper Networks device to accept MGCP traffic containing unknown message types in both Network Address Translation (NAT) mode and route mode.

This feature enables you to specify how unidentified MGCP messages are handled by the Juniper Networks device. The default is to drop unknown (unsupported) messages. Unknown messages can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown MGCP messages can help you get your network operational so that you can later analyze your voice-over-IP (VoIP) traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

Related Documentation

- Understanding the MGCP ALG on page 183
- MGCP ALG Configuration Overview on page 189
- Example: Allowing Unknown MGCP ALG Message Types on page 214

Example: Allowing Unknown MGCP ALG Message Types

This example shows how to configure the MGCP ALG to allow unknown MGCP message types in both NAT mode and route mode.

- Requirements on page 215
- Overview on page 215
- Configuration on page 215
- Verification on page 215

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, determine whether to accommodate new and unknown MGCP message types for the device. See "Understanding MGCP ALG Unknown Message Types" on page 214.

Overview

This feature enables you to specify how unidentified MGCP messages are handled by a Juniper Networks device. The default is to drop unknown (unsupported) messages, because unknown messages can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment.

Configuration

GUI Step-by-Step Procedure

To configure the MGCP ALG to allow unknown message types:

1. Select **Configure>Security>ALG**.
2. Select the **MGCP** tab.
3. Select the **Enable Permit NAT applied** check box.
4. Select the **Enable Permit routed** check box.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure the MGCP ALG to allow unknown message types:

1. Allow unknown message types to pass if the session is in either NAT mode or in route mode.

```
[edit]
user@host# set security alg mgcp application-screen unknown-message
  permit-nat-applied permit-routed
```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg mgcp** command.

Related Documentation

- Understanding MGCP ALG Unknown Message Types on page 214
- MGCP ALG Configuration Overview on page 189

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 17

Configuring the SCCP ALG

- Understanding SCCP ALGs on page 217
- SCCP ALG Configuration Overview on page 223
- Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone on page 223
- Understanding SCCP ALG Inactive Media Timeouts on page 233
- Example: Setting SCCP ALG Inactive Media Timeouts on page 233
- Understanding SCCP ALG Unknown Message Types on page 234
- Example: Allowing Unknown SCCP ALG Message Types on page 235
- Understanding SCCP ALG DoS Attack Protection on page 236
- Example: Configuring SCCP ALG DoS Attack Protection on page 236
- Verifying SCCP ALG Configurations on page 237

Understanding SCCP ALGs

Skinny Client Control Protocol (SCCP) is a Cisco proprietary protocol for call signaling. Skinny is based on a call-agent-based call-control architecture. The control protocol uses binary-coded frames encoded on TCP frames sent to well-known TCP port number destinations to set up and tear down RTP media sessions.

The SCCP protocol, in the same way as other call control protocols, negotiates media endpoint parameters—specifically the Real-Time Transport Protocol (RTP) port number and the IP address of media termination—by embedding information in the control packets. The SCCP Application Layer Gateway (ALG) parses these control packets and facilitates media and control packets to flow through the system.

The SCCP ALG also implements rate limiting of calls and helps protect critical resources from overloading and denial-of-service (DoS) attacks.

The following functions are implemented by the SCCP ALG in Junos OS:

- Validation of SCCP protocol data units
- Translation of embedded IP address and port numbers
- Allocation of firewall resources (pinholes and gates) to pass media
- Aging out idle calls

Copyright © 2016, Juniper Networks, Inc.

- Configuration API for SCCP ALG parameters
- Operational mode API for displaying counters, status and statistics

In the SCCP architecture, a proxy, known as the Call Manager, does most of the processing. IP phones, also called End Stations, run the SCCP client and connect to a primary (and, if available, a secondary) Call Manager over TCP on port 2000 and register with the primary Call Manager. This connection is then used to establish calls coming to or from the client.

The SCCP ALG supports the following:

- Call flow from a SCCP client, through the Call Manager, to another SCCP client.
- Seamless failover—Switches over all calls in process to the standby firewall during failure of the primary.
- Voice-over-IP (VoIP) signaling payload inspection—Fully inspects the payload of incoming VoIP signaling packets. Any malformed packet attack is blocked by the ALG.
- SCCP signaling payload inspection—Fully inspects the payload of incoming SCCP signaling packets. Any malformed-packet attack is blocked by the ALG.
- Stateful processing—Invokes the corresponding VoIP-based state machines to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- Network Address Translation (NAT)—Translates any embedded IP address and port information in the payload, based on the existing routing information and network topology, with the translated IP address and port number, if necessary.
- Pinhole creation and management for VoIP traffic—Identifies IP address and port information used for media or signaling and dynamically opens (and closes) pinholes to securely stream the media.

This topic includes the following sections:

- SCCP Security on page 218
- SCCP Components on page 219
- SCCP Transactions on page 220
- SCCP Version on page 221
- SCCP Control Messages and RTP Flow on page 221
- SCCP Messages on page 221

SCCP Security

The SCCP ALG includes the following security features:

- Stateful inspection of SCCP control messages over TCP and validation of the message format, and message validity for the current call state. Invalid messages are dropped.
- Security policy enforcement between Cisco IP phones and Cisco Call Manager.

Copyright © 2016, Juniper Networks, Inc.

- Protect against call flooding by rate limiting the number of calls processed by the ALG.
- Seamless failover of calls, including the ones in progress in case of device failure in a clustered deployment.

SCCP Components

The principal components of the SCCP VoIP architecture include the following:

- SCCP Client on page 219
- Call Manager on page 219
- Cluster on page 219

SCCP Client

The SCCP client runs on an IP phone, also called an *End Station*, which uses SCCP for signaling and for making calls. For an SCCP client to make a call, it must first register with a Primary Call Manager (and a secondary, if available). The connection between the client and the Call Manager is over TCP on port 2000. This connection is then used to establish calls to or from the client. Transmission of media is over RTP, UDP, and IP.

Call Manager

The Call Manager implements SCCP call control server software and has overall control of all devices and communication in the SCCP VoIP network. Its functions include defining, monitoring and controlling SCCP groups, regions of numbers, and route plans; providing initialization, admission, and registration of devices on the network; providing a redundant database that contains addresses, phone numbers, and number formats; and initiating contact with called devices or their agents to establish logical sessions in which voice communication can flow.

Cluster

A *cluster* is a collection of SCCP clients and a Call Manager. The Call Manager in the cluster detects all SCCP clients in the cluster. There can be more than one Call Manager for backup in a cluster. Call Manager behavior varies in each of the following cluster scenarios:

- Intra-Cluster, in which the Call Manager detects each SCCP client, and the call is between SCCP clients of the same cluster.
- Inter-Cluster, in which the Call Manager needs to communicate with another Call Manager using H.323 for call setup.
- Inter-Cluster calls using the gatekeeper for admission control and address resolution.

  Call Manager behavior also varies with calls between an SCCP client and a phone in a public switched telephone network (PSTN), and with calls between an SCCP client and a phone in another administrative domain that is using H.323.

Copyright © 2016, Juniper Networks, Inc.

SCCP Transactions

SCCP transactions are the processes that need to take place in order for an SCCP call to proceed. SCCP transactions include the following processes:

- Client Initialization on page 220
- Client Registration on page 220
- Call Setup on page 220
- Media Setup on page 220

Client Initialization

To initialize, the SCCP client needs to determine the IP address of the Call Manager, its own IP address, and other information about the IP gateway and DNS servers. Initialization takes place on the local LAN. The client sends a Dynamic Host Control Protocol (DHCP) request to get an IP address, the DNS server address, and the TFTP server name and address. The client needs the TFTP server name to download the configuration file called sep*macaddr*.cnf. If the TFTP name is not given, the client uses the default filename in the IP phone. The client then downloads the .cnf (xml) configuration file from TFTP server. CNF files contain the IP address or addresses of the primary and secondary Cisco Call Manager. With this information, the client contacts the Call Manager to register.

Client Registration

The SCCP client, after initialization, registers with the Call Manager over a TCP connection on well-known default port 2000. The client registers by providing the Call Manager with its IP address, the MAC address of the phone, and other information, such as protocol and version. The client cannot initiate or receive calls until it is registered. Keepalive messages keep this TCP connection open between the client and Call Manager so that the client can initiate or receive calls at any time, provided that a policy on the device allows this.

Call Setup

IP phone-to-IP phone call setup using SCCP is always handled by the Call Manager. Messages for call setup are sent to the Call Manager, which returns messages appropriate to the status of the call. If call setup is successful, and a policy on the device allows the call, the Call Manager sends the media setup messages to the client.

Media Setup

The Call Manager sends the IP address and port number of the called party to the calling party. The Call Manager also sends the media IP address and port number of the calling party to the called party. After media setup, media is transmitted directly between clients. When the call ends, the Call Manager is informed and terminates the media streams. At no time during this process does the Call Manager hand over call-setup function to the client. Media is streamed directly between clients through RTP/UDP/IP.

Copyright © 2016, Juniper Networks, Inc.

SCCP Version

Starting in Junos OS Release 12.1X46-D10, the SCCP ALG supports SCCP versions 16, 17, and 20 and several SCCP messages have been updated with a new format. Cisco Call Manager (CM) version 7 uses SCCP version 20.

SCCP Control Messages and RTP Flow

Figure 17 shows the SCCP control messages used to set up and tear down a simple call between Phone 1 and Phone 2. Except for the OffHook message initiating the call from Phone1 and the OnHook message signaling the end of the call, all aspects of the call are controlled by the Call Manager.

Figure 17: Call Setup and Teardown

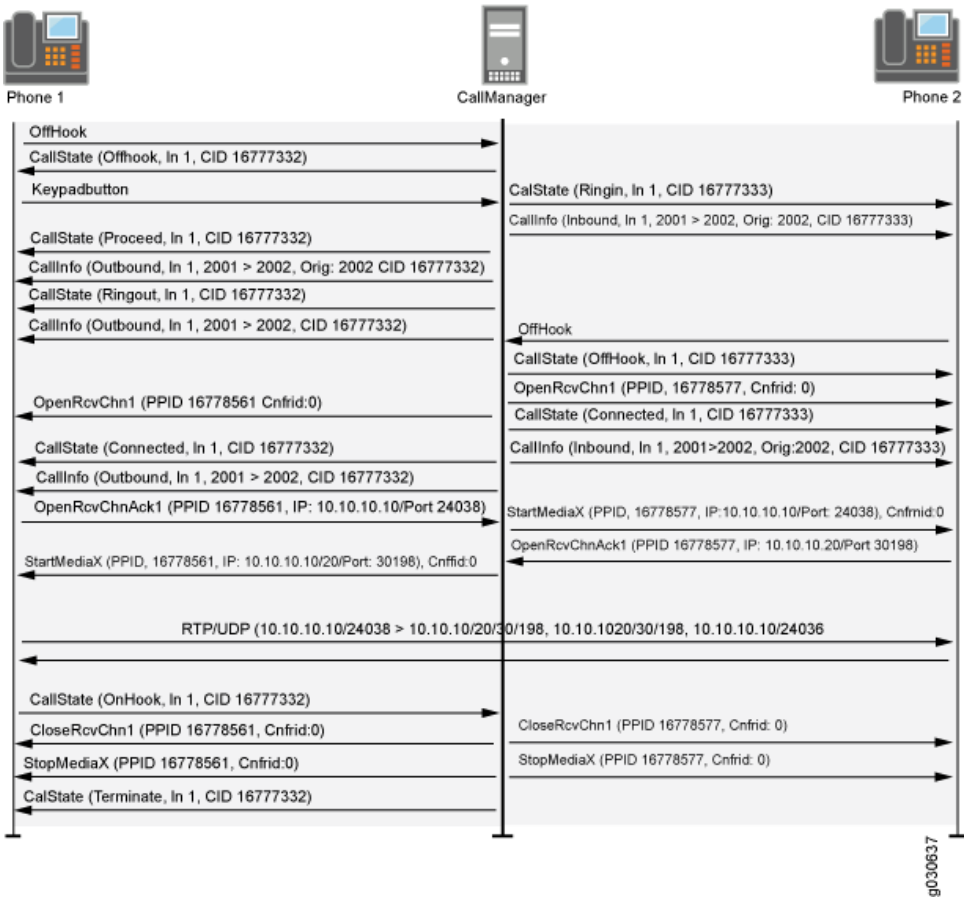

SCCP Messages

Table 8, Table 9, Table 10, and Table 11 list the SCCP call message IDs in the four intervals allowed by the device.

Table 8: Station to Call Manager Messages

#define STATION_REGISTER_MESSAGE	0x00000001

Copyright © 2016, Juniper Networks, Inc.

Table 8: Station to Call Manager Messages *(continued)*

#define STATION_IP_PORT_MESSAGE	0x00000002
#define STATION_ALARM_MESSAGE	0x00000020
#define STATION_OPEN_RECEIVE_CHANNEL_ACK	0x00000022

Table 9: Call Manager to Station Messages

#define STATION_START_MEDIA_TRANSMISSION	0x00000001
#define STATION_STOP_MEDIA_TRANSMISSION	0x00000002
#define STATION_CALL_INFO_MESSAGE	0x00000020
#define STATION_OPEN_RECEIVE_CHANNEL_ACK	0x00000022
#define STATION_CLOSE_RECEIVE_CHANNEL	0x00000106

Table 10: Call Manager 4.0 Messages and Post Sccp 6.2

#define STATION_REGISTER_TOKEN_REQ_MESSAGE	0x00000029
#define STATION_MEDIA_TRANSMISSION_FAILURE	0x0000002A
#define STATION_OPEN_MULTIMEDIA_RECEIVE_CHANNEL_ACK	0x00000031

Table 11: Call Manager to Station

#define STATION_OPEN_MULTIMEDIA_RECEIVE_CHANNEL	0x00000131
#define STATION_START_MULTIMEDIA_TRANSMISSION	0x00000132
#define STATION_STOP_MULTIMEDIA_TRANSMISSION	0x00000133
#define STATION_CLOSE_MULTIMEDIA_RECEIVE_CHANNEL	0x00000136

- The limitations for SCCP ALGs are as follows:
  - The SCCP is a Cisco proprietary protocol. So, any changes to the protocol by Cisco cause the SCCP ALG implementation to break. However, workarounds are provided to bypass strict decoding and allow any protocol changes to be handled gracefully.
  - Any changes to the policies will drop the sessions and impact already established SCCP calls.

Copyright © 2016, Juniper Networks, Inc.

- The SCCP ALG opens pinholes that are collapsed during traffic or media inactivity. This means that during a temporary loss of connectivity, media sessions are not reestablished.
- Call Manager (CM) version 6.x and later does not support TCP probe packets in chassis cluster mode. As a result, the existing SCCP sessions will break when there is a failover. You can still create new SCCP sessions during failover.

Related Documentation

- ALG Overview on page 3
- SCCP ALG Configuration Overview on page 223
- Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone on page 223

SCCP ALG Configuration Overview

The Skinny Client Control Protocol Application Layer Gateway (SCCP ALG) is enabled by default on the device—no action is required to enable it. However, you might choose to fine-tune SCCP ALG operations by using the following instructions:

1. Conserve network resources and maximize throughput. For instructions, see "Example: Setting SCCP ALG Inactive Media Timeouts" on page 233.
2. Enable unknown messages to pass when the session is in Network Address Translation (NAT) mode and route mode. For instructions, see "Example: Allowing Unknown SCCP ALG Message Types" on page 235.
3. Protect the SCCP clients from denial-of-service (DoS) flood attacks. For instructions, see "Example: Configuring SCCP ALG DoS Attack Protection" on page 236.

Related Documentation

- Understanding SCCP ALGs on page 217
- Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone on page 223
- Verifying SCCP ALG Configurations on page 237

Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone

This example shows how to configure static NAT on the outgoing interface of a Juniper Networks device to allow callers in a public zone to register with an SCCP ALG Call Manager or a TFTP server located in a private zone.

- Requirements on page 224
- Overview on page 224
- Configuration on page 225
- Verification on page 230

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, understand NAT support with SCCP ALG. See "Understanding SCCP ALGs" on page 217.

Overview

In this example (see Figure 18), a single device is serving as a Call Manager or a TFTP server. The Call Manager or TFTP server and phone1 are attached to the private zone, and phone2 is attached to the public zone. You configure a static NAT rule set for the Call Manager or TFTP server so that when phone2 boots up it contacts the TFTP server and obtains the IP address of the Call Manager. You then create a policy called in-pol to allow SCCP traffic from the public to the private zone and a policy called out-pol to allow phone1 to call out.

NOTE: We recommend that you change the IP address of the Call Manager, which resides in the TFTP server configuration file (**sep <mac_addr>.cnf**), to the NAT IP address of the Call Manager.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 18 shows call manager or TFTP server in the private zone.

Figure 18: Call Manager or TFTP Server in the Private Zone

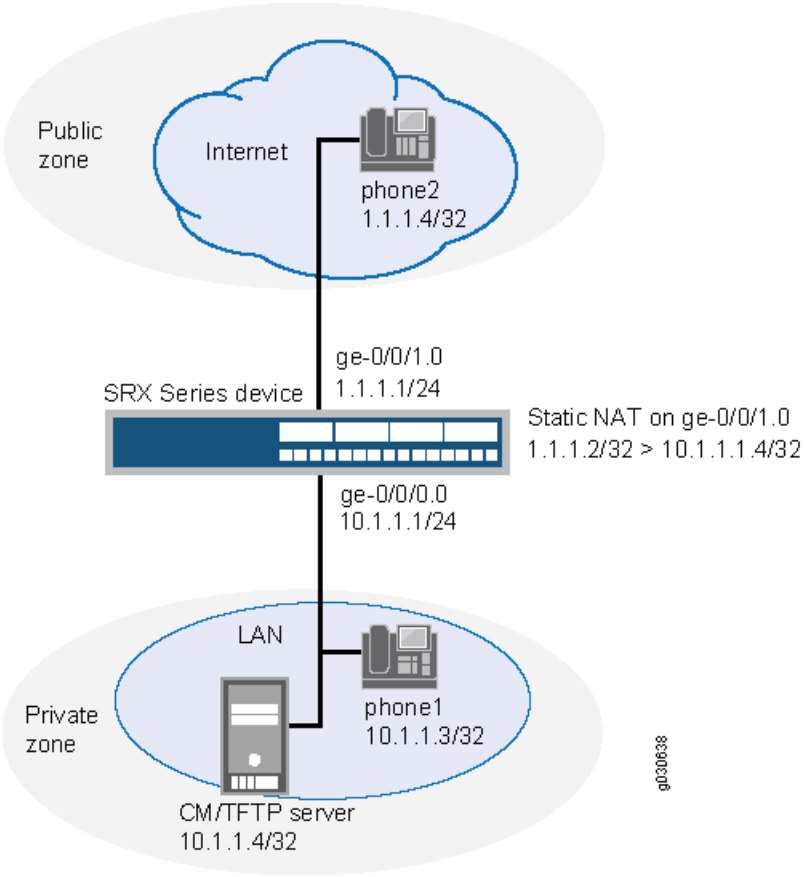

In this example, you configure NAT as follows:

- Create a static NAT rule set called to-proxy with a rule called phone2 to match packets from the public zone with the destination address 1.1.1.2/32. For matching packets, the destination IP address is translated to the private address 10.1.1.4/32.
- Configure proxy ARP for the address 1.1.1.2/32 on interface ge-0/0/1.0. This allows the system to respond to ARP requests received on the interface for these addresses.
- Configure a second rule set called phones with a rule called phone1 to enable interface NAT for communication from phone1 to phone2.

Configuration

CLI Quick Configuration

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

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public interfaces ge-0/0/1.0
set security address-book book1 address phone1 10.1.1.3/32
set security address-book book1 address cm-tftp_server 10.1.1.4/32
set security address-book book1 attach zone private
set security address-book book2 address phone2 1.1.1.4/32
set security address-book book2 attach zone public
set security nat source rule-set phones from zone private
set security nat source rule-set phones to zone public
set security nat source rule-set phones rule phone1 match source-address 10.1.1.3/32
set security nat source rule-set phones rule phone1 then source-nat interface
set security nat static rule-set to-proxy from zone public
set security nat static rule-set to-proxy rule phone2 match destination-address 1.1.1.2/32
set security nat static rule-set to-proxy rule phone2 then static-nat prefix 10.1.1.4/32
set security nat proxy-arp interface ge-0/0/1.0 address 1.1.1.2/32
set security policies from-zone public to-zone private policy in-pol match source-address
  phone2
set security policies from-zone public to-zone private policy in-pol match
  destination-address cm-tftp_server
set security policies from-zone public to-zone private policy in-pol match
  destination-address phone1
set security policies from-zone public to-zone private policy in-pol match application
  junos-sccp
set security policies from-zone public to-zone private policy in-pol then permit
set security policies from-zone private to-zone public policy out-pol match source-address
  any
set security policies from-zone private to-zone public policy out-pol match
  destination-address phone2
set security policies from-zone private to-zone public policy out-pol match application
  junos-sccp
set security policies from-zone private to-zone public policy out-pol then permit
set security policies from-zone private to-zone private policy tftp-pol match
  source-address any
set security policies from-zone private to-zone private policy tftp-pol match
  destination-address any
set security policies from-zone private to-zone private policy tftp-pol match application
  junos-tftp
set security policies from-zone private to-zone private policy tftp-pol then permit
```

Step-by-Step Procedure

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

To configure NAT for an SCCP ALG Call Manager or a TFTP server located in a private zone:

1. Configure interfaces.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
   user@host# set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
   ```

2. Create security zones.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security zones]
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone public interfaces ge-0/0/1.0
```

3. Create address books and attach zones to them.

```
[edit security address-book book1]
user@host# set address phone1 10.1.1.3/32
user@host# set address cm-tftp_server 10.1.1.4/32
user@host# set attach zone private

[edits security address-book book2]
user@host# set address phone2 1.1.1.4/32
user@host# set attach zone public
```

4. Create a rule set for static NAT and assign a rule to it.

```
[edit security nat static rule-set to-proxy]
user@host# set from zone public
user@host# set rule phone2 match destination-address 1.1.1.2/32
user@host# set rule phone2 then static-nat prefix 10.1.1.4/32
```

5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/1.0 address 1.1.1.2/32
```

6. Configure interface NAT for communication from phone1 to phone2.

```
[edit security nat source rule-set phones]
user@host# set from zone private
user@host# set to zone public
user@host# set rule phone1 match source-address 10.1.1.3/32
user@host# set rule phone1 then source-nat interface
```

7. Configure a policy to allow traffic from the public zone to the private zone.

```
[edit security policies from-zone public to-zone private policy in-pol]
user@host# set match source-address phone2
user@host# set match destination-address cm-tftp_server
user@host# set match destination-address phone1
user@host# set match application junos-sccp
user@host# set then permit
```

8. Configure a policy to allow traffic from the private zone to the public zone.

```
[edit security policies from-zone private to-zone public policy out-pol]
user@host# set match source-address any
user@host# set match destination-address phone2
user@host# set match application junos-sccp
user@host# set then permit
```

9. Configure a policy to allow traffic from phone1 to the CM/TFTP server.

```
[edit security policies from-zone private to-zone private policy tftp-pol]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application junos-tftp
user@host# set then permit
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security address-book**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
    ge-0/0/1 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }
[edit]
user@host# show security zones
security-zone private {
  interfaces {
    ge-0/0/0.0;
  }
}
  security-zone public {
    interfaces {
      ge-0/0/1.0;
    }
  }
[edit]
user@host# show security address-book
book1 {
  address phone1 10.1.1.3/32;
  address cm-tftp_server 10.1.1.4/32;
  attach {
    zone private;
  }
}
  book2 {
    address phone2 1.1.1.4/32;
    attach {
      zone public;
    }
  }
[edit]
user@host# show security nat
source {
  rule-set phones {
    from zone private;
    to zone public;
    rule phone1 {
```


Copyright © 2016, Juniper Networks, Inc.

```
        match {
          source-address 10.1.1.3/32;
        }
        then {
          source-nat {
            interface;
          }
        }
      }
    }
  }
    static {
      rule-set to-proxy {
        from zone public;
        rule phone2 {
          match {
            destination-address 1.1.1.2/32;
          }
          then {
            static-nat prefix 10.1.1.4/32;
          }
        }
      }
    }
    proxy-arp {
      interface ge-0/0/1.0 {
        address {
          1.1.1.2/32;
        }
      }
    }
[edit]
user@host# show security policies
from-zone public to-zone private {
  policy in-pol {
    match {
      source-address phone2;
      destination-address cm-tftp_server;
      destination-address phone1;
      application junos-sccp;
    }
    then {
      permit;
    }
  }
}
  from-zone private to-zone public {
    policy out-pol {
      match {
        source-address any;
        destination-address phone2;
        application junos-sccp;
      }
      then {
        permit;
      }
```

```
    }
  }
  from-zone private to-zone private {
    policy tftp-pol {
      match {
        source-address any;
        destination-address any;
        application junos-tftp;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Rule Usage on page 230
- Verifying Static NAT Configuration on page 231
- Verifying Static NAT Configuration on page 231
- Verifying SCCP ALG on page 232
- Verifying the Security Polices of SIP ALG on page 232

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command.

```
user@host> show security nat source rule all

Total rules: 2
Total referenced IPv4/IPv6 ip-prefixes: 3/0
source NAT rule: r1                   Rule-set: rs1
  Rule-Id                   : 1
  Rule position             : 1
  From interface            : ge-0/0/5.0
                            : ge-0/0/6.0
  To interface              : ge-0/0/2.0
  Match
    Source addresses        : 30.0.0.0        - 30.0.0.255
                              40.0.0.0        - 40.0.0.255
    Destination port        : 0               - 0
  Action                      : interface
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
    Max session number        : 0
  Translation hits          : 0
source NAT rule: phone1               Rule-set: phones
  Rule-Id                   : 3
  Rule position             : 2
```


Copyright © 2016, Juniper Networks, Inc.

```
  From zone                 : private
  To zone                   : public
  Match
    Source addresses        : 10.1.1.3        - 10.1.1.3
    Destination port        : 0               - 0
  Action                       : interface
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
    Max session number        : 0
  Translation hits          : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule.

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule phone1** command.

```
user@host> show security nat static rule phone1
```

```
source NAT rule: phone1               Rule-set: phones
  Rule-Id                 : 3
  Rule position           : 2
  From zone               : private
  To zone                 : public
  Match
    Source addresses      : 10.1.1.3        - 10.1.1.3
    Destination port      : 0               - 0
  Action                       : interface
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
    Max session number        : 0
  Translation hits          : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule set.

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule phone2** command.

```
user@host> show security nat static rule phone2
```

```
Static NAT rule: phone2               Rule-set: to-proxy
  Rule-Id                 : 1
  Rule position           : 1
  From zone               : public
  Destination addresses   : 1.1.1.2
  Host addresses          : 10.1.1.4
  Netmask                 : 32
  Host routing-instance   : N/A
  Translation hits        : 0
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule set.

Verifying SCCP ALG

Purpose Verify that the SCCP ALG is enabled.

Action From operational mode, enter the **show security alg status | match sccp** command.

```
user@host> show security alg status | match sccp
SCCP      : Enabled
```

Meaning The output shows the SCCP ALG status as follows:

- Enabled—Shows the SCCP ALG is enabled.
- Disabled—Shows the SCCP ALG is disabled.

Verifying the Security Polices of SIP ALG

Purpose Verify that the static NAT between public zone and private zone is set.

Action From operational mode, enter the **show security policies** command.

```
user@host> show security policies
from-zone private to-zone public {
       policy out-pol {
        match {
            source-address any;
            destination-address phone2;
            application junos-sccp;
        }
        then {
            permit;
        }
    }
}
from-zone public to-zone private {
       policy in-pol {
        match {
            source-address phone2;
            destination-address [ cm-tftp_server phone1 ];
            application junos-sccp;
        }
        then {
            permit;
        }
    }
}
from-zone private to-zone private {
    policy tftp-pol {
        match {
            source-address any;
            destination-address any;
            application junos-tftp;
        }
        then {
```


Copyright © 2016, Juniper Networks, Inc.

```
                permit;
            }
        }
    }
```

Meaning The sample output shows that the static NAT between public zone and private zone is set.

Related Documentation
- SCCP ALG Configuration Overview on page 223

Understanding SCCP ALG Inactive Media Timeouts

The inactive media timeout feature helps you to conserve network resources and maximize throughput.

This parameter indicates the maximum length of time (in seconds) a call can remain active without any media traffic within a group. Each time a Real-Time Transport Protocol (RTP) or Real-Time Control Protocol (RTCP) packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the gates the Skinny Client Control Protocol (SCCP) opened for media are closed. The default setting is 120 seconds, and the range is from 10 to 2550 seconds. Note that upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

Related Documentation
- Understanding SCCP ALGs on page 217
- SCCP ALG Configuration Overview on page 223
- Example: Setting SCCP ALG Inactive Media Timeouts on page 233

Example: Setting SCCP ALG Inactive Media Timeouts

This example shows how to set the inactive media timeout value for the SCCP ALG.

- Requirements on page 233
- Overview on page 233
- Configuration on page 234
- Verification on page 234

Requirements

Before you begin, review the parameter used to indicate the maximum length of time (in seconds) a call can remain active without any media traffic within a group. See "Understanding SCCP ALG Inactive Media Timeouts" on page 233.

Overview

Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the gates the SCCP opened for media are closed. This example sets the media inactivity timeout to 90 seconds.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To set the inactive media timeout for the SCCP ALG:

1. Select **Configure>Security>ALG**.
2. Select the **SCCP** tab.
3. In the Inactive Media Timeout box, enter **90**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To set the inactive media timeout for the SCCP ALG:

1. Configure the SCCP ALG inactive media timeout value.

   ```
   [edit]
   user@host# set security alg sccp inactive-media-timeout 90
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sccp** command.

Related Documentation

- Understanding SCCP ALG Inactive Media Timeouts on page 233
- SCCP ALG Configuration Overview on page 223
- Verifying SCCP ALG Configurations on page 237

Understanding SCCP ALG Unknown Message Types

To accommodate on-going development of the Skinny Client Control Protocol (SCCP), you might want to allow traffic containing new SCCP message types. The unknown SCCP message type feature enables you to configure the device to accept SCCP traffic containing unknown message types in both Network Address Translation (NAT) mode and route mode.

This feature enables you to specify how unidentified SCCP messages are handled by the device. The default is to drop unknown (unsupported) messages. We do not recommend permitting unknown messages because they can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown SCCP messages can help you get your network operational so that you can later analyze your voice-over-IP (VoIP) traffic to determine why some messages were being dropped.

Note that this command applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as

Copyright © 2016, Juniper Networks, Inc.

a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

Related Documentation

- Understanding SCCP ALGs on page 217
- SCCP ALG Configuration Overview on page 223
- Example: Allowing Unknown SCCP ALG Message Types on page 235

Example: Allowing Unknown SCCP ALG Message Types

This example shows how to configure the SCCP ALG to allow unknown SCCP message types in both NAT mode and route mode.

- Requirements on page 235
- Overview on page 235
- Configuration on page 235
- Verification on page 236

Requirements

Before you begin, determine whether to accommodate new and unknown SCCP message types for the device. See "Understanding SCCP ALG Unknown Message Types" on page 234.

Overview

This feature enables you to specify how unidentified SCCP messages are handled by a Juniper Networks device. The default is to drop unknown (unsupported) messages because unknown messages can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment.

Configuration

GUI Step-by-Step Procedure

To configure the SCCP ALG to allow unknown message types:

1. Select **Configure>Security>ALG**.
2. Select the **SCCP** tab.
3. Select the **Enable Permit NAT applied** check box.
4. Select the **Enable Permit routed** check box.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure the SCCP ALG to allow unknown message types:

1. Allow unknown message types to pass if the session is in either NAT mode or in route mode.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host#  set security alg sccp application-screen unknown-message
  permit-nat-applied permit-routed
```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sccp** command.

Related Documentation

- Understanding SCCP ALG Unknown Message Types on page 234
- SCCP ALG Configuration Overview on page 223
- Verifying SCCP ALG Configurations on page 237

Understanding SCCP ALG DoS Attack Protection

You can protect Skinny Client Control Protocol Application Layer Gateway (SCCP ALG) clients from denial-of-service (DoS) flood attacks by limiting the number of calls they attempt to process.

When you configure SCCP call flood protection, the SCCP ALG drops any calls exceeding the threshold you set. The range is 2 to 1000 calls per second per client, the default is 20.

Related Documentation

- Understanding SCCP ALGs on page 217
- SCCP ALG Configuration Overview on page 223
- Example: Configuring SCCP ALG DoS Attack Protection on page 236

Example: Configuring SCCP ALG DoS Attack Protection

This example shows how to configure connection flood protection for the SCCP ALG.

- Requirements on page 236
- Overview on page 236
- Configuration on page 237
- Verification on page 237

Requirements

Before you begin, determine whether to protect the SCCP media gateway from DoS flood attacks. See “Understanding SCCP ALG DoS Attack Protection” on page 236.

Overview

In this example, the device is configured to drop any calls exceeding 500 per second per client.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To configure call flood protection for the SCCP ALG:

1. Select **Configure>Security>ALG**.
2. Select the **SCCP** tab.
3. In the Call flood threshold box, type **500**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure call flood protection for the SCCP ALG:

1. Configure the DoS attack protection:

   ```
   [edit]
   user@host# set security alg sccp application-screen call-flood threshold 500
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sccp** command.

Related Documentation

- Understanding SCCP ALG DoS Attack Protection on page 236
- SCCP ALG Configuration Overview on page 223
- Verifying SCCP ALG Configurations on page 237

Verifying SCCP ALG Configurations

- Verifying SCCP ALG on page 237
- Verifying SCCP ALG Calls on page 238
- Verifying SCCP ALG Call Details on page 238
- Verifying SCCP ALG Counters on page 239

Verifying SCCP ALG

Purpose Display SCCP verification options.

Action From the CLI, enter the **show security alg sccp** command.

```
user@host> show security alg sccp ?
Possible completions:
  calls                Show SCCP calls
  counters             Show SCCP counters
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The output shows a list of all SCCP verification parameters. Verify the following information:

- All SCCP calls
- Counters for all SCCP calls

Verifying SCCP ALG Calls

Purpose Display a list of all SCCP calls

Action From the CLI, enter the **show security alg sccp calls** command.

```
user@host> show security alg sccp calls
Possible completions:
  calls                Show SCCP calls
  counters             Show SCCP counters
  endpoints            Show SCCP endpoints
```

Meaning The output shows a list of all SCCP verification parameters. Verify the following information:

- All SCCP calls
- Counters for all SCCP c alls
- Information about all SCCP endpoints

Verifying SCCP ALG Call Details

Purpose Display details about all SCCP calls.

Action From the CLI, enter the **show security alg sccp calls detail** command.

```
user@host> show security alg sccp calls detail
Client IP address: 11.0.102.91
  Client zone: 7
  Call Manager IP: 13.0.99.226
  Conference ID: 16789504
  Resource manager group: 2048
  SCCP channel information:
    Media transmit channel address (IP address/Port): 0.0.0.0:0
    Media transmit channel translated address (IP address/Port): 0.0.0.0:0
    Media transmit channel pass-through party ID (PPID): 0
    Media transmit channel resource ID: 0
    Media receive channel address (IP address/Port): 11.0.102.91:20060
    Media receive channel translated address (IP address/Port): 25.0.0.1:1032
    Media receive channel pass-through party ID (PPID): 16934451
    Media receive channel resource ID: 8185
    Multimedia transmit channel address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel pass-through party ID (PPID): 0
    Multimedia transmit channel resource ID: 0
    Multimedia receive channel address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel pass-through party ID (PPID): 0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Multimedia receive channel resource ID: 0
Total number of calls = 1
```

Meaning The output shows a list of all SCCP verification parameters. Verify the following information:

- Client zone
- Call Manager IP address: 13.0.99.226
- Conference ID
- Resource manager group
- SCCP channel information
- Total number of calls

Verifying SCCP ALG Counters

Purpose Display a list of all SCCP counters

Action From the J-Web interface, select **Monitor>ALGs>SCCP>Counters**. Alternatively, from the CLI, enter the **show security alg sccp counters** command.

```
user@host> show security alg sccp ounters
SCCP call statistics:
  Active client sessions         : 0
  Active calls                   : 0
  Total calls                    : 0
  Packets received               : 0
  PDUs processed                 : 0
  Current call rate              : 0
Error counters:
  Packets dropped                     : 0
  Decode errors                       : 0
  Protocol errors                     : 0
  Address translation errors          : 0
  Policy lookup errors                : 0
  Unknown PDUs                        : 0
  Maximum calls exceeded              : 0
  Maximum call rate exceeded          : 0
  Initialization errors               : 0
  Internal errors                     : 0
  Nonspecific error                   : 0
  No active calls to delete           : 0
  No active client sessions to delete : 0
  Session cookie create errors        : 0
  Invalid NAT cookie detected         : 0
```

Meaning The output shows a list of all SCCP verification parameters. Verify the following information:

- SCCP call statistics
- Error counters

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- SCCP ALG Configuration Overview on page 223
- Example: Configuring the SCCP ALG Call Manager or TFTP Server in the Private Zone on page 223

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 18

Configuring the SIP ALG

- Understanding the SIP ALG on page 241
- Understanding IPv6 Support for SIP ALG on page 246
- Understanding SIP ALG Request Methods on page 247
- SIP ALG Configuration Overview on page 248
- Understanding SIP ALG Call Duration and Timeouts on page 249
- Example: Setting SIP ALG Call Duration and Timeouts on page 250
- Understanding SIP ALG DoS Attack Protection on page 251
- Example: Configuring SIP ALG DoS Attack Protection on page 252
- Understanding SIP ALG Unknown Message Types on page 253
- Example: Allowing Unknown SIP ALG Message Types on page 253
- Understanding SIP ALG Hold Resources on page 254
- Retaining SIP ALG Hold Resources (CLI Procedure) on page 255
- Retaining SIP ALG Hold Resources (J-Web Procedure) on page 255
- Understanding the SIP ALG and NAT on page 256
- Understanding Incoming SIP ALG Call Support Using the SIP Registrar and NAT on page 266
- Example: Configuring Interface Source NAT for Incoming SIP Calls on page 267
- Decreasing Network Complexity by Configuring a Source NAT Pool for Incoming SIP Calls on page 273
- Example: Configuring Static NAT for Incoming SIP Calls on page 281
- Example: Configuring the SIP Proxy in the Private Zone and NAT in the Public Zone on page 288
- Example: Configuring a Three-Zone SIP ALG and NAT Scenario on page 294
- Verifying SIP ALG Configurations on page 302

Understanding the SIP ALG

Session Initiation Protocol (SIP) is an Internet Engineering Task Force (IETF)-standard protocol for initiating, modifying, and terminating multimedia sessions over the Internet.

Copyright © 2016, Juniper Networks, Inc.

Such sessions might include conferencing, telephony, or multimedia, with features such as instant messaging and application-level mobility in network environments.

Junos OS supports SIP as a service, allowing and denying it based on a policy that you configure. SIP is a predefined service in Junos OS and uses port 5060 as the destination port.

One of SIP's function is to distribute session-description information, and during the session, to negotiate and modify the parameters of the session. SIP is also used to terminate a multimedia session, signal a call establishment, provide failure indication, and provide methods for endpoint registration.

Session-description information is included in INVITE and 200-OK messages or 200-OK and ACK messages and indicates the multimedia type of the session; for example, whether it is voice or video. Although SIP can use different description protocols to describe the session, the Juniper Networks SIP Application Layer Gateway (ALG) supports only the Session Description Protocol (SDP).

SDP provides information that a system can use to join a multimedia session. SDP might include information such as IP addresses, port numbers, times, and dates. Note that the IP address and port number in the SDP header (the c= and m= fields, respectively) are the address and port where the client wants to receive the media streams and not the IP address and port number from which the SIP request originates (although they can be the same).

SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call). A user agent (UA) is an application that runs at the endpoints of the call and consists of two parts:

- user agent client (UAC), which sends SIP requests on behalf of the user
- user agent server (UAS), which listens to the responses and notifies the user when they arrive

UAC and UAS are defined in relation to the role a particular agent is playing in a negotiation.

Examples of UAs are SIP proxy servers and phones.

This topic contains the following sections:

- SIP ALG Operation on page 242
- SDP Session Descriptions on page 244
- Pinhole Creation on page 244

SIP ALG Operation

There are two types of SIP traffic, the signaling and the media stream. SIP signaling traffic consists of request and response messages between client and server and uses transport protocols such as UDP or TCP. The media stream carries the data (audio data, for example) using transport protocols.

Copyright © 2016, Juniper Networks, Inc.

TCP support over SIP ALG reduces traffic to the server by eliminating the need to reregister or refresh the server frequently.

By default, Junos OS supports SIP signaling messages on port 5060. You can configure the port by creating a policy that permits SIP service, and the software filters SIP signaling traffic like any other type of traffic, permitting or denying it. The media stream, however, uses dynamically assigned port numbers that can change several times during the course of a call. Without fixed ports, it is insecure to create a static policy to control media traffic. In this case, the device invokes the SIP ALG. The device transport ports used for the media sessions are not known in advance; however, the ports used for the SIP negotiation are well-known (or predefined). The ALG registers interest in packets from the control session, which it can easily distinguish from the other packets, and inspects the negotiation for the transport information used for the media session (both IP addresses and ports).

NOTE: The SIP ALG creates a pinhole when it determines a matching IP, port, transport address, and protocol, which are identified with whatever information is known at the time when the pinhole is opened.

The SIP ALG monitors SIP transactions and dynamically creates and manages pinholes based on the information it extracts from these transactions. The Juniper Networks SIP ALG supports all SIP methods and responses. You can allow SIP transactions to traverse the Juniper Networks firewall by creating a static policy that permits SIP service. If the policy is configured to inspect SIP traffic (or, more appropriately, if the policy sends some traffic to the SIP ALG for inspection), the allowed actions are to permit the traffic (in which case the appropriate pinholes are opened) or to deny the traffic.

The SIP ALG intercepts SIP messages that contain SDP and, using a parser, extracts the information it requires to create pinholes. The SIP ALG examines the SDP portion of the packet, and a parser extracts information such as IP addresses and port numbers, which the SIP ALG records in a pinhole table. The SIP ALG uses the IP addresses and port numbers recorded in the pinhole table to open pinholes and allow media streams to traverse the device.

NOTE: When the device is performing NAT, the transport addresses that the UAs employ are incorrect. The SIP ALG modifies the transport addresses based on the translated ports and addresses allocated by the device translating network addresses. When SDP is encrypted, the device cannot either extract or modify the contents of the message and therefore cannot correct the transport addresses. To provide a workaround, the STUN protocol has been deployed (requiring NAT devices to do some form of cone-NAT), which allows the clients to determine the translated addresses and use those newly discovered addresses in the SDP messages.

NEC SIP products are conditionally supported.

Copyright © 2016, Juniper Networks, Inc.

SDP Session Descriptions

An SDP session description is a well-defined format for conveying sufficient information to discover and participate in a multimedia session. A session is described by a series of attribute/value pairs, one per line. The attribute names are single characters, followed by =, and a value. Optional values are specified with =*. Values are either an ASCII string, or a sequence of specific types separated by spaces. Attribute names are only unique within the associated syntactic construct, such as within the session, time, or media only.

NOTE: In the SDP session description, the media-level information begins with the m= field.

Of the many fields in the SDP description, two are particularly useful to the SIP ALG because they contain Transport Layer information.

- **c=** for connection information

  This field can appear at the session or media level. It appears in this format:

  c=<network-type><address-type><connection-address>

  Currently, Junos OS supports only “IN” (for Internet) as the network type, “IPv4” and “IPv6” as the address type, and a unicast IP address or domain name as the destination (connection) IP address.

  If the destination IP address is a unicast IP address, the SIP ALG creates pinholes using the IP address and port numbers specified in the media description field m=.

- **m=** for media announcement

  This field appears at the media level and contains the description of the media. It appears in this format:

  m=<media><port><transport><fmt list>

  Currently, Junos OS supports ”RTP” as the Application Layer transport protocol. The port number indicates the destination port of the media stream (the origin is allocated by the remote UA). The format list (fmt list) provides information on the Application Layer protocol that the media uses.

  The software opens ports only for RTP and Real-Time Control Protocol (RTCP). Every RTP session has a corresponding RTCP session. Therefore, whenever a media stream uses RTP, the SIP ALG must reserve ports (create pinholes) for both RTP and RTCP traffic. By default, the port number for RTCP is one higher than the RTP port number.

Pinhole Creation

Each pinhole (one for RTP traffic and the other for RTCP traffic) share the same destination IP address. The IP address comes from the c= field in the SDP session description. Because the c= field can appear in either the session-level or the media-level portion of the SDP session description, the parser determines the IP address based on the following rules (in accordance with SDP conventions):

Copyright © 2016, Juniper Networks, Inc.

- First, the SIP ALG parser looks for a c= field containing an IP address in the media level. If there is such a field, the parser extracts that IP address, and the SIP ALG uses that address to create a pinhole for the media.
- If there is no c= field in the media level, the SIP ALG parser extracts the IP address from the c= field in the session level, and the SIP ALG uses that IP address to create a pinhole for the media. If the session description does not contain a c= field in either level, this indicates an error in the protocol stack, and the device drops the packet and logs the event.

The SIP ALG also opens pinholes for signal traffic. These signal pinholes are useful after the previous signal session timeout, and they are also useful for the signal traffic sent to a third-party address that does not match with the previous signal session. The SIP ALG signal pinholes never age out, unlike RTP or RTCP pinholes, where only the destination IP and destination port are specified.

The SIP ALG opens signal pinholes for following headers, if needed:

- VIA
- CONTACT
- ROUTE
- RECORD-ROUTE

The SIP ALG needs the following information to create a pinhole. This information either comes from the SDP session description or from the SIP headers (as listed above).

- Protocol—UDP or TCP.
- Source IP—Unknown.
- Source port—Unknown.
- Destination IP—The parser extracts the destination IP address from the c= field at the media or session level.
- Destination port—The parser extracts the destination port number for RTP from the m= field in the media level and calculates the destination port number for RTCP using the following formula:

  RTP port number + one
- Lifetime—This value indicates the length of time (in seconds) during which a pinhole is open to allow a packet through. A packet must go through the pinhole before the lifetime expires. When the lifetime expires, the SIP ALG removes the pinhole.

  When a packet goes through the pinhole within the lifetime period, immediately afterwards the SIP ALG removes the pinhole for the direction from which the packet came.

  Figure 19 describes a call setup between two SIP clients and how the SIP ALG creates pinholes to allow RTP and RTCP traffic. The illustration assumes that the device has a policy that permits SIP, thus opening port 5060 for SIP signaling messages.

Copyright © 2016, Juniper Networks, Inc.

Figure 19: SIP ALG Call Setup

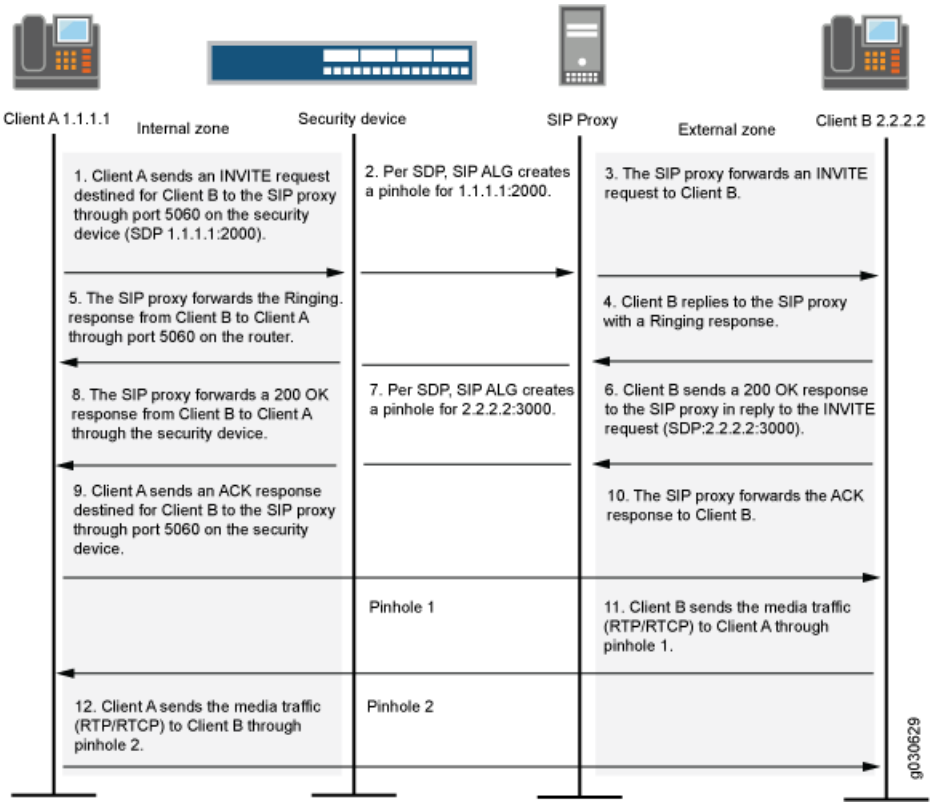

NOTE: The SIP ALG does not create pinholes for RTP and RTCP traffic when the destination IP address is 0.0.0.0, which indicates that the session is on hold. To put a session on hold during a telephone communication, for example, Client A sends Client B a SIP message in which the destination IP address is 0.0.0.0. Doing so indicates to Client B that it should not send any media until further notice. If Client B sends media anyway, the device drops the packets.

Related Documentation

- ALG Overview on page 3
- Understanding SIP ALG Request Methods on page 247
- Understanding the SIP ALG and NAT on page 256
- SIP ALG Configuration Overview on page 248

Understanding IPv6 Support for SIP ALG

IPv6 is supported on the SIP ALG along with NAT-PT mode and NAT64 address translation.

The SIP ALG processes the IPv6 address in the same way it processes the IPv4 address for updating the payload if NAT is configured and opening pinholes for future traffic.

Copyright © 2016, Juniper Networks, Inc.

Special processing occurs for the following formats:

- **IPv6 in SIP URIs**—The SIP URI looks the same as a URI with IPv4 addresses. As in all URIs, an IPv6 address is enclosed in square brackets. The IPv6 address blocks are separated by colons. In many notations, a colon separates the hostname or IP address from the protocol port. To parse the full IPv6 address and separate the port, the address is encapsulated within square brackets
- **IPv6 in SDP**—IPv6 addresses in the Session Description Protocol (SDP) have the IP6 marker.
- The SIP ALG with IPv6 support has the following limitation:
  - When NAT64 with persistent NAT is implemented, the SIP ALG adds the NAT translation to the persistent NAT binding table if NAT is configured on the Address of Record (AOR). Because persistent NAT cannot duplicate the address configured, coexistence of NAT66 and NAT64 configured on the same address is not supported.

    Only one binding is created for the same source IP address.

Related Documentation

- Understanding the SIP ALG on page 241

Understanding SIP ALG Request Methods

The Session Initiation Protocol (SIP) transaction model includes a number of request and response messages, each of which contains a *method* field that denotes the purpose of the message.

Junos OS supports the following method types and response codes:

- INVITE—A user sends an INVITE request to invite another user to participate in a session. The body of an INVITE request can contain the description of the session.
- ACK—The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE request. If the original INVITE request did not contain the session description, the ACK request must include it.
- OPTIONS—The User Agent (UA) obtains information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports.
- BYE—A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session.
- CANCEL—A user sends a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE had sent a final response for the INVITE before it received the CANCEL.
- REGISTER—A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user.

Copyright © 2016, Juniper Networks, Inc.

- Info—Used to communicate mid-session signaling information along the signaling path for the call.
- Subscribe—Used to request current state and state updates from a remote node.
- Notify—Sent to inform subscribers of changes in state to which the subscriber has a subscription.
- Refer—Used to refer the recipient (identified by the Request-URI) to a third party by the contact information provided in the request.

  For example, if user A in a private network refers user B, in a public network, to user C, who is also in the private network, the SIP Application Layer Gateway (ALG) allocates a new IP address and port number for user C so that user C can be contacted by user B. If user C is registered with a registrar, however, its port mapping is stored in the ALG Network Address Translation (NAT) table and is reused to perform the translation.
- Update—Used to open pinhole for new or updated SDP information. The Via:, From:, To:, Call-ID:, Contact:, Route:, and Record-Route: header fields are modified.
- 1*xx*, 202, 2*xx*, 3*xx*, 4*xx*, 5*xx*, 6*xx* Response Codes—Used to indicate the status of a transaction. Header fields are modified.

Related Documentation

- ALG Overview on page 3
- Understanding the SIP ALG on page 241

SIP ALG Configuration Overview

The Session Initiation Protocol Application Layer Gateway (SIP ALG) is disabled by default on high-end devices—it should be enabled using the CLI if required. On other devices, it is enabled by default. To fine-tune SIP ALG operations use the following instructions:

1. Control SIP call activity. For instructions, see "Example: Setting SIP ALG Call Duration and Timeouts" on page 250.
2. Protect the SIP proxy server from denial-of-service (DoS) flood attacks. For instructions, see "Example: Configuring SIP ALG DoS Attack Protection" on page 252.
3. Enable unknown messages to pass when the session is in Network Address Translation (NAT) mode and route mode. For instructions, see "Example: Allowing Unknown SIP ALG Message Types" on page 253.
4. Accommodate proprietary SIP call flows. For instructions, see:
   - Retaining SIP ALG Hold Resources (J-Web Procedure) on page 255
   - Retaining SIP ALG Hold Resources (CLI Procedure) on page 255

Related Documentation

- Understanding the SIP ALG on page 241
- Understanding the SIP ALG and NAT on page 256

Copyright © 2016, Juniper Networks, Inc.

- Verifying SIP ALG Configurations on page 302

Understanding SIP ALG Call Duration and Timeouts

The call duration and timeout features give you control over Session Initiation Protocol (SIP) call activity and help you to manage network resources.

Typically a call ends when one of the clients sends a BYE or CANCEL request. The SIP Application Layer Gateway (ALG) intercepts the BYE or CANCEL request and removes all media sessions for that call. There could be reasons or problems preventing clients in a call from sending BYE or CANCEL requests, for example, a power failure. In this case, the call might go on indefinitely, consuming resources on the device.

A call can have one or more voice channels. Each voice channel has two sessions (or two media streams), one for Real-Time Transport Protocol (RTP) traffic and one for Real-Time Control Protocol (RTCP) signaling. When managing the sessions, the device considers the sessions in each voice channel as one group. Timeouts and call duration settings apply to a group as opposed to each session.

The following parameters govern SIP call activity:

- **inactive-media-timeout**—This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the temporary openings (pinholes) in the firewall the SIP ALG opened for media are closed. The default setting is 120 seconds, and the range is 10 through 2550 seconds. Note that upon timeout, resources for media (sessions and pinholes) are removed and SIP calls on the device will also be terminated if all the media resources of this call are removed.
- **maximum-call-duration**—This parameter sets the absolute maximum length of a call. When a call exceeds this parameter setting, the SIP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, and the range is 3 through 720 minutes.
- **t1-interval**—This parameter specifies the roundtrip time estimate, in seconds, of a transaction between endpoints. The default is 500 milliseconds. Because many SIP timers scale with the t1-interval (as described in RFC 3261), when you change the value of the t1-interval timer, those SIP timers also are adjusted.
- **t4-interval**—This parameter specifies the maximum time a message remains in the network. The default is 5 seconds and the range is 5 through 10 seconds. Because many SIP timers scale with the t4-interval (as described in RFC 3261), when you change the value of the t4-interval timer, those SIP timers also are adjusted.
- **c-timeout**—This parameter specifies the INVITE transaction timeout at the proxy, in minutes; the default is 3. Because the SIP ALG is in the middle, instead of using the INVITE transaction timer value B (which is (64 * T1) = 32 seconds), the SIP ALG gets its timer value from the proxy.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding the SIP ALG on page 241
- SIP ALG Configuration Overview on page 248
- Example: Setting SIP ALG Call Duration and Timeouts on page 250

Example: Setting SIP ALG Call Duration and Timeouts

This example shows how to set the call duration and the media inactivity timeout.

Requirements

Before you begin, review the call duration and timeout features used to control SIP call activity. See "Understanding SIP ALG Call Duration and Timeouts" on page 249.

Overview

The call duration and inactivity media timeout features help you to conserve network resources and maximize throughput.

The **maximum-call-duration** parameter sets the maximum allowable length of time a call can be active. When the duration is exceeded, the SIP ALG tears down the call and releases the media sessions. The default setting is 720 minutes, and the range is 3 through 720 minutes. This setting also frees up bandwidth in cases where calls fail to properly terminate.

The **inactive-media-timeout** parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTPC) traffic within a group. Each time an RTP or RTCP packet occurs within a call, this timeout resets. When the period of inactivity exceeds this setting, the SIP ALG temporary openings (pinholes) for media in the firewall are closed. The default setting is 120 seconds, and the range is 10 through 2550 seconds. Upon timeout, while resources for media (sessions and pinholes) are removed, the call is not terminated.

In this example, the call duration is set to 36000 seconds and the media inactivity timeout is set to 90 seconds.

Configuration

GUI Step-by-Step Procedure

To set the SIP ALG call duration and the media inactivity timeout:

1. Select **Configure** >**Security** >**ALG**.
2. Select the **SIP** tab.
3. In the Maximum call duration field, type **600**.
4. In the Inactive media timeout field, enter **90**.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options** >**Commit**.

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To set the SIP ALG call duration and the media inactivity timeout:

1. Configure the SIP ALG call duration.

   ```
   [edit]
   user@host# set security alg sip maximum-call-duration 600
   ```

2. Configure the SIP ALG inactivity media timeout.

   ```
   [edit]
   user@host# set security alg sip inactive-media-timeout 90
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sip** command.

Related Documentation

- SIP ALG Configuration Overview on page 248
- Verifying SIP ALG Configurations on page 302

Understanding SIP ALG DoS Attack Protection

The ability of the Session Initiation Protocol (SIP) proxy server to process calls can be impacted by repeat SIP INVITE requests—requests that it initially denied. The denial-of-service (DoS) protection feature enables you to configure the device to monitor INVITE requests and proxy server replies to them. If a reply contains a 3*xx,* 4*xx*, or 5*xx* response code other than 401, 407, 487, and 488 that are not real failure responses, then the request should not be blocked. See "Classes of SIP Responses" on page 263. The ALG stores the source IP address of the request and the IP address of the proxy server in a table. Subsequently, the device checks all INVITE requests against this table and, for a configurable number of seconds (the default is 3), discards any packets that match entries in the table. You can configure the device to monitor and deny repeat INVITE requests to all proxy servers, or you can protect a specific proxy server by specifying the destination IP address. SIP attack protection is configured globally.

NOTE: IPv6 is supported on the SIP ALG along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation.

The type of the <destination-ip-address> is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Related Documentation

- Understanding the SIP ALG on page 241
- SIP ALG Configuration Overview on page 248

Copyright © 2016, Juniper Networks, Inc.

• Example: Configuring SIP ALG DoS Attack Protection on page 252

Example: Configuring SIP ALG DoS Attack Protection

This example shows how to configure the DoS attack protection feature.

Requirements

Before you begin, review the DoS attack protection feature used to control SIP call activity. See "Understanding SIP ALG DoS Attack Protection" on page 251.

Overview

The ability of the SIP proxy server to process calls can be impacted by repeat SIP INVITE requests—requests that the server initially denied. The DoS protection feature enables you to configure the device to monitor INVITE requests and proxy server replies to them.

In this example, the device is configured to protect a single SIP proxy server (1.1.1.3) from repeat INVITE requests to which it has already been denied service. Packets are dropped for a period of 5 seconds, after which the device resumes forwarding INVITE requests from those sources.

Configuration

GUI Step-by-Step Procedure

To configure SIP ALG DoS attack protection:

1. Select **Configure>Security>ALG**.
2. Select the **SIP** tab.
3. In the Enable attack protection area, click the **Selected servers** option.
4. In the Destination IP box, enter **1.1.1.3** and click **Add**.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To configure SIP ALG DoS attack protection:

1. Configure the device to protect a single SIP proxy server.

```
[edit]
user@host# set security alg sip application-screen protect deny destination-ip 1.1.1.3
```

NOTE: IPv6 is supported on the SIP ALG along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation.

The type of the <destination-ip-address> is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Copyright © 2016, Juniper Networks, Inc.

2. Configure the device for the deny timeout period.

```
[edit]
user@host# set security alg sip application-screen protect deny timeout 5
```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sip** command.

Related Documentation

- SIP ALG Configuration Overview on page 248
- Verifying SIP ALG Configurations on page 302

Understanding SIP ALG Unknown Message Types

This feature enables you to specify how unidentified Session Initiation Protocol (SIP) messages are handled by the device. The default is to drop unknown (unsupported) messages.

We do not recommend permitting unknown messages because they can compromise security. However, in a secure test or production environment, this command can be useful for resolving interoperability issues with disparate vendor equipment. Permitting unknown SIP messages can help you get your network operational so you can later analyze your voice-over-IP (VoIP) traffic to determine why some messages were being dropped. The unknown SIP message type feature enables you to configure the device to accept SIP traffic containing unknown message types in both Network Address Translation (NAT) mode and route mode.

NOTE: This option applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol and you have configured the device to permit unknown message types, the message is forwarded without processing.

Related Documentation

- Understanding the SIP ALG on page 241
- SIP ALG Configuration Overview on page 248
- Example: Allowing Unknown SIP ALG Message Types on page 253

Example: Allowing Unknown SIP ALG Message Types

This example shows how to allow unknown message types.

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, review how unidentified SIP messages are handled by the device. See "Understanding SIP ALG Unknown Message Types" on page 253.

Overview

In this example, you configure the device to allow unknown message types in SIP traffic in both NAT mode and route mode. The default is to drop unknown (unsupported) messages.

Configuration

GUI Step-by-Step Procedure

To allow unknown SIP ALG message types:

1. Select **Configure>Security>ALG**.
2. Select the **SIP** tab.
3. Select the **Enable Permit NAT applied** check box.
4. Select the **Enable Permit routed** check box.
5. Click **OK** to check your configuration and save it as a candidate configuration.
6. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

To allow unknown SIP ALG message types:

1. Configure the device to allow unknown message types in SIP traffic.

```
[edit]
user@host# set security alg sip application-screen unknown-message
  permit-nat-applied permit-routed
```

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

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

Verification

To verify the configuration is working properly, enter the **show security alg sip** command.

Related Documentation

- SIP ALG Configuration Overview on page 248
- Verifying SIP ALG Configurations on page 302

Understanding SIP ALG Hold Resources

When a user puts a call on hold, the Session Initiation Protocol Application Layer Gateway (SIP ALG) releases Session Description Protocol (SDP) media resources, such as pinholes and translation contexts. When the user resumes the call, an INVITE request message negotiates a new SDP offer and answer and the SIP ALG reallocates resources for the media stream. This can result in new translated IP address and port numbers for the

Copyright © 2016, Juniper Networks, Inc.

media description even when the media description is the same as the previous description. This is compliant with *RFC 3264 An Offer/Answer Model with the Session Description Protocol (SDP)*.

Some proprietary SIP implementations have designed call flows so that the User Agent (UA) module ignores the new SDP INVITE offer and continues to use the SDP offer of the previous negotiation. To accommodate this functionality, you must configure the device to retain SDP media resources when a call is put on hold for reuse when the call is resumed.

Related Documentation

- Understanding the SIP ALG on page 241
- SIP ALG Configuration Overview on page 248
- Retaining SIP ALG Hold Resources (J-Web Procedure) on page 255
- Retaining SIP ALG Hold Resources (CLI Procedure) on page 255

Retaining SIP ALG Hold Resources (CLI Procedure)

To accommodate proprietary SIP call flows:

```
user@host# set security alg sip retain-hold-resource
```

Related Documentation

- Understanding SIP ALG Hold Resources on page 254
- SIP ALG Configuration Overview on page 248
- Retaining SIP ALG Hold Resources (J-Web Procedure) on page 255
- Verifying SIP ALG Configurations on page 302

Retaining SIP ALG Hold Resources (J-Web Procedure)

To accommodate proprietary SIP call flows:

1. Select **Configure>Security>ALG**.
2. Select the **SIP** tab.
3. Select the **Enable retail hold resource** check box.
4. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.

Related Documentation

- Understanding SIP ALG Hold Resources on page 254
- SIP ALG Configuration Overview on page 248
- Retaining SIP ALG Hold Resources (CLI Procedure) on page 255
- Verifying SIP ALG Configurations on page 302

Copyright © 2016, Juniper Networks, Inc.

Understanding the SIP ALG and NAT

The Network Address Translation (NAT) protocol enables multiple hosts in a private subnet to share a single public IP address to access the Internet. For outgoing traffic, NAT replaces the private IP address of the host in the private subnet with the public IP address. For incoming traffic, the public IP address is converted back into the private address, and the message is routed to the appropriate host in the private subnet.

Using NAT with the Session Initiation Protocol (SIP) service is more complicated because SIP messages contain IP addresses in the SIP headers as well as in the SIP body. When using NAT with the SIP service, the SIP headers contain information about the caller and the receiver, and the device translates this information to hide it from the outside network. The SIP body contains the Session Description Protocol (SDP) information, which includes IP addresses and port numbers for transmission of the media. The device translates SDP information for allocating resources to send and receive the media.

How IP addresses and port numbers in SIP messages are replaced depends on the direction of the message. For an outgoing message, the private IP address and port number of the client are replaced with the public IP address and port number of the Juniper Networks firewall. For an incoming message, the public address of the firewall is replaced with the private address of the client.

When an INVITE message is sent out across the firewall, the SIP Application Layer Gateway (ALG) collects information from the message header into a call table, which it uses to forward subsequent messages to the correct endpoint. When a new message arrives, for example an ACK or 200 OK, the ALG compares the “From:, To:, and Call-ID:” fields against the call table to identify the call context of the message. If a new INVITE message arrives that matches the existing call, the ALG processes it as a REINVITE.

When a message containing SDP information arrives, the ALG allocates ports and creates a NAT mapping between them and the ports in the SDP. Because the SDP requires sequential ports for the Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) channels, the ALG provides consecutive even-odd ports. If it is unable to find a pair of ports, it discards the SIP message.

IPv6 is supported on the SIP ALG along with NAT-PT mode and NAT64 address translation.

This topic contains the following sections:

- Outgoing Calls on page 257
- Incoming Calls on page 257
- Forwarded Calls on page 258
- Call Termination on page 258
- Call Re-INVITE Messages on page 258
- Call Session Timers on page 258
- Call Cancellation on page 258
- Forking on page 259

Copyright © 2016, Juniper Networks, Inc.

- SIP Messages on page 259
- SIP Headers on page 259
- SIP Body on page 261
- SIP NAT Scenario on page 261
- Classes of SIP Responses on page 263
- NAT Mode in Pure IPv6 Mode (NAT66) for SIP IPv6 ALG on page 264
- NAT-PT on page 264
- NAT64 on page 265
- STUN and SIP ALG on page 265

Outgoing Calls

When a SIP call is initiated with a SIP request message from the internal to the external network, NAT replaces the IP addresses and port numbers in the SDP and binds the IP addresses and port numbers to the Juniper Networks firewall. Via, Contact, Route, and Record-Route SIP header fields, if present, are also bound to the firewall IP address. The ALG stores these mappings for use in retransmissions and for SIP response messages.

The SIP ALG then opens pinholes in the firewall to allow media through the device on the dynamically assigned ports negotiated based on information in the SDP and the Via, Contact, and Record-Route header fields. The pinholes also allow incoming packets to reach the Contact, Via, and Record-Route IP addresses and ports. When processing return traffic, the ALG inserts the original Contact, Via, Route, and Record-Route SIP fields back into packets.

Incoming Calls

Incoming calls are initiated from the public network to public static NAT addresses or to interface IP addresses on the device. Static NATs are statically configured IP addresses that point to internal hosts; interface IP addresses are dynamically recorded by the ALG as it monitors REGISTER messages sent by internal hosts to the SIP registrar. When the device receives an incoming SIP packet, it sets up a session and forwards the payload of the packet to the SIP ALG.

The ALG examines the SIP request message (initially an INVITE) and, based on information in the SDP, opens gates for outgoing media. When a 200 OK response message arrives, the SIP ALG performs NAT on the IP addresses and ports and opens pinholes in the outbound direction. (The opened gates have a short time-to-live, and they time out if a 200 OK response message is not received quickly.)

When a 200 OK response arrives, the SIP proxy examines the SDP information and reads the IP addresses and port numbers for each media session. The SIP ALG on the device performs NAT on the addresses and port numbers, opens pinholes for outbound traffic, and refreshes the timeout for gates in the inbound direction.

When the ACK arrives for the 200 OK, it also passes through the SIP ALG. If the message contains SDP information, the SIP ALG ensures that the IP addresses and port numbers are not changed from the previous INVITE—if they are, the ALG deletes old pinholes and creates new pinholes to allow media to pass through. The ALG also monitors the Via,

Copyright © 2016, Juniper Networks, Inc.

Contact, and Record-Route SIP fields and opens new pinholes if it determines that these fields have changed.

Forwarded Calls

A forwarded call is when, for example, user A outside the network calls user B inside the network, and user B forwards the call to user C outside the network. The SIP ALG processes the INVITE from user A as a normal incoming call. But when the ALG examines the forwarded call from B to C outside the network and notices that B and C are reached using the same interface, it does not open pinholes in the firewall, because media will flow directly between user A and user C.

Call Termination

The BYE message terminates a call. When the device receives a BYE message, it translates the header fields just as it does for any other message. But because a BYE message must be acknowledged by the receiver with a 200 OK, the ALG delays call teardown for 5 seconds to allow time for transmission of the 200 OK.

Call Re-INVITE Messages

Re-INVITE messages add new media sessions to a call and remove existing media sessions. When new media sessions are added to a call, new pinholes are opened in the firewall and new address bindings are created. The process is identical to the original call setup. When all the media sessions or media pinholes are removed from a call, the call is removed when a BYE message is received.

Call Session Timers

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the device is protected should one of the following events occur:

- End systems crash during a call and a BYE message is not received.
- Malicious users never send a BYE in an attempt to attack a SIP ALG.
- Poor implementations of SIP proxy fail to process Record-Route and never send a BYE message.
- Network failures prevent a BYE message from being received.

Call Cancellation

Either party can cancel a call by sending a CANCEL message. Upon receiving a CANCEL message, the SIP ALG closes pinholes through the firewall—if any have been opened—and releases address bindings. Before releasing the resources, the ALG delays the control channel age-out for approximately 5 seconds to allow time for the final 200 OK to pass through. The call is terminated when the 5-second timeout expires, regardless of whether a 487 or non-200 response arrives.

Copyright © 2016, Juniper Networks, Inc.

Forking

Forking enables a SIP proxy to send a single INVITE message to multiple destinations simultaneously. When the multiple 200 OK response messages arrive for the single call, the SIP ALG parses but updates call information with the first 200 OK messages it receives.

SIP Messages

The SIP message format consists of a SIP header section and the SIP body. In request messages, the first line of the header section is the request line, which includes the method type, request-URI, and protocol version. In response messages, the first line is the status line, which contains a status code. SIP headers contain IP addresses and port numbers used for signaling. The SIP body, separated from the header section by a blank line, is reserved for session description information, which is optional. Junos OS currently supports the SDP only. The SIP body contains IP addresses and port numbers used to transport the media.

SIP Headers

In the following sample SIP request message, NAT replaces the IP addresses in the header fields to hide them from the outside network.

```
INVITE bob@10.150.20.5 SIP/2.0
Via: SIP/2.0/UDP 10.150.20.3:5434
From: alice@10.150.20.3
To: bob@10.150.20.5
Call-ID: a12abcde@10.150.20.3
Contact: alice@10.150.20.3:5434
Route: <sip:netscreen@10.150.20.3:5060>
Record-Route: <sip:netscreen@10.150.20.3:5060>
```

How IP address translation is performed depends on the type and direction of the message. A message can be any of the following:

- Inbound request
- Outbound response
- Outbound request
- Inbound response

Table 12 shows how NAT is performed in each of these cases. Note that for several of the header fields the ALG determine more than just whether the messages comes from inside or outside the network. It must also determine what client initiated the call, and whether the message is a request or response.

Copyright © 2016, Juniper Networks, Inc.

Table 12: Requesting Messages with NAT Table

Inbound Request (from public to private)	To:	Replace domain with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	Replace ALG address with local address
	Contact:	None
	Record-Route:	None
	Route:	None
Outbound Response (from private to public)	To:	Replace ALG address with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	N/A
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	None
Outbound Request (from private to public)	To:	None
	From:	Replace local address with ALG address
	Call-ID:	None
	Via:	Replace local address with ALG address
	Request-URI:	None
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	Replace local address with ALG address

Copyright © 2016, Juniper Networks, Inc.

Table 12: Requesting Messages with NAT Table *(continued)*

Outbound Response (from public to private)	To:	None
	From:	Replace ALG address with local address
	Call-ID:	None
	Via:	Replace ALG address with local address
	Request-URI:	N/A
	Contact:	None
	Record-Route:	Replace ALG address with local address
	Route:	Replace ALG address with local address

SIP Body

The SDP information in the SIP body includes IP addresses the ALG uses to create channels for the media stream. Translation of the SDP section also allocates resources, that is, port numbers to send and receive the media.

The following excerpt from a sample SDP section shows the fields that are translated for resource allocation.

```
o=user 2344234 55234434 IN IP4 10.150.20.3
c=IN IP4 10.150.20.3
m=audio 43249 RTP/AVP 0
```

SIP messages can contain more than one media stream. The concept is similar to attaching multiple files to an e-mail message. For example, an INVITE message sent from a SIP client to a SIP server might have the following fields:

```
c=IN IP4 10.123.33.4
m=audio 33445 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33447 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33449 RTP/AVP 0
```

Junos OS supports up to 6 SDP channels negotiated for each direction, for a total of 12 channels per call. For more information, see "SDP Session Descriptions" on page 244.

SIP NAT Scenario

Figure 20 and Figure 21 show a SIP call INVITE and 200 OK. In Figure 20, ph1 sends a SIP INVITE message to ph2. Note how the IP addresses in the header fields—shown in bold font—are translated by the device.

The SDP section of the INVITE message indicates where the caller is willing to receive media. Note that the Media Pinhole contains two port numbers, 52002 and 52003, for RTCP and RTP. The Via/Contact Pinhole provides port number 5060 for SIP signaling.

Copyright © 2016, Juniper Networks, Inc.

Observe how, in the 200 OK response message in Figure 21, the translations performed in the INVITE message are reversed. The IP addresses in this message, being public, are not translated, but gates are opened to allow the media stream access to the private network.

Figure 20: SIP NAT Scenario 1

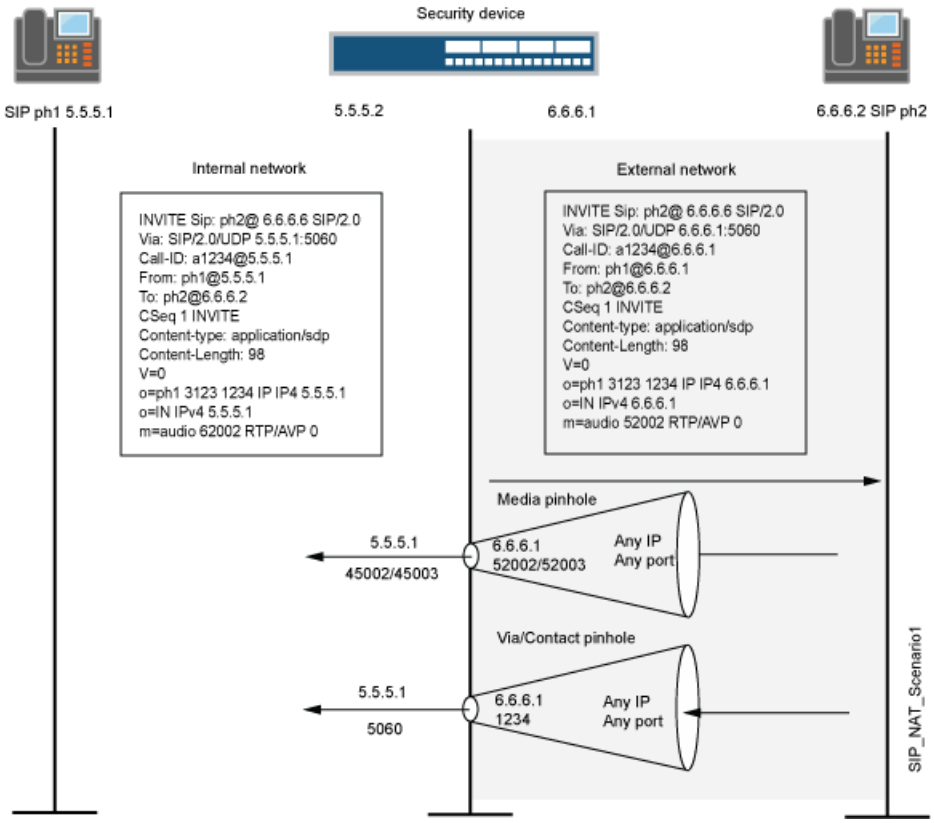

Copyright © 2016, Juniper Networks, Inc.

Figure 21: SIP NAT Scenario 2

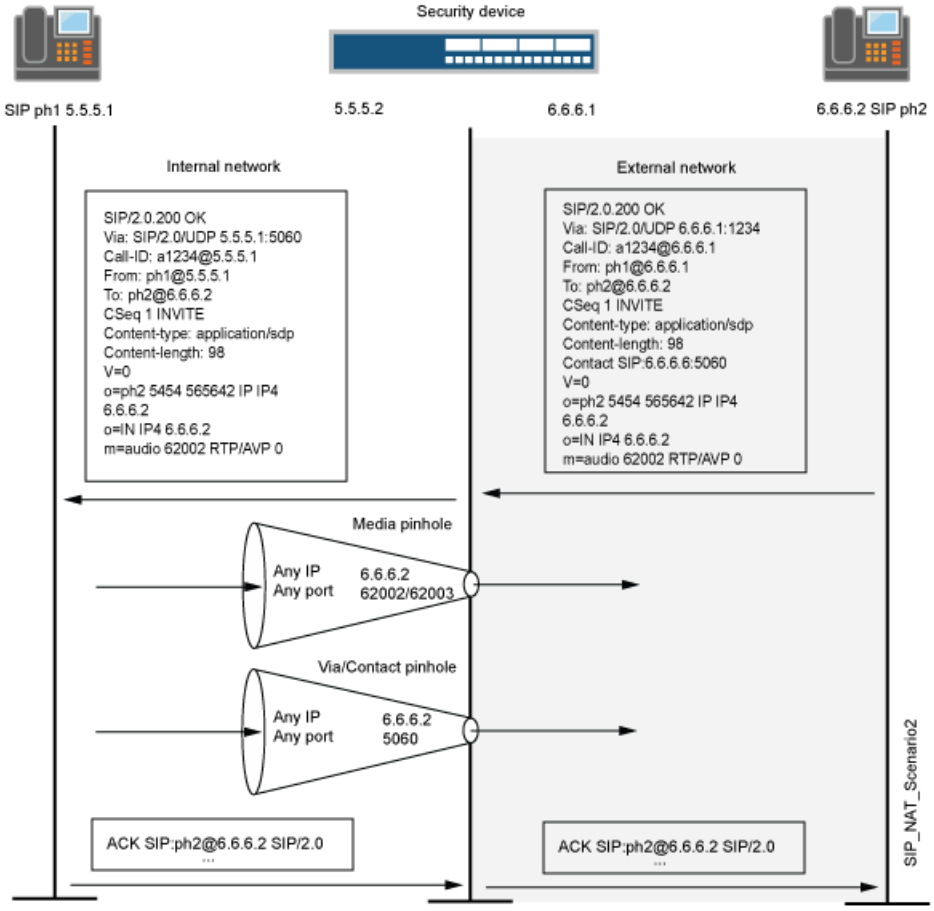

Classes of SIP Responses

SIP responses provide status information about SIP transactions and include a response code and a reason phrase. SIP responses are grouped into the following classes:

- Informational (100 to 199)—Request received, continuing to process the request.
- Success (200 to 299)—Action successfully received, understood, and accepted.
- Redirection (300 to 399)—Further action required to complete the request.
- Client Error (400 to 499)—Request contains bad syntax or cannot be fulfilled at this server.
- Server Error (500 to 599)—Server failed to fulfill an apparently valid request.
- Global Failure (600 to 699)—Request cannot be fulfilled at any server.

Table 13 provides a complete list of current SIP responses.

Copyright © 2016, Juniper Networks, Inc.

Table 13: SIP Responses

Informational	100 Trying	180 Ringing	181 Call is being forwarded
	182 Queued	183 Session progress	
Success	200 OK	202 Accepted	
Redirection	300 Multiple choices	301 Moved permanently	302 Moved temporarily
	305 Use proxy	380 Alternative service	
Client Error	400 Bad request	401 Unauthorized	402 Payment required
	403 Forbidden	404 Not found	405 Method not allowed
	406 Not acceptable	407 Proxy authentication required	408 Request time-out
	409 Conflict	410 Gone	411 Length required
	413 Request entity too large	414 Request URL too large	415 Unsupported media type
	420 Bad extension	480 Temporarily not available	481 Call leg/transaction does not exist
	482 Loop detected	483 Too many hops	484 Address incomplete
	485 Ambiguous	486 Busy here	487 Request canceled
	488 Not acceptable here		
Server Error	500 Server internal error	501 Not implemented	502 Bad gateway
	502 Service unavailable	504 Gateway time-out	505 SIP version not supported
Global Failure	600 Busy everywhere	603 Decline	604 Does not exist anywhere
	606 Not acceptable		

NAT Mode in Pure IPv6 Mode (NAT66) for SIP IPv6 ALG

The SIP IPv6 ALG supports NAT66 just like NAT44. NAT66 (IPv6 NAT) provides source NAT and static NAT functions similar to NAT44 (IPv4 NAT).

NAT-PT

Network Address Translation Protocol Translation (NAT-PT) (RFC 2766) is a protocol translation mechanism that allows communication between IPv6-only and IPv4-only nodes through protocol-independent translation of IPv4 and IPv6 datagrams, requiring no state information for the session.

Copyright © 2016, Juniper Networks, Inc.

NAT-PT is implemented by normal NAT from IPv6 address to IPv4 address and vice versa. The SIP ALG processes those address translations in the payload just as the addresses are processed in normal NAT.

NAT-PT binds the addresses in the IPv6 network with addresses in the IPv4 network and vice versa to provide transparent routing for the datagrams traversing between address realms.

The main advantage of NAT-PT is that the end devices and networks can run either IPv4 addresses or IPv6 addresses and traffic can be started from any side.

NAT64

NAT64 is a mechanism to allow IPv6 hosts to communicate with IPv4 servers. NAT64 is required to keep the IPv6 to IPv4 address mapping. Such address mapping is either statically configured by the system administrator (stateless translation), or more frequently, created automatically when the first packet from the IPv6 network reaches NAT64 to be translated (stateful).

On SRX Series devices, NAT64 is implemented by persistent NAT. When the first SIP request message (first packet should be only from IPv6) transverses the DUT, address binding is created and then the packets can flow in both directions.

The NAT64 mechanism translates IPv6 packets to IPv4 packets and vice versa, which allows IPv6 clients to contact to the IPv4 servers using unicast UDP, TCP, or ICMP. The NAT-PT and NAT64 behavior seems similar, but these mechanisms are implemented differently.

When NAT64 with persistent NAT is implemented, the SIP ALG with IPv6 support adds the NAT translation to the persistent NAT binding table if NAT is configured on the address of record. Because persistent NAT cannot duplicate the address configured, coexistence of NAT66 and NAT64 configured on the same address is not supported.

Only one binding is created for the same source IP address.

STUN and SIP ALG

Session Traversal Utilities for NAT (STUN) is a solution to make VoIP work through NAT and firewall.

Previously STUN worked without the SIP ALG. This means that the SIP ALG was not involved when persistent NAT was configured.

STUN can coexist with the SIP ALG and SIP ALG is involved when persistent NAT is configured.

Related Documentation

- Understanding the SIP ALG on page 241
- Understanding Incoming SIP ALG Call Support Using the SIP Registrar and NAT on page 266
- Example: Configuring Interface Source NAT for Incoming SIP Calls on page 267

Copyright © 2016, Juniper Networks, Inc.

- Decreasing Network Complexity by Configuring a Source NAT Pool for Incoming SIP Calls on page 273
- Example: Configuring Static NAT for Incoming SIP Calls on page 281
- Example: Configuring the SIP Proxy in the Private Zone and NAT in the Public Zone on page 288
- Example: Configuring a Three-Zone SIP ALG and NAT Scenario on page 294

Understanding Incoming SIP ALG Call Support Using the SIP Registrar and NAT

Session Initiation Protocol (SIP) registration provides a discovery capability by which SIP proxies and location servers can identify the location or locations where users want to be contacted. A user registers one or more contact locations by sending a REGISTER message to the registrar. The To and Contact fields in the REGISTER message contain the address-of-record Uniform Resource Identifier (URI) and one or more contact URIs, as shown in Figure 22. Registration creates bindings in a location service that associates the address-of-record with the contact address or addresses.

The device monitors outgoing REGISTER messages, performs Network Address Translation (NAT) on these addresses, and stores the information in an Incoming NAT table. Then, when an INVITE message is received from outside the network, the device uses the Incoming NAT table to identify which internal host to route the INVITE message to. You can take advantage of SIP proxy registration service to allow incoming calls by configuring interface source NAT or NAT pools on the egress interface of the device. Interface source NAT is adequate for handling incoming calls in a small office, whereas we recommend setting up source NAT pools for larger networks or an enterprise environment.

NOTE: Incoming call support using interface source NAT or a source NAT pool is supported for SIP and H.323 services only. For incoming calls, Junos OS currently supports UDP and TCP only. Domain name resolution is also currently not supported; therefore, URIs must contain IP addresses, as shown in Figure 22.

Copyright © 2016, Juniper Networks, Inc.

Figure 22: Using the SIP Registrar

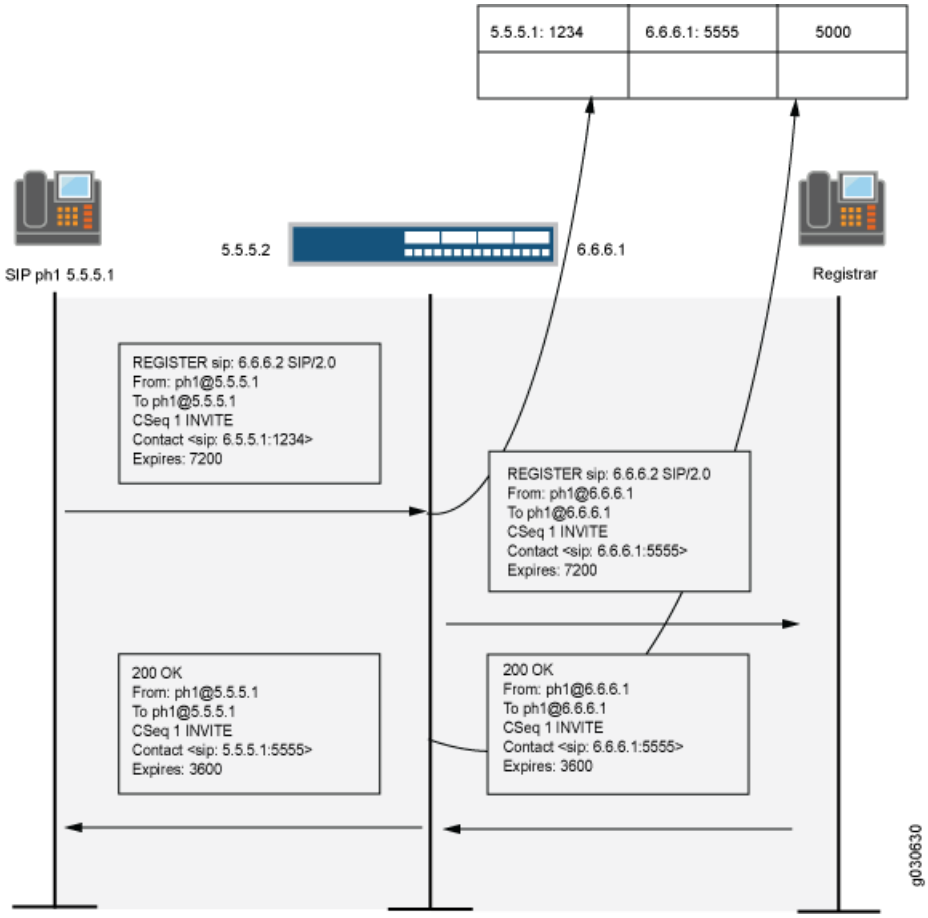

Related Documentation

- ALG Overview on page 3
- Understanding the SIP ALG and NAT on page 256
- SIP ALG Configuration Overview on page 248

Example: Configuring Interface Source NAT for Incoming SIP Calls

This example shows how to configure a source NAT rule on a public zone interface allowing NAT to be used for incoming SIP calls.

- Requirements on page 268
- Overview on page 268
- Configuration on page 269
- Verification on page 272

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, understand how NAT works with the SIP ALG. See "Understanding the SIP ALG and NAT" on page 256.

Overview

In a two-zone scenario with the SIP proxy server in an external zone, you can use NAT for incoming calls by configuring a source NAT rule on the interface in the public or external zone.

In this example (see Figure 23), phone1 is on the ge-0/0/0 interface in the private zone, and phone2 and the proxy server are on the ge-0/0/2 interface in the public zone. You configure a source NAT rule on the public interface ge-0/0/2.0.

Topology

Figure 23 shows source NAT for incoming SIP calls.

Figure 23: Source NAT for Incoming SIP Calls

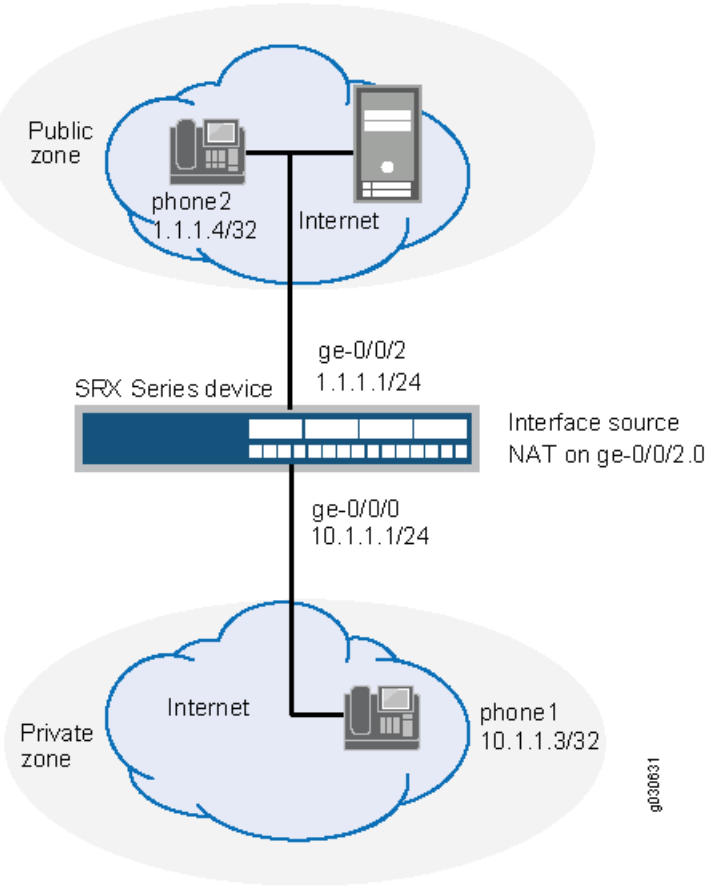

Copyright © 2016, Juniper Networks, Inc.

In this example, after creating zones called private and public and assigning them to interfaces, you configure address books to be used in the source NAT rule set. Then you configure source NAT by defining a rule set called sip-phones and a rule called phone1 that matches any packets from the source address 10.1.1.3/32.

Finally, you create security policies to allow all SIP traffic between the private and public zones.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address phone1 10.1.1.3/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address proxy 1.1.1.3/32
set security zones security-zone public address-book address phone2 1.1.1.4/32
set security zones security-zone public interfaces ge-0/0/2.0
set security nat source rule-set sip-phones from zone private
set security nat source rule-set sip-phones to zone public
set security nat source rule-set sip-phones rule phone1 match source-address 10.1.1.3/32
set security nat source rule-set sip-phones rule phone1 then source-nat interface
set security policies from-zone private to-zone public policy outgoing match
  source-address phone1
set security policies from-zone private to-zone public policy outgoing match
  destination-address phone2
set security policies from-zone private to-zone public policy outgoing match
  destination-address proxy
set security policies from-zone private to-zone public policy outgoing match application
  junos-sip
set security policies from-zone private to-zone public policy outgoing then permit
set security policies from-zone public to-zone private policy incoming match
  source-address phone2
set security policies from-zone public to-zone private policy incoming match
  destination-address phone1
set security policies from-zone public to-zone private policy incoming match
  source-address proxy
set security policies from-zone public to-zone private policy incoming match application
  junos-sip
set security policies from-zone public to-zone private policy incoming then permit
```

Step-by-Step Procedure

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

To configure a source NAT rule on a public zone interface:

1. Configure interfaces.

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

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure zones and assign them to the interfaces.

```
[edit security zones]
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone public interfaces ge-0/0/2.0
```

3. Configure address books and create addresses.

```
[edit security zones]
user@host# set security-zone private address-book address phone1 10.1.1.3/32
user@host# set security-zone public address-book address proxy 1.1.1.3/32
user@host# set security-zone public address-book address phone2 1.1.1.4/32
```

4. Configure a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set sip-phones from zone private
user@host# set rule-set sip-phones to zone public
user@host# set rule-set sip-phones rule phone1 match source-address 10.1.1.3/32
user@host# set rule-set sip-phones rule phone1 then source-nat interface
```

5. Enable persistent source NAT translation.

```
[edit security nat source]
user@host# set address-persistent
```

6. Configure a security policy to allow outgoing SIP traffic.

```
[edit security policies from-zone private to-zone public policy outgoing]
user@host# set match source-address phone1
user@host# set match destination-address phone2
user@host# set match destination-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

7. Configure a security policy to allow incoming SIP traffic.

```
[edit security policies from-zone public to-zone private policy incoming]
user@host# set match source-address phone2
user@host# set match destination-address phone1
user@host# set match source-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security policies**, and **show security nat** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
   unit 0 {
      family inet {
         address 10.1.1.1/24;
      }
   }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
    ge-0/0/2 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }

[edit]
user@host# show security zones
security-zone private {
  address-book {
    address phone1 10.1.1.3/32;
  }
  interfaces {
    ge-0/0/0.0;
  }
}
  security-zone public {
    address-book {
      address proxy 1.1.1.3/32;
      address phone2 2.2.2.4/32;
    }
    interfaces {
      ge-0/0/2.0;
    }
  }
[edit]
user@host# show security nat
source {
  rule-set sip-phones {
    from zone private;
    to zone public;
    rule phone1 {
      match {
        source-address 10.1.1.3/32;
      }
      then {
        source-nat {
          interface;
        }
      }
    }
  }
}
[edit]
user@host# show security policies
from-zone private to-zone public {
  policy outgoing {
    match {
      source-address phone1;
      destination-address [ phone2 proxy ];
      application junos-sip;
    }
    then {
```

```
        permit;
      }
    }
  }
  from-zone public to-zone private {
    policy incoming {
      match {
        source-address [ phone2 proxy ];
        destination-address phone1 ;
        application junos-sip;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Rule Usage on page 272
- Verifying SIP ALG Status on page 273

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat source rule all
source NAT rule: phone1      Rule-set: sip-phones
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : private
  To zone                    : public
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Destination port         : 0               - 0
  Action                       : interface
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Verifying SIP ALG Status

Purpose Verify that SIP ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the SIP ALG status as follows:

- Enabled—Shows the SIP ALG is enabled.
- Disabled—Shows the SIP ALG is disabled.

Related Documentation
- Verifying SIP ALG Configurations on page 302

Decreasing Network Complexity by Configuring a Source NAT Pool for Incoming SIP Calls

This example shows how to decrease network complexity by configuring a source NAT pool on an external interface to enable NAT for incoming SIP calls.

- Requirements on page 273
- Overview on page 273
- Configuration on page 275
- Verification on page 279

Requirements

Before you begin, understand how NAT works with the SIP ALG. See "Understanding the SIP ALG and NAT" on page 256.

Overview

In a two-zone scenario with the SIP proxy server in an external or public zone, you can use NAT for incoming calls by configuring a NAT pool on the interface to the public zone.

Copyright © 2016, Juniper Networks, Inc.

In this example (see Figure 24), phone1 is in the private zone, and phone2 and the proxy server are in the public zone. You configure a source NAT pool to do NAT. You also create a policy that permits SIP traffic from the private to the public zone. This enables phone1 in the private zone to register with the proxy server in the public zone, and it also enables incoming calls from the public zone to the private zone.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 24 shows source NAT pool for incoming calls.

Figure 24: Source NAT Pool for Incoming SIP Calls

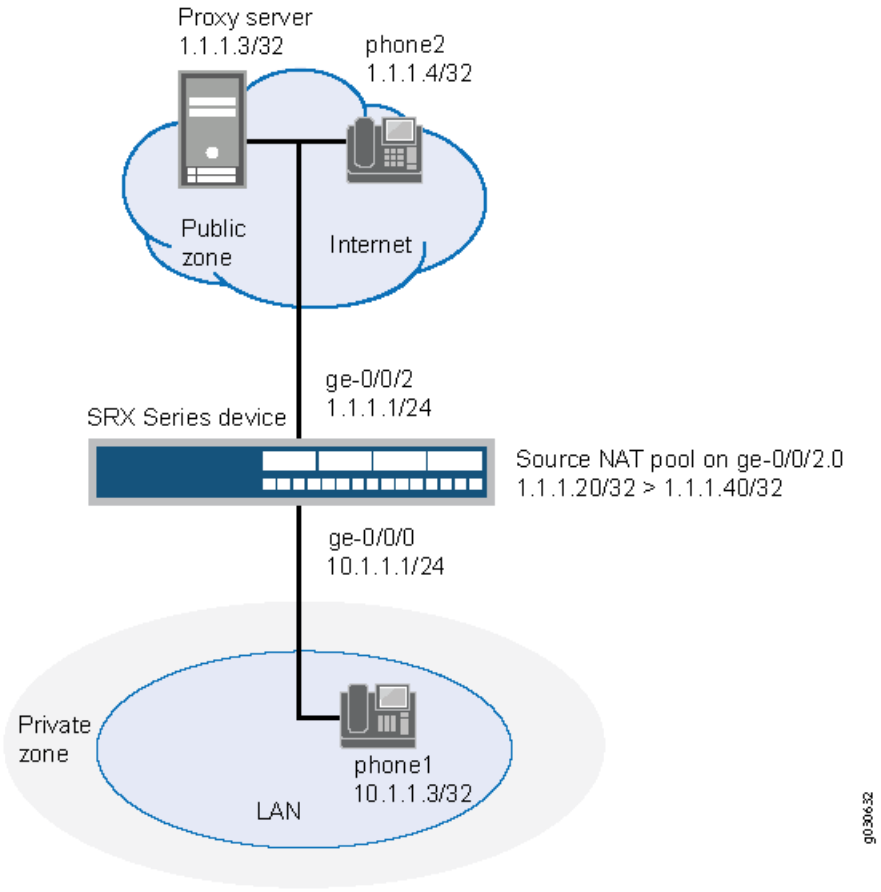

In this example, you configure source NAT as follows:

- Define source NAT pool called sip-nat-pool to contain the IP address range from 1.1.1.20/32 through 1.1.1.40/32.
- Create a source NAT rule set called sip-nat with a rule sip-r1 to match packets from the private zone to the public zone with the source IP address 10.1.1.3/24. For matching packets, the source address is translated to one of the IP address in sip-nat-pool.
- Configure proxy ARP for the addresses 1.1.1.20/32 through 1.1.1.40/32 on interface ge-0/0/2.0. This allows the system to respond to ARP requests received on the interface for these addresses.

Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

Copyright © 2016, Juniper Networks, Inc.

network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address phone1 10.1.1.3/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address proxy 1.1.1.3/32
set security zones security-zone public address-book address phone2 1.1.1.4/32
set security zones security-zone public interfaces ge-0/0/2.0
set security nat source pool sip-nat-pool address 1.1.1.20/32 to 1.1.1.40/32
set security nat source address-persistent
set security nat source rule-set sip-nat from zone private
set security nat source rule-set sip-nat to zone public
set security nat source rule-set sip-nat rule sip-r1 match source-address 10.1.1.3/24
set security nat source rule-set sip-nat rule sip-r1 then source-nat pool sip-nat-pool
set security nat proxy-arp interface ge-0/0/2.0 address 1.1.1.20/32 to 1.1.1.40/32
set security policies from-zone private to-zone public policy outgoing match
  source-address phone1
set security policies from-zone private to-zone public policy outgoing match
  destination-address any
set security policies from-zone private to-zone public policy outgoing match application
  junos-sip
set security policies from-zone private to-zone public policy outgoing then permit
set security policies from-zone public to-zone private policy incoming match
  source-address phone2
set security policies from-zone public to-zone private policy incoming match
  destination-address phone1
set security policies from-zone public to-zone private policy incoming match application
  junos-sip
set security policies from-zone public to-zone private policy incoming then permit
```

Step-by-Step Procedure

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

To configure a source NAT pool for incoming calls:

1. Configure interfaces.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
   user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
   ```

2. Configure zones and assign interfaces to them.

   ```
   [edit security zones]
   user@host# set security-zone private interfaces ge-0/0/0.0
   user@host# set security-zone public interfaces ge-0/0/2.0
   ```

3. Configure address books.

   ```
   [edit security zones]
   user@host# set security-zone private address-book address phone1 10.1.1.3/32
   user@host# set security-zone public address-book address proxy 1.1.1.3/32
   user@host# set security-zone public address-book address phone2 1.1.1.4/32
   ```


Copyright © 2016, Juniper Networks, Inc.

4. Configure a source NAT pool.

```
[edit security nat]
user@host# set source pool sip-nat-pool address 1.1.1.20/32 to 1.1.1.40/32
```

5. Configure a source NAT rule set with a rule.

```
[edit security nat source rule-set sip-nat]
user@host# set from zone private
user@host# set to zone public
user@host# set rule sip-r1 match source-address 10.1.1.3/24
user@host# set rule sip-r1 then source-nat pool sip-nat-pool
```

6. Enable persistent NAT.

```
[edit security nat]
user@host# set source address-persistent
```

7. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/2.0 address 1.1.1.20/32 to 1.1.1.40/32
```

8. Configure a security policy to allow outgoing SIP traffic.

```
[edit security policies from-zone private to-zone public policy outgoing]
set match source-address phone1
set match destination-address any
set match application junos-sip
set then permit
```

9. Configure a security policy to allow incoming SIP traffic.

```
[edit security policies from-zone public to-zone private policy incoming]
set match source-address phone2
set match destination-address phone1
set match application junos-sip
set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
    }
    ge-0/0/2 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
[edit]
user@host# show security zones
  security-zone private {
    address-book {
      address phone1 10.1.1.3/32;
    }
    interfaces {
      ge-0/0/0.0;
    }
  }
    security-zone public {
      address-book {
        address proxy 1.1.1.3/32;
        address phone2 1.1.1.4/32;
      }
      interfaces {
        ge-0/0/2.0;
      }
    }
user@host# show security nat
  source {
    pool sip-nat-pool {
      address {
        1.1.1.20/32 to 1.1.1.40/32;
      }
    }
    address-persistent;
    rule-set sip-nat {
      from zone private;
      to zone public;
      rule sip-r1 {
        match {
          source-address 10.1.1.3/24;
        }
        then {
          source-nat {
            pool {
              sip-nat-pool;
            }
          }
        }
      }
    }
  }
    proxy-arp {
      interface ge-0/0/2.0 {
        address {
          1.1.1.20/32 to 1.1.1.40/32;
        }
      }
    }
[edit]
user@host# show security policies
  from-zone private to-zone public {
    policy outgoing {
```


Copyright © 2016, Juniper Networks, Inc.

```
        match {
            source-address phone1;
            destination-address any;
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
from-zone public to-zone private {
    policy incoming {
        match {
            source-address phone2;
            destination-address phone1;
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 279
- Verifying Source NAT Rule Usage on page 280
- Verifying SIP ALG Status on page 280
- Verifying the Security Polices of SIP ALG on page 281

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command.

```
user@host> show security nat source pool all
Total pools: 1
Pool name          : sip-nat-pool
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
port overloading   : 1
Total addresses    : 21
Translation hits   : 0
Address range                        Single Ports   Twin Ports
          1.1.1.20 - 1.1.1.40            0              0
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The **Translation hits** field shows that there is no traffic used by IP addresses from the source NAT pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command.

```
user@host> show security nat source rule all

source NAT rule: sip-r1      Rule-set: sip-nat
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : private
  To zone                    : public
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Destination port         : 0               - 0
  Action                       : interface
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule.

Verifying SIP ALG Status

Purpose Verify that SIP ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status

ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the SIP ALG status as follows:

Copyright © 2016, Juniper Networks, Inc.

- ·Enabled—Shows the SIP ALG is enabled.
- ·Disabled—Shows the SIP ALG is disabled.

Verifying the Security Polices of SIP ALG

Purpose Verify that the source NAT between public zone and private zone is set.

Action From operational mode, enter the **show security policies** command.

```
user@host> show security policies
from-zone private to-zone public {
    policy outgoing {
        match {
            source-address phone1;
            destination-address any;
            application junos-sip;
        }
        then {
            permit;
        }
    }

from-zone public to-zone private {
    policy incoming {
        match {
            source-address phone2;
            destination-address phone1;
            application junos-sip;
        }
        then {
            permit;
        }
    }
```

Meaning The sample output shows that the source NAT between public zone and private zone is set.

Related Documentation
- Verifying SIP ALG Configurations on page 302

Example: Configuring Static NAT for Incoming SIP Calls

This example shows how to configure a static NAT mapping that allows callers in the private zone to register with the proxy server in the public zone.

- Requirements on page 282
- Overview on page 282
- Configuration on page 283
- Verification on page 287

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, understand how NAT works with the SIP ALG. See "Understanding the SIP ALG and NAT" on page 256.

Overview

When a SIP proxy server is located in an external or public zone, you can configure static NAT on the public interface to enable callers in the private zone to register with the proxy server.

In this example (see Figure 25), phone1 is on the ge-0/0/0 interface in the private zone, and phone2 and the proxy server are on the ge-0/0/2 interface in the public zone. You create a static NAT rule set called incoming-sip with a rule called phone1 to match packets from the public zone with the destination address 1.1.1.3/32. For matching packets, the destination IP address is translated to the private address 10.1.1.3/32. You also create proxy ARP for the address 1.1.1.3/32 on interface ge-0/0/2.0. This allows the system to respond to ARP requests received on the interface for these addresses. Finally, you create a security policy called incoming that allows SIP traffic from the public zone to the private zone.

NOTE: When configuring static NAT for incoming SIP calls, make sure to configure one public address for each private address in the private zone.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 25 shows static NAT for incoming calls.

Figure 25: Static NAT for Incoming Calls

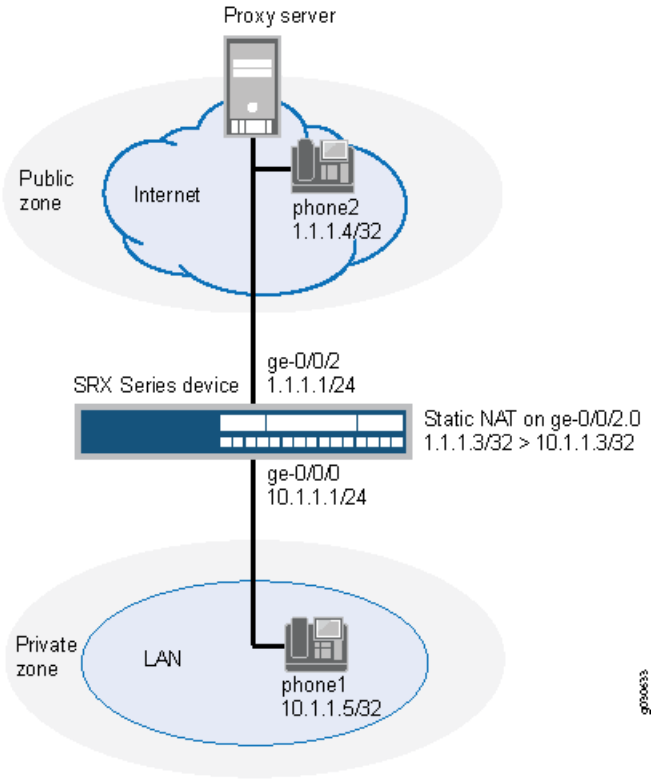

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone private address-book address phone1 10.1.1.5/32
set security zones security-zone public interfaces ge-0/0/2.0
set security zones security-zone public address-book address proxy 1.1.1.3/32
set security zones security-zone public address-book address phone2 1.1.1.4/32
set security nat static rule-set incoming-sip from zone public
set security nat static rule-set incoming-sip rule phone1 match destination-address
    1.1.1.3/32
set security nat static rule-set incoming-sip rule phone1 then static-nat prefix 10.1.1.3/32
set security nat proxy-arp interface ge-0/0/2.0 address 1.1.1.3/32
set security policies from-zone public to-zone private policy incoming match
    source-address phone2
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone public to-zone private policy incoming match
  source-address proxy
set security policies from-zone public to-zone private policy incoming match
  destination-address phone1
set security policies from-zone public to-zone private policy incoming match application
  junos-sip
set security policies from-zone public to-zone private policy incoming then permit
set security policies from-zone private to-zone public policy outgoing match
  source-address phone1
set security policies from-zone private to-zone public policy outgoing match
  destination-address phone2
set security policies from-zone private to-zone public policy outgoing match
  destination-address proxy
set security policies from-zone private to-zone public policy outgoing match application
  junos-sip
set security policies from-zone private to-zone public policy outgoing then permit
```

Step-by-Step Procedure

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

To configure static NAT for incoming calls:

1. Configure interfaces.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
   user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
   ```

2. Create security zones.

   ```
   [edit security zones]
   user@host# set security-zone private interfaces ge-0/0/0.0
   user@host# set security-zone public interfaces ge-0/0/2.0
   ```

3. Assign addresses to the security zones.

   ```
   [edit security zones]
   user@host# set security-zone private address-book address phone1 10.1.1.5/32
   user@host# set security-zone public address-book address proxy 1.1.1.3/32
   user@host# set security-zone public address-book address phone2 1.1.1.4/32
   ```

4. Create a static NAT rule set with a rule.

   ```
   [edit security nat static rule-set incoming-sip]
   user@host# set from zone public
   user@host# set rule phone1 match destination-address 1.1.1.3/32
   user@host# set rule phone1 then static-nat prefix 10.1.1.3/32
   ```

5. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/2.0 address 1.1.1.3/32
   ```

6. Define a security policy to allow incoming SIP traffic.

   ```
   [edit security policies from-zone public to-zone private policy incoming]
   user@host#  set match source-address phone2
   user@host# set match source-address proxy
   user@host# set match destination-address phone1
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set match application junos-sip
user@host# set then permit
```

7. Define a security policy to allow outgoing SIP traffic.

```
[edit security policies from-zone private to-zone public policy outgoing]
user@host#  set match source-address phone1
user@host# set match destination-address phone2
user@host# set match destination-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 1.1.1.1/24;
      }
    }
  }

[edit]
user@host# show security zones
security-zone private {
  address-book {
    address phone1 10.1.1.5/32;
  }
  interfaces {
    ge-0/0/0.0;
  }
}
  security-zone public {
    address-book {
      address proxy 1.1.1.3/32;
      address phone2 1.1.1.4/32;
    }
    interfaces {
      ge-0/0/2.0;
    }
  }

[edit]
user@host# show security nat
```

Copyright © 2016, Juniper Networks, Inc.

```
static {
    rule-set incoming-sip {
        from zone public;
        rule phone1 {
            match {
                destination-address 1.1.1.3/32;
            }
            then {
                static-nat prefix 10.1.1.3/32;
            }
        }
    }
}
    proxy-arp {
        interface ge-0/0/2.0 {
            address {
                1.1.1.3/32;
            }
        }
    }
```

```
[edit]
user@host# show security policies
from-zone public to-zone private {
    policy incoming {
        match {
            source-address phone2;
            destination-address phone1;
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
from-zone private to-zone public {
    policy outgoing {
        match {
            source-address phone1;
            destination-address [phone2 proxy];
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static NAT Configuration on page 287
- Verifying SIP ALG Status on page 287
- Verifying the Security Polices of SIP ALG on page 288

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule all** command.

```
user@host> show security nat static rule all
```

```
Static NAT rule: phone1 Rule-set: incoming-sip
Rule-Id : 1
Rule position : 1
From zone : trust
Source addresses : 40.10.10.0 - 40.10.10.3
: addr1
Source ports : 200 - 300
Destination addresses : 20.1.1.0
Host addresses : 3.3.3.0
Netmask : 24
Host routing-instance : N/A
Translation hits : 4
Successful sessions : 4
Failed sessions : 0
Number of sessions : 4
```

Meaning The **Translation hits** field shows that there is traffic matching the static NAT rule set.

Verifying SIP ALG Status

Purpose Verify that SIP ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
```

```
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
```

Copyright © 2016, Juniper Networks, Inc.

```
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Meaning The output shows the SIP ALG status as follows:

- ·Enabled—Shows the SIP ALG is enabled.
- ·Disabled—Shows the SIP ALG is disabled.

Verifying the Security Polices of SIP ALG

Purpose Verify that the static NAT between public zone and private zone is set.

Action From operational mode, enter the **show security policies** command.

```
user@host> show security policies
from-zone public to-zone private {
    policy incoming {
        match {
            source-address [ phone2 proxy ];
            destination-address phone1;
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
  from-zone private to-zone public {
    policy outgoing {
        match {
            source-address phone1;
            destination-address [ phone2 proxy ];
            application junos-sip;
        }
        then {
            permit;
        }
    }
}
```

Meaning The sample output shows that the static NAT between public zone and private zone is set.

Related Documentation

- Verifying SIP ALG Configurations on page 302

Example: Configuring the SIP Proxy in the Private Zone and NAT in the Public Zone

This example shows how to configure a SIP proxy server in a private zone and static NAT in a public zone to allow callers in the public zone to register with the proxy server.

- Requirements on page 289
- Overview on page 289

 Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 290
- Verification on page 293

Requirements

Before you begin, understand how NAT works with the SIP ALG. See "Understanding the SIP ALG and NAT" on page 256.

Overview

With the SIP proxy server in the private zone, you can configure static NAT on the external, or public, interface to allow callers in the public zone to register with the proxy server.

In this example (see Figure 26), phone1 and the SIP proxy server are on the ge-0/0/0 interface in the private zone, and phone2 is on the ge-0/0/2 interface in the public zone. You configure a static NAT rule for the proxy server to allow phone2 to register with the proxy server, and then create a policy called outgoing that allows SIP traffic from the public to the private zone to enable callers in the public zone to register with the proxy server. You also configure a policy called incoming from the private to the public zone to allow phone1 to call out.

Topology

Figure 26 shows configuring SIP proxy in the private zone and NAT in a public zone.

Figure 26: Configuring SIP Proxy in the Private Zone and NAT in a Public Zone

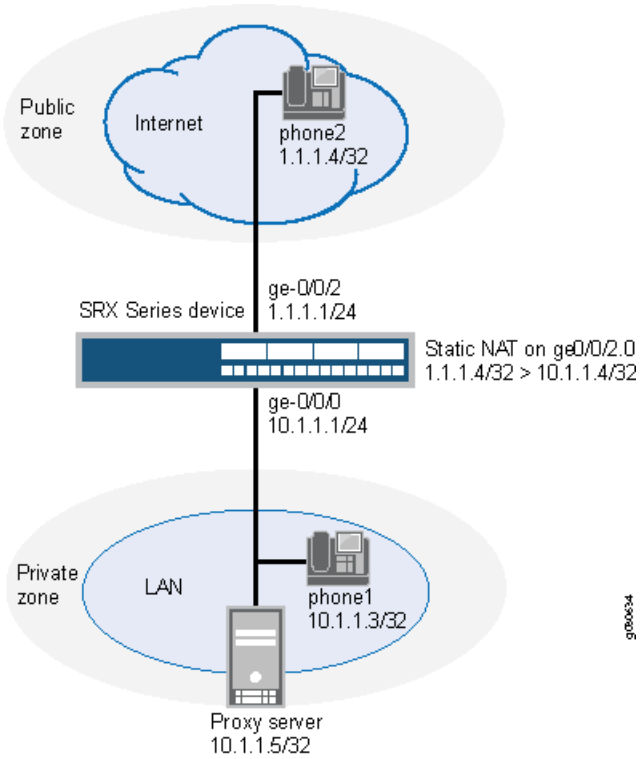

In this example, you configure NAT as follows:

- Configure static NAT on the ge-0/0/2 interface to the proxy server with a rule set called incoming-sip with a rule called proxy to match packets from the public zone with the destination address 1.1.1.4/32. For matching packets, the destination IP address is translated to the private address 10.1.1.4/32.
- Configure a second rule set called sip-phones with a rule called phone1 to enable interface NAT for communication from phone1 to phone2.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address phone1 10.1.1.3/32
set security zones security-zone private address-book address proxy 10.1.1.5/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address phone2 1.1.1.4/32
set security zones security-zone public interfaces ge-0/0/2.0
set security nat source rule-set sip-phones from zone private
set security nat source rule-set sip-phones to zone public
set security nat source rule-set sip-phones rule phone1 match source-address 10.1.1.3/32
set security nat source rule-set sip-phones rule phone1 then source-nat interface
set security nat static rule-set incoming-sip from zone public
set security nat static rule-set incoming-sip rule proxy match destination-address 1.1.1.4/32
set security nat static rule-set incoming-sip rule proxy then static-nat prefix 10.1.1.4/32
set security policies from-zone private to-zone public policy outgoing match
  source-address any
set security policies from-zone private to-zone public policy outgoing match
  destination-address phone2
set security policies from-zone private to-zone public policy outgoing match application
  junos-sip
set security policies from-zone private to-zone public policy outgoing then permit
set security policies from-zone public to-zone private policy incoming match
  source-address phone2
set security policies from-zone public to-zone private policy incoming match
  destination-address proxy
set security policies from-zone public to-zone private policy incoming match application
  junos-sip
set security policies from-zone public to-zone private policy incoming then permit
```

Step-by-Step Procedure

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

To configure static NAT for incoming calls:

1. Configure interfaces.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure security zones.

```
[edit security zones]
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone public interfaces ge-0/0/2.0
```

3. Assign addresses to the security zones.

```
[edit security zones]
user@host# set security-zone private address-book address phone1 10.1.1.3/32
user@host# set security-zone private address-book address proxy 10.1.1.5/32
user@host# set security-zone public address-book address phone2 1.1.1.4/32
```

4. Create a rule set for static NAT and assign a rule to it.

```
[edit security nat static rule-set incoming-sip]
user@host# set from zone public
user@host# set rule proxy match destination-address 1.1.1.4/32
user@host# set rule proxy then static-nat prefix 10.1.1.4/32
```

5. Configure the second rule set and assign a rule to it.

```
[edit security nat source rule-set sip-phones]
user@host# set from zone private
user@host# set to zone public
user@host# set rule phone1 match source-address 10.1.1.3/32
user@host# set rule phone1 then source-nat interface
```

6. Configure a security policy for outgoing traffic.

```
[edit security policies from-zone private to-zone public policy outgoing]
user@host# set match source-address any
user@host# set match destination-address phone2
user@host# set match application junos-sip
user@host# set then permit
```

7. Configure a security policy for incoming traffic.

```
[edit security policies from-zone public to-zone private policy incoming]
user@host# set match source-address phone2
user@host# set match destination-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
   unit 0 {
      family inet {
         address 10.1.1.1/24;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    }
    ge-0/0/2 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }
[edit]
user@host# show security zones
security-zone private {
  address-book {
    address phone1 10.1.1.3/32;
    address proxy 10.1.1.5/32;
  }
  interfaces {
    ge-0/0/0.0;
  }
}
  security-zone public {
    address-book {
      address phone2 1.1.1.4/32;
    }
    interfaces {
      ge-0/0/2.0;
    }
  }
[edit]
user@host# show security nat
source {
  rule-set sip-phones {
    from zone private;
    to zone public;
    rule phone1 {
      match {
        source-address 10.1.1.3/32;
      }
      then {
        source-nat {
          interface;
        }
      }
    }
  }
}
  static {
    rule-set incoming-sip {
      from zone public;
      rule proxy {
        match {
          destination-address 1.1.1.4/32;
        }
        then {
          static-nat prefix 10.1.1.4/32;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
[edit]
user@host# show security policies
from-zone private to-zone public {
  policy outgoing {
    match {
      source-address any;
      destination-address phone2;
      application junos-sip;
    }
    then {
      permit;
    }
  }
}
  from-zone public to-zone private {
    policy incoming {
      match {
        source-address phone2;
        destination-address proxy;
        application junos-sip;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static NAT Configuration on page 293
- Verifying SIP ALG Status on page 294

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule all** command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat source rule all
source NAT rule: phone1      Rule-set: sip-phones
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : private
  To zone                    : public
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Destination port         : 0               - 0
  Action                       : interface
```

```
    Persistent NAT type         : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout          : 0
    Max session number          : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Meaning The **Translation hits** field shows that, there is no traffic matching the source NAT rule.

Verifying SIP ALG Status

Purpose Verify that SIP ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
ALG Status :
  DNS       : Enabled
  FTP       : Enabled
  H323      : Disabled
  MGCP      : Disabled
  MSRPC     : Enabled
  PPTP      : Enabled
  RSH       : Disabled
  RTSP      : Disabled
  SCCP      : Disabled
  SIP       : Enabled
  SQL       : Enabled
  SUNRPC    : Enabled
  TALK      : Enabled
  TFTP      : Enabled
  IKE-ESP   : Disabled
```

Meaning The output shows the SIP ALG status as follows:

- Enabled—Shows the SIP ALG is enabled.
- Disabled—Shows the SIP ALG is disabled.

Related Documentation

- Verifying SIP ALG Configurations on page 302

Example: Configuring a Three-Zone SIP ALG and NAT Scenario

This example shows how to configure a SIP proxy server in a private zone and static NAT in a public zone to allow callers in the public zone to register with the proxy server.

- Requirements on page 295
- Overview on page 295
- Configuration on page 296
- Verification on page 300

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, understand how NAT works with the SIP ALG. See "Understanding the SIP ALG and NAT" on page 256.

Overview

In a three-zone SIP configuration, the SIP proxy server is typically in a different zone from the calling and called systems. Such a scenario requires additional address and zone configuration, and policies to ensure that all systems have access to each other and to the proxy server.

In this example, phone1 is on the ge-0/0/0.0 interface in the private zone, phone2 is on the ge-0/0/2.0 interface in the public zone, and the proxy server is on the ge-0/0/1.0 interface in the DMZ. You configure static NAT rule for phone1 in the private zone. You then create policies for traffic traversing from the private zone to the DMZ and from the DMZ to the private zone, from the public zone to the DMZ and from the DMZ to the public zone, and from the private zone to the public zone. The arrows in Figure 27 show the flow of SIP signaling traffic when phone2 in the public zone places a call to phone1 in the private zone. After the session is initiated, the data flows directly between phone1 and phone2.

Figure 27: Three-Zone SIP Configuration with Proxy in the DMZ

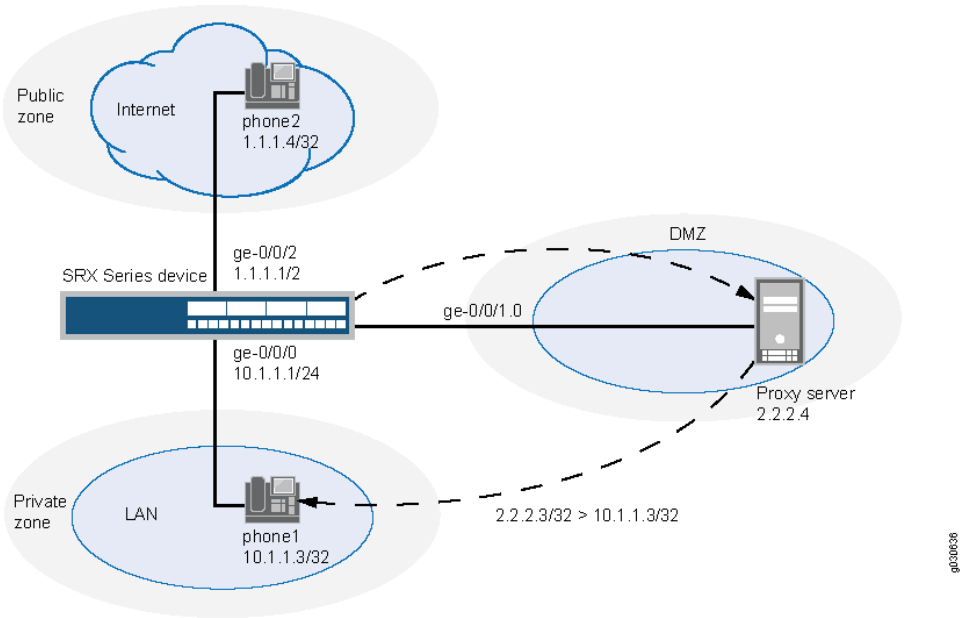

In this example, you configure NAT as follows:

- Configure a static NAT rule set called incoming-sip with a rule phone1 to match packets from the public zone with the destination address 2.2.2.3/32. For matching packets, the destination IP address is translated to the private address 10.1.1.3/32.
- Configure proxy ARP for the address 2.2.2.3/32 on interface ge-0/0/1.0 allowing the system to respond to ARP requests received on the interface for this address.

Copyright © 2016, Juniper Networks, Inc.

- Configure a second rule set called sip-phones with a rule r1 to enable interface NAT for communication from phone1 to the proxy server and from phone1 to phone2.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24
set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
set security zones security-zone private address-book address phone1 10.1.1.3/32
set security zones security-zone private interfaces ge-0/0/0.0
set security zones security-zone public address-book address phone2 1.1.1.4/32
set security zones security-zone public interfaces ge-0/0/2.0
set security zones security-zone dmz address-book address proxy 2.2.2.4/32
set security zones security-zone dmz interfaces ge-0/0/1.0
set security nat source rule-set sip-phones from zone private
set security nat source rule-set sip-phones to zone dmz
set security nat source rule-set sip-phones rule r1 match source-address 10.1.1.3/32
set security nat source rule-set sip-phones rule r1 then source-nat interface
set security policies from-zone private to-zone dmz policy private-to-proxy match
  source-address phone1
set security policies from-zone private to-zone dmz policy private-to-proxy match
  destination-address proxy
set security policies from-zone private to-zone dmz policy private-to-proxy match
  application junos-sip
set security policies from-zone private to-zone dmz policy private-to-proxy then permit
set security policies from-zone public to-zone dmz policy public-to-proxy match
  source-address phone2
set security policies from-zone public to-zone dmz policy public-to-proxy match
  destination-address proxy
set security policies from-zone public to-zone dmz policy public-to-proxy match application
  junos-sip
set security policies from-zone public to-zone dmz policy public-to-proxy then permit
set security policies from-zone private to-zone public policy private-to-public match
  source-address phone1
set security policies from-zone private to-zone public policy private-to-public match
  destination-address phone2
set security policies from-zone private to-zone public policy private-to-public match
  application junos-sip
set security policies from-zone private to-zone public policy private-to-public then permit
set security policies from-zone dmz to-zone private policy proxy-to-private match
  source-address proxy
set security policies from-zone dmz to-zone private policy proxy-to-private match
  destination-address phone1
set security policies from-zone dmz to-zone private policy proxy-to-private match
  application junos-sip
set security policies from-zone dmz to-zone private policy proxy-to-private then permit
set security policies from-zone dmz to-zone public policy proxy-to-public match
  source-address proxy
set security policies from-zone dmz to-zone public policy proxy-to-public match
  destination-address phone2
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone dmz to-zone public policy proxy-to-public match application
    junos-sip
set security policies from-zone dmz to-zone public policy proxy-to-public then permit
```

Step-by-Step Procedure

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

To configure a SIP proxy server in a private zone and static NAT in a public zone:

1. Configure interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
```

2. Configure security zones.

```
[edit security zones]
user@host# set security-zone private interfaces ge-0/0/0.0
user@host# set security-zone public interfaces ge-0/0/2.0
user@host# set security-zone dmz interfaces ge-0/0/1.0
```

3. Assign addresses to the security zones.

```
[edit security zones]
user@host# set security-zone private address-book address phone1 10.1.1.3/32
user@host# set security-zone public address-book address phone2 1.1.1.4/32
user@host# set security-zone dmz address-book address proxy 2.2.2.4/32
```

4. Configure interface NAT for communication from phone1 to proxy.

```
[edit security nat source rule-set sip-phones]
user@host# set from zone private
user@host# set to zone dmz
user@host# set rule r1 match source-address 10.1.1.3/32
user@host# set rule r1 then source-nat interface
```

5. Configure a security policy to allow traffic from zone private to zone DMZ.

```
[edit security policies from-zone private to-zone dmz policy private-to-proxy]
user@host# set match source-address phone1
user@host# set match destination-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

6. Configure a security policy to allow traffic from zone public to zone DMZ.

```
[edit security policies from-zone public to-zone dmz policy public-to-proxy]
user@host# set match source-address phone2
user@host# set match destination-address proxy
user@host# set match application junos-sip
user@host# set then permit
```

7. Configure a security policy to allow traffic from zone private to zone public.

```
[edit security policies from-zone private to-zone public policy private-to-public]
user@host# set match source-address phone1
user@host#  set match destination-address phone2
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set match application junos-sip
user@host# set then permit
```

8. Configure a security policy to allow traffic from zone DMZ to zone private.

```
[edit security policies from-zone dmz to-zone private policy proxy-to-private]
user@host# set match source-address proxy
user@host# set match destination-address phone1
user@host# set match application junos-sip
user@host# set then permit
```

9. Configure a security policy to allow traffic from zone DMZ to zone public.

```
[edit security policies from-zone dmz to-zone public policy proxy-to-public]
user@host# set match source-address proxy
user@host# set match destination-address phone2
user@host# set match application junos-sip
user@host# set then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.1.1/24;
        }
    }
}
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 2.2.2.2/24;
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            family inet {
                address 1.1.1.1/24;
            }
        }
    }
```

```
[edit]
user@host# show security zones
security-zone private {
    address-book {
        address phone1 10.1.1.3/32;
    }
    interfaces {
        ge-0/0/0.0;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    security-zone public {
      address-book {
        address phone2 1.1.1.4/32;
      }
      interfaces {
        ge-0/0/2.0;
      }
    }
    security-zone dmz {
      address-book {
        address proxy 2.2.2.4/32;
      }
      interfaces {
        ge-0/0/1.0;
      }
    }
[edit]
user@host# show security nat
source {
  rule-set sip-phones {
    from zone private;
    to zone dmz;
    rule r1 {
      match {
        source-address 10.1.1.3/32;
      }
      then {
        source-nat {
          interface;
        }
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/1.0 {
      address {
        2.2.2.3/32;
      }
    }
  }

[edit]
user@host# show security policies
from-zone private to-zone dmz {
  policy private-to-proxy {
    match {
      source-address phone1;
      destination-address proxy;
      application junos-sip;
    }
    then {
      permit;
    }
  }
```

```
}
  from-zone public to-zone dmz {
    policy public-to-proxy {
      match {
        source-address phone2;
        destination-address proxy;
        application junos-sip;
      }
      then {
        permit;
      }
    }
  }
  from-zone private to-zone public {
    policy private-to-public {
      match {
        source-address phone1;
        destination-address phone2;
        application junos-sip;
      }
      then {
        permit;
      }
    }
  }
  from-zone dmz to-zone private {
    policy proxy-to-private {
      match {
        source-address proxy;
        destination-address phone2;
        application junos-sip;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Rule Usage on page 300
- Verifying Static NAT Configuration on page 301
- Verifying SIP ALG Status on page 301

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat source rule all
source NAT rule: r1      Rule-set: sip-phones
  Rule-Id                    : 1
  Rule position              : 1
  From zone                  : private
  To zone                    : public
  Match
    Source addresses         : 0.0.0.0         - 255.255.255.255
    Destination port         : 0               - 0
  Action                       : interface
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule r1** command. View the Translation hits field to check for traffic that matches the rule.

Verifying SIP ALG Status

Purpose Verify that SIP ALG is enabled on your system.

Action From operational mode, enter the **show security alg status** command.

```
user@host> show security alg status
ALG Status :
  DNS      : Enabled
  FTP      : Enabled
  H323     : Disabled
  MGCP     : Disabled
  MSRPC    : Enabled
  PPTP     : Enabled
  RSH      : Disabled
  RTSP     : Disabled
  SCCP     : Disabled
  SIP      : Enabled
  SQL      : Enabled
  SUNRPC   : Enabled
  TALK     : Enabled
  TFTP     : Enabled
  IKE-ESP  : Disabled
```

Related Documentation

- Verifying SIP ALG Configurations on page 302

Copyright © 2016, Juniper Networks, Inc.

Verifying SIP ALG Configurations

- Verifying SIP ALG on page 302
- Verifying SIP ALG Calls on page 302
- Verifying SIP ALG Call Details on page 302
- Verifying SIP ALG Counters on page 303
- Verifying the Rate of SIP ALG Messages on page 304

Verifying SIP ALG

Purpose Verify SIP ALG verification options.

Action From the CLI, enter the **show security alg sip ?** command.

```
user@host> show security alg sip ?
Possible completions:
  calls                Show SIP calls
  counters             Show SIP counters
  rate                 Show SIP rate
```

Meaning The output shows a list of all SIP verification parameters. Verify the following information:

- Calls—Lists all SIP calls.
- Counters—Provides counters of response codes for each SIP request method and error type.
- Rate—Provides speed and periodicity of SIP signaling messages.

Verifying SIP ALG Calls

Purpose Display information about active calls.

Action From the J-Web interface, select **Monitor>ALGs>SIP>Calls**. Alternatively, from the CLI, enter the **show security alg sip calls** command.

```
user@host> show security alg sip calls
Total number of calls: 1
       Call ID: 47090a32@30.2.20.5
   Method: INVITE
```

Meaning The output shows a list of all active SIP calls. Verify the User Agent Server (UAS) call ID and local and remote tags, and the state of the call.

Verifying SIP ALG Call Details

Purpose Display address and SDP about active calls.

Copyright © 2016, Juniper Networks, Inc.

Action From the J-Web interface, select **Monitor>ALGs>SIP>Details**. Alternatively, from the CLI, enter the **show security alg sip calls detail** command.

```
user@host> show security alg sip calls detail
Total number of calls: 1
       Call ID    : 47090a32@30.2.20.5
Method   : INVITE
  State    : SETUP
  Group ID : 24575
```

Meaning The output provides details about all active SIP calls. Verify the following information:

- The total number of calls, their ID and tag information, and state
- Remote group ID
- The IP addresses and port numbers and SDP connection and media details

Verifying SIP ALG Counters

Purpose Display information about SIP counters.

Action From the J-Web interface, select **Monitor>ALGs>SIP>Counters**. Alternatively, from the CLI, enter the **show security alg sip counters** command.

```
user@host> show security alg sip counters
Method          T        1xx        2xx        3xx        4xx        5xx        6xx
                RT         RT         RT         RT         RT         RT         RT
INVITE          4          4          3          0          0          0          0
                0          0          0          0          0          0          0
CANCEL          0          0          0          0          0          0          0
                0          0          0          0          0          0          0
ACK             3          0          0          0          0          0          0
                0          0          0          0          0          0          0
BYE             3          0          3          0          0          0          0
                0          0          0          0          0          0          0
REGISTER        7          0          7          0          0          0          0
                0          0          0          0          0          0          0
OPTIONS         0          0          0          0          0          0          0
                0          0          0          0          0          0          0
INFO            0          0          0          0          0          0          0
                0          0          0          0          0          0          0
MESSAGE         0          0          0          0          0          0          0
                0          0          0          0          0          0          0
NOTIFY          0          0          0          0          0          0          0
                0          0          0          0          0          0          0
PRACK           0          0          0          0          0          0          0
                0          0          0          0          0          0          0
PUBLISH         0          0          0          0          0          0          0
                0          0          0          0          0          0          0
REFER           0          0          0          0          0          0          0
                0          0          0          0          0          0          0
SUBSCRIBE       0          0          0          0          0          0          0
                0          0          0          0          0          0          0
UPDATE          0          0          0          0          0          0          0
                0          0          0          0          0          0          0
BENOTIFY        0          0          0          0          0          0          0
                0          0          0          0          0          0          0
```

```
SERVICE          0          0          0          0          0          0          0
                 0          0          0          0          0          0          0
OTHER            0          0          0          0          0          0          0
                 0          0          0          0          0          0          0
    SIP Error Counters
    ----------------------------
       Total Pkt-in                        :34
       Total Pkt dropped on error          :0
              Call error                            :0
       IP resolve error                    :0
       NAT error                           :0
       Resource manager error              :0
       RR header exceeded max              :0
       Contact header exceeded max         :0
       Call Dropped due to limit           :0
       SIP stack error                     : 0
       SIP decode error                    : 0
      SIP unknown method error             : 0
       RTO message sent                    : 0
       RTO message received                 : 0
       RTO buffer allocation failure        : 0
       RTO buffer transmit failure          : 0
       RTO send processing error            : 0
        RTO receive processing error         : 0
       RTO receive invalid length           : 0
        RTO receive call process error        : 0
        RTO receive call allocation error     : 0
        RTO receive call register error       : 0
        RTO receive invalid status error      : 0
```

Meaning The output provides a count of all SIP response codes transmitted and received, and of SIP errors. Verify the following information:

- A count of transmissions of response codes for each SIP request method
- A count of all possible error types

Verifying the Rate of SIP ALG Messages

Purpose Display information about SIP message rate.

Action From the J-Web interface, select **Monitor>ALGs>SIP>Rate**. Alternatively, from the CLI, enter the **show security alg sip rate** command.

```
user@host> show security alg sip rate
CPU ticks per microseconds is 3735928559
    Time taken for the last message is 0 microseconds
    Total time taken for 0 messages is 0 microseconds(in less than 10 minutes)
    Rate: 3735928559 messages/second
```

Meaning The output provides information about CPU usage for messages, and speed and periodicity of SIP signaling messages. Verify the following information:

- CPU ticks per US
- Passage time for last message, for all messages, and the rate at which messages transit the network

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- SIP ALG Configuration Overview on page 248
- Example: Configuring Interface Source NAT for Incoming SIP Calls on page 267
- Decreasing Network Complexity by Configuring a Source NAT Pool for Incoming SIP Calls on page 273
- Example: Configuring Static NAT for Incoming SIP Calls on page 281
- Example: Configuring the SIP Proxy in the Private Zone and NAT in the Public Zone on page 288
- Example: Configuring a Three-Zone SIP ALG and NAT Scenario on page 294

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 4

Configuration Statements and Operational Commands

- Configuration Statements on page 309
- Operational Commands on page 415

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 19

Configuration Statements

- Applications Configuration Statement Hierarchy on page 311
- [edit security alg] Hierarchy Level on page 312
- [edit security nat] Hierarchy Level on page 316
- [edit security policies] Hierarchy Level on page 320
- [edit security zones] Hierarchy Level on page 324
- [edit security traceoptions] Hierarchy Level on page 325
- address (Security Destination NAT) on page 326
- alg on page 327
- alg (Applications) on page 332
- alg-manager on page 333
- allow-dns-reply on page 333
- application (Security Policies) on page 334
- application-protocol (Applications) on page 335
- application-screen (Security H323) on page 336
- application-screen (Security MGCP) on page 337
- application-screen (Security SCCP) on page 338
- application-screen (Security SIP) on page 339
- banner (Access FTP HTTP Telnet Authentication) on page 340
- call-flood on page 340
- c-timeout on page 341
- deny (Security SIP) on page 342
- destination-address (Security Destination NAT) on page 343
- destination-address (Security Policies) on page 344
- destination-address (Security Source NAT) on page 345
- destination-address (Security Static NAT) on page 345
- destination-nat on page 346
- destination-port (Applications) on page 347
- dns on page 351

Copyright © 2016, Juniper Networks, Inc.

- dns (System Services) on page 352
- dscp-rewrite on page 353
- endpoint-registration-timeout on page 354
- family inet (Interfaces) on page 355
- ftp (Access) on page 357
- ftp (Security ALG) on page 358
- gatekeeper on page 359
- h323 on page 360
- host-inbound-traffic on page 361
- ike-esp-nat on page 362
- ike (Security) on page 363
- inactive-media-timeout (Security MGCP) on page 365
- inactive-media-timeout (Security SCCP) on page 365
- inactive-media-timeout (Security SIP) on page 366
- map-entry-timeout on page 366
- maximum-call-duration (Security) on page 367
- maximum-message-length on page 367
- media-source-port-any on page 368
- message-flood (Security H323) on page 368
- mgcp on page 369
- msrpc on page 370
- nat on page 371
- nat-pat-address on page 375
- policy (Security Policies) on page 376
- pptp on page 378
- protect on page 379
- protocols (Security Zones Interfaces) on page 380
- retain-hold-resource on page 381
- rsh on page 382
- rtsp on page 383
- sccp on page 384
- security-zone on page 385
- sip (Security) on page 387
- source-address (Security Destination NAT) on page 388
- source-address (Security Policies) on page 389
- source-nat on page 390
- sql on page 391

Copyright © 2016, Juniper Networks, Inc.

- static-nat on page 392
- sunrpc on page 393
- support-lib on page 394
- system-services (Security Zones Host Inbound Traffic) on page 395
- t1-interval on page 397
- t4-interval on page 397
- talk on page 398
- term (Applications) on page 399
- tftp (Security ALG) on page 400
- traceoptions (Security ALG) on page 401
- traceoptions (Security H323 ALG) on page 403
- traceoptions (Security MGCP ALG) on page 404
- traceoptions (Security SCCP ALG) on page 405
- traceoptions (Security SIP ALG) on page 406
- traceoptions (System Services DNS) on page 407
- transaction-timeout on page 409
- unknown-message (Security H323 ALG) on page 410
- unknown-message (Security MGCP ALG) on page 411
- unknown-message (Security SCCP ALG) on page 412
- unknown-message (Security SIP ALG) on page 413

Applications Configuration Statement Hierarchy

Use the statements in the **applications** configuration hierarchy to configure applications properties and group applications objects.

```
applications {
  application application-name {
    application-protocol (dns | ftp |gprs-gtp-c| gprs-gtp-u | gprs-gtp-v0 |gprs-sctp | http
      | ignore | ike-esp-nat | mgcp-ca | mgcp-ua | ms-rpc | q931 | ras | realaudio | rsh | rtsp
      | sccp | sip | sqlnet-v2 | sun-rpc | talk | tftp);
    description text;
    destination-port port-identifier;
    do-not-translate-A-query-to-AAAA-query;
    do-not-translate-AAAA-query-to-A-query;
    ether-type hex-value;
    icmp-code value;
    icmp-type value;
    icmp6-code value;
    icmp6-type value;
    inactivity-timeout (seconds | never);
    protocol number;
    rpc-program-number number;
    source-port port-number;
    term term-name {
      alg application;
```

Copyright © 2016, Juniper Networks, Inc.

```
        destination-port port-identifier;
        icmp-code value;
        icmp-type value;
        icmp6-code value;
        icmp6-type value;
        inactivity-timeout (seconds | never);
        protocol number;
        rpc-program-number number;
        source-port port-number;
        uuid hex-value;
      }
      uuid hex-value;
    }
    application-set application-set-name {
      application application-name;
      application-set application-set-name;
      description text;
    }
  }
```

Related Documentation

- application (Applications) on page 1224

[edit security alg] Hierarchy Level

```
security {
  alg {
    alg-manager {
      traceoptions {
        flag {
          all <extensive>;
        }
      }
    }
    alg-support-lib {
      traceoptions {
        flag {
          all <extensive>;
        }
      }
    }
    dns {
      disable;
      doctoring (none | sanity-check);
      maximum-message-length number;
      traceoptions {
        flag {
          all <extensive>;
        }
      }
    }
    ftp {
      allow-mismatch-ip-address;
      disable;
      ftps-extension;
```


Copyright © 2016, Juniper Networks, Inc.

```
    line-break-extension;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  h323 {
    application-screen {
      message-flood {
        gatekeeper {
          threshold rate;
        }
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
    endpoint-registration-timeout value-in-seconds;
    media-source-port-any;
    traceoptions {
      flag flag <detail | extensive | terse>;
    }
  }
  ike-esp-nat {
    enable;
    esp-gate-timeout value-in-seconds;
    esp-session-timeout value-in-seconds;
    state-timeout value-in-seconds;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  mgcp {
    application-screen {
      connection-flood {
        threshold rate;
      }
      message-flood {
        threshold rate;
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
```

```
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    traceoptions {
      flag flag <extensive>;
    }
    transaction-timeout value-in-seconds;
  }
  msrpc {
    disable;
    map-entry-timeout map-entry-timeout;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  pptp {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  rsh {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  rtsp {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  sccp {
    application-screen {
      call-flood {
        threshold rate;
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
    inactive-media-timeout value-in-seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
    traceoptions {
      flag flag <extensive>;
    }
  }
  sip {
    application-screen {
      protect {
        deny {
          all {
            timeout value-in-seconds;
          }
          destination-ip address;
          timeout value-in-seconds;
        }
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    c-timeout value-in-minutes;
    disable;
    dscp-rewrite {
      code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    retain-hold-resource;
    t1-interval value-in-milliseconds;
    t4-interval value-in-seconds;
    traceoptions {
      flag flag <detail | extensive | terse>;
    }
  }
  sql {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  sunrpc {
    disable;
    map-entry-timeout map-entry-timeout;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  talk {
    disable;
    traceoptions {
      flag {
        all <extensive>;
```

```
        }
      }
    }
    tftp {
      disable;
      traceoptions {
        flag {
          all <extensive>;
        }
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      level (brief | detail | extensive | verbose);
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security nat] Hierarchy Level

```
security {
  nat {
    destination {
      pool pool-name {
        address <ip-address> {
          (port port-number | to ip-address);
        }
        description text;
        routing-instance (routing-instance-name | default);
      }
      rule-set rule-set-name {
        description text;
        from {
          interface [interface-name];
          routing-instance [routing-instance-name];
          zone [zone-name];
        }
        rule rule-name {
          description text;
          match {
            application {
              [application];
              any;
            }
            (destination-address ip-address | destination-address-name address-name);
```


Copyright © 2016, Juniper Networks, Inc.

```
            destination-port (port-or-low <to high>);
            protocol [protocol-name-or-number];
            source-address [ip-address];
            source-address-name [address-name];
          }
          then {
            destination-nat (off | pool pool-name | rule-session-count-alarm
              (clear-threshold value | raise-threshold value));
          }
        }
      }
    }
    proxy-arp interface interface-name address ip-address;
      to ip-address;
    }
    proxy-ndp interface interface-name address ip-address;
      to ip-address;
    }
    source {
      address-persistent;
      interface (port-overloading off | port-overloading-factor number);
      pool pool-name {
        address ip-address {
          to ip-address;
        }
        address-persistent subscriber ipv6-prefix-length prefix-length;
        address-pooling (paired | no-paired);
        address-shared;
        description text;
        host-address-base ip-address;
        overflow-pool (pool-name | interface);
        pool-utilization-alarm (clear-threshold value | raise-threshold value);
        port {
          block-allocation {
            active-block-timeouttimeout-interval;
            block-size block-size;
            log disable;
            maximum-blocks-per-host maximum-block-number
          }
          deterministic {
            block-size block-size;
            host {
              address ip-address;
              address-name address-name;
          }
            no-translation;
            port-overloading-factor number;
            range {
              port-low <to port-high>;
              to port-high;
              twin-port port-low <to port-high>;
            }
        }
        routing-instance routing-instance-name;
      }
      pool-default-port-range lower-port-range to upper-port-range;
```

```
    pool-default-twin-port-range lower-port-range to upper-port-range;
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port-randomization disable;
    port-round-robin disable;
    rule-set rule-set-name {
      description text;
      from {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
      rule rule-name {
        description text;
        match {
          application {
            [application];
            any;
          }
          (destination-address <ip-address> | destination-address-name
            <address-name>);
          destination-port (port-or-low <to high>);
          protocol [protocol-name-or-number];
          source-address [ip-address];
          source-address-name [address-name];
          source-port (port-or-low <to high>);
        }
        then source-nat;
          interface {
            persistent-nat {
              address-mapping;
              inactivity-timeout seconds;
              max-session-number value;
              permit (any-remote-host | target-host | target-host-port);
            }
          off;
          pool <pool-name>
            persistent-nat
              address-mapping;
              inactivity-timeout seconds;
              max-session-number number;
              permit (any-remote-host | target-host | target-host-port);
            }
          rule-session-count-alarm (clear-threshold value | raise-threshold value);
        }
      }
      to {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
    }
  }
  static rule-set rule-set-name;
      description text;
      from {
        interface [interface-name];
```


Copyright © 2016, Juniper Networks, Inc.

```
            routing-instance [routing-instance-name];
            zone [zone-name];
          }
          rule rule-name {
            description text;
            match {
              (destination-address <ip-address> | destination-address-name
                <address-name>);
              destination-port (port-or-low  | <to high>);
              source-address [ip-address];
              source-address-name [address-name];
              source-port (port-or-low  <to high>);
            }
            then static-nat;
              inet {
                routing-instance (routing-instance-name | default);
              }
              nptv6-prefix {
                address-prefix;
                  routing-instance routing-instance-name;
              nptv6-prefix-name {
                address-prefix-name;
                  routing-instance routing-instance-name;
              [
              prefix {
                address-prefix;
                mapped-port lower-port-range to upper-port-range;
                routing-instance (routing-instance-name| default);
              }
              prefix-name {
                address-prefix-name;
                mapped-port lower-port-range to upper-port-range;
                routing-instance (routing-instance-name | default);
              }
              rule-session-count-alarm (clear-threshold value | raise-threshold value);
              }
            }
          }
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (world-readable | no-world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Introduction to NAT on page 4839

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          destination-address-excluded;
          source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          source-address-excluded;
          source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
          }
        }
        scheduler-name scheduler-name;
        then {
          count {
            alarm {
              per-minute-threshold number;
              per-second-threshold number;
            }
          }
          deny;
          log {
            session-close;
            session-init;
          }
```


Copyright © 2016, Juniper Networks, Inc.

```
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

```
global {
   policy policy-name {
      description description;
      match {
         application {
            [application];
            any;
         }
         destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
         }
         from-zone {
            [zone-name];
            any;
         }
         source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
         }
         source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
         }
         to-zone {
            [zone-name];
            any;
         }
      }
      scheduler-name scheduler-name;
      then {
         count {
            alarm {
               per-minute-threshold number;
               per-second-threshold number;
            }
         }
         deny;
         log {
            session-close;
            session-init;
         }
         permit {
            application-services {
               application-firewall {
                  rule-set rule-set-name;
               }
               application-traffic-control {
                  rule-set rule-set-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
        }
        reject;
      }
    }
  }
  policy-rematch;
  policy-stats {
    system-wide (disable | enable);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
```

```
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

[edit security zones] Hierarchy Level

```
security {
  zones {
    functional-zone {
      management {
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
        interfaces interface-name  {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
      }
    }
    security-zone zone-name {
      address-book {
        address address-name {
          ip-prefix {
            description text;
          }
          description text;
          dns-name domain-name {
            ipv4-only;
            ipv6-only;
          }
          range-address lower-limit to upper-limit;
          wildcard-address ipv4-address/wildcard-mask;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        address-set address-set-name {
          address address-name;
          address-set address-set-name;
          description text;
        }
      }
      application-tracking;
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
      interfaces interface-name {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
      tcp-rst;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Security Zones and Interfaces Overview on page 1049

[edit security traceoptions] Hierarchy Level

```
security {
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
      size maximum-file-size;
    }
    flag flag;
    no-remote-trace;
    rate-limit messages-per-second;
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

address (Security Destination NAT)

Syntax

```
address <ip-address> {
   (port port-number | to ip-address);
}
```

Hierarchy Level [edit security nat destination pool *pool-name*]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a single address or an address range of the destination NAT pool.

Options
- ***ip-address*** —IP address of a pool.
- **port** ***port-number***—Specify the port number.
- **to**—Specify the upper limit of the address range.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

alg

Syntax

```
alg {
  alg-manager {
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  alg-support-lib {
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  dns {
    disable;
    doctoring (none | sanity-check);
    maximum-message-length number;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  ftp {
    allow-mismatch-ip-address;
    disable;
    ftps-extension;
    line-break-extension;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  h323 {
    application-screen {
      message-flood {
        gatekeeper {
          threshold rate;
        }
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
```

```
    endpoint-registration-timeout value-in-seconds;
    media-source-port-any;
    traceoptions {
      flag flag <detail | extensive | terse>;
    }
  }
  ike-esp-nat {
    enable;
    esp-gate-timeout value-in-seconds;
    esp-session-timeout value-in-seconds;
    state-timeout value-in-seconds;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  mgcp {
    application-screen {
      connection-flood {
        threshold rate;
      }
      message-flood {
        threshold rate;
      }
      unknown-message {
        permit-nat-applied;
        permit-routed;
      }
    }
    disable;
    dscp-rewrite {
      code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    traceoptions {
      flag flag <extensive>;
    }
    transaction-timeout value-in-seconds;
  }
  msrpc {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
  }
  pptp {
    disable;
    traceoptions {
      flag {
        all <extensive>;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
real {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
rsh {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
rtsp {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
sccp {
    application-screen {
        call-flood {
            threshold rate;
        }
        unknown-message {
            permit-nat-applied;
            permit-routed;
        }
    }
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    traceoptions {
        flag flag <extensive>;
    }
}
sip {
    application-screen {
        protect {
            deny {
                all {
                    timeout value-in-seconds;
                }
                destination-ip address;
                timeout value-in-seconds;
            }
        }
        unknown-message {
```

```
            permit-nat-applied;
            permit-routed;
        }
    }
    c-timeout value-in-minutes;
    disable;
    dscp-rewrite {
        code-point string;
    }
    inactive-media-timeout value-in-seconds;
    maximum-call-duration value-in-minutes;
    retain-hold-resource;
    t1-interval value-in-milliseconds;
    t4-interval value-in-seconds;
    traceoptions {
        flag flag <detail | extensive | terse>;
    }
}
sql {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
sunrpc {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
talk {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
tftp {
    disable;
    traceoptions {
        flag {
            all <extensive>;
        }
    }
}
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
```


Copyright © 2016, Juniper Networks, Inc.

```
        size maximum-file-size;
      }
      level (brief | detail | extensive | verbose);
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure an Application Layer Gateway (ALG) on the device. An ALG runs as a service and can be associated in policies with specified types of traffic. ALGs are enabled by default.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *Network Monitoring and Troubleshooting Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

alg (Applications)

Syntax alg *application*;

Hierarchy Level [edit applications application *application-name* <term *term-name*>]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define individual Application Layer Gateway (ALG).

Options *application* —Name of the application. The following protocols are supported:

- **dns** —Domain Name Service
- **ftp** —File Transfer Protocol
- **http** —Hypertext Transfer Protocol
- **https** —Hypertext Transfer Protocol
- **ignore**—Ignore application type
- **ike-esp-nat** —IKE ESP NAT application protocol
- **mgcp-ca** —Media Gateway Control Protocol with Call Agent
- **mgcp-ua** —MGCP with User Agent
- **ms-rpc** —Microsoft RPC
- **pptp** —Point-to-Point Tunneling Protocol
- **q931** —ISDN connection control protocol (Q.931)
- **ras** —Remote Access Service
- **realaudio** —RealAudio
- **rsh** —UNIX remote shell services
- **rtsp** —Real-Time Streaming Protocol
- **sccp** —Skinny Client Control Protocol
- **sip**—Session Initiation Protocol
- **smtp** —Simple Mail Transfer Protocol
- **sqlnet-v2** —Oracle SQLNET v2
- **ssh** —Secure Shell Protocol
- **sun-rpc** —Sun Microsystems RPC
- **talk**—TALK program
- **telnet** —Telnet Protocol
- **tftp** —Trivial File Transfer Protocol

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- term (Applications) on page 399

alg-manager

Syntax

```
alg-manager {
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level
[edit security alg]

Description
Configure the Application Layer Gateway (ALG) manager.

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- traceoptions (Security ALG) on page 401

allow-dns-reply

Syntax

```
allow-dns-reply;
```

Hierarchy Level
[edit security flow]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Allow an incoming Domain Name Service (DNS) reply packet without a matched request. By default, if an incoming UDP first-packet has dst-port 53, the device checks the DNS message packet header to verify that the query bit (QR) is 0, which denotes a query message. If the QR bit is 1, which denotes a response message, the device drops the packet, does not create a session, and increments the illegal packet flow counter for the interface. Using the **allow-dns-reply** statement directs the device to skip the check.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

application (Security Policies)

Syntax

```
application {
  [application];
  any;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP or remote procedure call (RPC) application or set of applications to be used as match criteria.

Options

application-name-or-set—Name of the predefined or custom application or application set used as match criteria.

any—Any predefined or custom applications or application sets.

NOTE: A custom application that does not use a well-known destination port for the application will not be included in the any option, and must be named explicitly.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application-protocol (Applications)

Syntax application-protocol (dns | ftp | http | https | ignore | ike-esp-nat | mgcp-ca | mgcp-ua | ms-rpc | pptp | q931 | ras | realaudio | rsh | rtsp | sccp | sip | smtp | sqlnet-v2 | ssh | sun-rpc | talk | telnet | tftp);

Hierarchy Level [edit applications **application** *application-name*]

Release Information Statement modified in Junos OS Release 8.5. The **ike-esp-nat** option introduced in Junos OS Release 10.2.

Description Identify the application protocol name. The following protocols are supported:

- **dns** —Domain Name Service
- **ftp** —File Transfer Protocol
- **http** —Hypertext Transfer Protocol
- **https** —Hypertext Transfer Protocol
- **ignore** —Ignore application type
- **ike-esp-nat** —IKE ESP NAT application protocol
- **mgcp-ca** —Media Gateway Control Protocol with Call Agent
- **mgcp-ua** —MGCP with User Agent
- **ms-rpc** —Microsoft RPC
- **pptp** —Point-to-Point Tunneling Protocol
- **q931** —ISDN connection control protocol (Q.931)
- **ras** —Remote Access Service
- **realaudio** —RealAudio
- **rsh** —UNIX remote shell services
- **rtsp**—Real-Time Streaming Protocol
- **sccp** —Skinny Client Control Protocol
- **sip**—Session Initiation Protocol
- **smtp** —Simple Mail Transfer Protocol
- **sqlnet-v2** —Oracle SQLNET v2
- **ssh** —Secure Shell Protocol
- **sun-rpc** —Sun Microsystems RPC
- **talk**—TALK program
- **telnet** —Telnet Protocol
- **tftp** —Trivial File Transfer Protocol

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

application-screen (Security H323)

Syntax

```
application-screen {
  message-flood {
    gatekeeper {
      threshold rate;
    }
  }
  unknown-message {
    permit-nat-applied;
    permit-routed;
  }
}
```

Hierarchy Level
[edit security alg h323]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Configure the security screens for the H.323 protocol Application Layer Gateway (ALG).

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- h323 on page 360

Copyright © 2016, Juniper Networks, Inc.

application-screen (Security MGCP)

Syntax

```
application-screen {
  connection-flood {
    threshold rate;
  }
  message-flood {
    threshold rate;
  }
  unknown-message {
    permit-nat-applied;
    permit-routed;
  }
}
```

Hierarchy Level [edit security alg mgcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the security screens for the Media Gateway Control Protocol (MGCP) Application Layer Gateway (ALG).

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369

Copyright © 2016, Juniper Networks, Inc.

application-screen (Security SCCP)

Syntax

```
application-screen {
  call-flood {
    threshold rate;
  }
  unknown-message {
    permit-nat-applied;
    permit-routed;
  }
}
```

Hierarchy Level [edit security alg sccp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the security screens for the Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG).

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sccp on page 384

Copyright © 2016, Juniper Networks, Inc.

application-screen (Security SIP)

Syntax

```
application-screen {
  protect {
    deny {
      all {
        timeout value-in-seconds;
      }
      destination-ip address;
      timeout value-in-seconds;
    }
  }
  unknown-message {
    permit-nat-applied;
    permit-routed;
  }
}
```

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the security screens for the Session Initiation Protocol (SIP) Application Layer Gateway (ALG).

NOTE: IPv6 is supported on the SIP ALG along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation.

The type of the <destination-ip-address> is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

banner (Access FTP HTTP Telnet Authentication)

Syntax

```
banner {
  fail string;
  login string;
  success string;
}
```

Hierarchy Level [edit access firewall-authentication pass-through (ftp | http | telnet)]

Release Information Statement introduced in Junos OS Release 8.5.
HTTPS for web-authentication is supported on high-end SRX Series Services Gateways from Junos OS Release 12.1X44-D10 and on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways from Junos OS Release 15.1X49-D40.

Description Configure the banners that appear to users during the FTP, HTTP, HTTPS, and Telnet firewall authentication process. The banners appear during login, after successful authentication, and after failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Firewall Authentication Banner Customization on page 5251

call-flood

Syntax

```
call-flood {
  threshold rate;
}
```

Hierarchy Level [edit security alg sccp application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Limit the number of calls per second allowed to Skinny Client Control Protocol (SCCP) client. Calls exceeding the threshold are dropped by the SCCP Application Layer Gateway (ALG).

Options
threshold *rate*—Number of calls per second per client.
Range: 2 through 1000
Default: 20

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- sccp on page 384

Copyright © 2016, Juniper Networks, Inc.

c-timeout

Syntax c-timeout *value-in-minutes*;

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the timeout interval for Session Initiation Protocol (SIP) transactions in minutes.

Options ***value-in-minutes***—Timeout interval.
Range: 3 through 10 minutes
Default: 3 minutes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- sip (Security) on page 387

Copyright © 2016, Juniper Networks, Inc.

deny (Security SIP)

Syntax

```
deny {
  all {
    timeout value-in-seconds;
  }
  destination-ip address;
  timeout value-in-seconds;
}
```

Hierarchy Level [edit security alg sip application-screen protect]

Release Information Statement introduced in Junos OS Release 8.5.

Description Protect servers against INVITE attacks.

Options

- **all**—Configure the Session Initiation Protocol (SIP) application screen to protect servers at all destination IP addresses against INVITE attacks.
- **destination-ip *address***—Configure the SIP application screen to protect the server at this destination IP address against INVITE attacks. You can include up to 16 destination IP addresses of servers to be protected. Enabling this option disables the all option.
- **timeout *value-in-seconds***—Amount of time (in ***seconds***) to make an attack table entry for each INVITE, which is listed in the application screen.

  Range: 1 through 3600 seconds

  Default: 5 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

 Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Destination NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure one address or a subnet.

NOTE:

- If the destination address is IPv4 and the pool is an IPv6 prefix, the length of the IPv6 prefix must be 96.
- If the destination address is an IPv6 prefix and the pool is an IPv6 prefix, their length must be the same.

Options ***ip-address***— Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Policies)

Syntax

```
destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options

address—IP address (**any**, **any-ipv4**, **any-ipv6**), IP address set, or address book entry, or wildcard address (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Source NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure multiple addresses or subnets.

Options ***ip-address***—Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

destination-address (Security Static NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure one address or a subnet.

Options ***ip-address***—Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-nat

Syntax destination-nat (off | pool *pool-name* | rule-session-count-alarm (clear-threshold *value* | raise-threshold *value*));

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. The **rule-session-count-alarm** option added in Junos OS Release 12.1X45-D10.

Description Specify the action of the destination NAT rule.

Options

off—Do not perform destination NAT operation.

pool—Use user-defined destination NAT pool to perform destination NAT.

rule-session-count-alarm—Define session count alarm thresholds for a specific destination NAT rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

NOTE: If you enter a value for **raise-threshold** but not for **clear-threshold**, **clear-threshold** is automatically set to 80 percent of **raise-threshold**.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-port (Applications)

Syntax destination-port *port-identifier*;

Hierarchy Level [edit applications application *application-name*],
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify a TCP or UDP destination port number.

Options ***port-identifier*** —Range of ports. You can use a numeric value or one of the text synonyms listed in Table 14.

Copyright © 2016, Juniper Networks, Inc.

Table 14: Port Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760

Copyright © 2016, Juniper Networks, Inc.

Table 14: Port Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162

Copyright © 2016, Juniper Networks, Inc.

Table 14: Port Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

dns

Syntax

```
dns {
  disable;
  doctoring (none | sanity-check);
  maximum-message-length number;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Domain Name Service (DNS) Application Layer Gateway (ALG) on the device.

Options

- **disable**—Disable the DNS ALG. By default, the DNS ALG is enabled. This option will enable or disable DNS ALG for both IPV4 and IPV6 mode.
- **doctoring**—Configure DNS ALG doctoring.
  - **none**— Disable all DNS ALG Doctoring.
  - **sanity-check**—Perform only DNS ALG sanity checks.
- **maximum-message-length**—A limit imposed on the size of individual DNS messages (see related section).
- **traceoptions**—Configure SQL ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

dns (System Services)

Syntax

```
dns {
    dns-proxy {
      cache hostname inet ip-address;
      default-domain domain-name {
        forwarders ip-address;
      }
      interface interface-name;
      propogate-setting (enable | disable);
      view view-name {
        domain domain-name {
          forward-only;
          forwarders ip-address;
        }
        match-clients subnet-address;
      }
    }
  }
  dnssec {
    disable;
    dlv {
      domain-name domain-name trusted-anchor  trusted-anchor;
      }
  secure-domains domain-name;
    trusted-keys (key dns-key | load-key-file url);
  forwarders {
    ip-address;
  }
  max-cache-ttl seconds;
  max-ncache-ttl seconds;
  traceoptions {
    category {
      category-type;
    }
    debug-level level;
    file {
      filename;
      files number;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the DNS server.

Copyright © 2016, Juniper Networks, Inc.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *DNS Overview*

dscp-rewrite

Syntax

```
dscp-rewrite {
  code-point 6 bit patterns;
}
```

Hierarchy Level
[edit security alg sip]
[edit security alg h323]
[edit security alg mgcp]
[edit security alg sccp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify a rewrite-rule for the traffic that passes through a voice over IP Application Layer Gateway (VoIP ALG).

Options **6-bit patterns**—Value of the code point in binary form.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding VoIP DSCP Rewrite Rules on page 143

Copyright © 2016, Juniper Networks, Inc.

endpoint-registration-timeout

Syntax endpoint-registration-timeout *value-in-seconds*;

Hierarchy Level [edit security alg h323]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the timeout value in seconds for entries in the NAT table.

Options ***value-in-seconds***—Timeout value.
Range: 10 through 50,000 seconds
Default: 3600 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- h323 on page 360

Copyright © 2016, Juniper Networks, Inc.

family inet (Interfaces)

Syntax

```
inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address (source–address/prefix) {
    arp destination-address {
      (mac mac-address | multicast-mac multicast-mac-address);
      publish publish-address;
    }
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key-value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds
      (preempt <hold-time seconds> | no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
          priority-cost value;
        }
      }
      virtual-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  dhcp {
    client-identifier {
```

```
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      user-id (ascii string| hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp arp-name;
    input input-name;
    output output-name;
  }
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
    simple-filter;
  }
  targeted-broadcast {
    (forward-and-send-to-re |forward-only);
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
}
```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement introduced in a prior release of Junos OS.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.

NOTE: You use family inet to assign an IPv4 address. You use family inet6 to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

ftp (Access)

Syntax

```
ftp {
  banner {
    fail string;
    login string;
    success string;
  }
}
```

Hierarchy Level [edit access firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure banners for the FTP login prompt, successful authentication, and failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Pass-Through Authentication on page 5205
- Example: Configuring Pass-Through Authentication on page 5207

Copyright © 2016, Juniper Networks, Inc.

ftp (Security ALG)

Syntax

```
ftp {
  allow-mismatch-ip-address;
  disable;
  ftps-extension;
  line-break-extension;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement modified in Junos OS Release 11.4.

Description Specify the FTP ALG on the device.

Options

- **disable**—Disable the FTP ALG. By default, the FTP ALG is enabled. This option will enable or disable FTP ALG for both IPV4 and IPV6 mode.
- **ftps-extension**—Enable secure FTP and FTP SSL protocols.
- **line-break-extension**—Enable line-break-extension. This option will enable the FTP ALG to recognize the LF as line break in addition to the standard CR+LF (carriage return, followed by line feed).
- **traceoptions**—Configure FTP ALG tracing options. To specify more than one trace operation, include multiple flag statements.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - **extensive**—(Optional) Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- FTP ALG Overview on page 25

Copyright © 2016, Juniper Networks, Inc.

gatekeeper

Syntax

```
gatekeeper {
  threshold rate;
}
```

Hierarchy Level [edit security alg h323 application-screen message-flood]

Release Information Statement introduced in Junos OS Release 8.5.

Description Limit the rate at which remote access server (RAS) requests to the gatekeeper are processed. Messages exceeding the threshold are dropped.

Options **threshold *rate***—Threshold measured in messages per second.

Range: 1 through 50,000

Default: 1000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- h323 on page 360

Copyright © 2016, Juniper Networks, Inc.

h323

Syntax

```
h323 {
  application-screen {
    message-flood {
      gatekeeper {
        threshold rate;
      }
    }
    unknown-message {
      permit-nat-applied;
      permit-routed;
    }
  }
  disable;
  dscp-rewrite {
    code-point string;
  }
  endpoint-registration-timeout value-in-seconds;
  media-source-port-any;
  traceoptions {
    flag flag <detail | extensive | terse>;
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the H.323 Application Layer Gateway (ALG) on the device. H.323 is a control-signaling protocol used to exchange messages between H.323 endpoints.

Options **disable**—Disable the H.323 ALG. By default, H.323 ALG is enabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- application-screen (Security H323) on page 336
- endpoint-registration-timeout on page 354
- dscp-rewrite on page 353

Copyright © 2016, Juniper Networks, Inc.

host-inbound-traffic

Syntax

```
host-inbound-traffic {
  protocols protocol-name {
    except;
  }
  system-services service-name {
    except;
  }
}
```

Hierarchy Level [edit security zones functional-zone management],
[edit security zones functional-zone management interfaces *interface-name*],
[edit security zones security-zone *zone-name*],
[edit security zones security-zone *zone-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Control the type of traffic that can reach the device from interfaces bound to the zone.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058

Copyright © 2016, Juniper Networks, Inc.

ike-esp-nat

Syntax

```
ike-esp-nat {
  enable;
  esp-gate-timeout seconds;
  esp-session-timeout seconds;
  state-timeout seconds;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure Application Layer Gateway (ALG) for Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) traffic with Network Address Translation (NAT).

Options

- **Enable** —Enable the IKE-ESP ALG.
- **esp-gate-timeout *seconds***—Set the timeout for the ESP gates created after an IKE Phase 2 exchange has completed.

  Range: 2 through 30 seconds.

  Default: 5 seconds.
- **esp-session-timeout *seconds***—Set the idle timeout for the ESP sessions created from the IPsec gates.

  Range: 60 through 2400 seconds.

  Default: 1800 seconds.
- **state-timeout *seconds***—Set the timeout for the ALG state information.

  Range: 180 through 86,400 seconds.

  Default: 14,400 seconds.
- **traceoptions**—Set the IKE-ESP ALG trace options.
  - **flag** —Specify which tracing operation to perform.
    - **all**—Trace all operations.
      - **extensive**—Set trace verbosity level to extensive.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the ALG for IKE and ESP on page 33

 Copyright © 2016, Juniper Networks, Inc.

ike (Security)

Syntax

```
ike {
  gateway gateway-name {
    address [ip-address-or-hostname];
    advpn {
      suggester {
        disable;
      }
      partner {
        connection-limit <number>;
        idle-threshold <packets/sec>;
        idle-time <seconds>;
        disable;
      }
    }
    dead-peer-detection {
      (always-send | optimized | probe-idle-tunnel);
      interval seconds;
      threshold number;
    }
    dynamic {
      connections-limit number;
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
      ike-user-type (group-ike-id | shared-ike-id);
    }
    external-interface external-interface-name;
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
      (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address |
        user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    remote-identity {
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    version (v1-only | v2-only);
    xauth {
      access-profile profile-name;
    }
  }
  policy policy-name {
    certificate {
      local-certificate certificate-id;
      peer-certificate-type (pkcs7 | x509-signature);
      policy-oids [ oid ];
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1);
    authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384
      | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
  }
  respond-bad-spi <max-responses>;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    rate-limit messages-per-second;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The **inet6** option added in Junos OS Release 11.1.

Description Define Internet Key Exchange (IKE) configuration.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- ALG Overview on page 3
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

inactive-media-timeout (Security MGCP)

Syntax inactive-media-timeout *value-in-seconds*;

Hierarchy Level [edit security alg mgcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options ***value-in-seconds***—Maximum amount of time that the pinholes remain open.

Range: 10 through 2550 seconds

Default: 120 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369

inactive-media-timeout (Security SCCP)

Syntax inactive-media-timeout *value-in-seconds*;

Hierarchy Level [edit security alg sccp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options ***value-in-seconds***—Maximum amount of time that the pinholes remain open.

Range: 10 through 600 seconds

Default: 120 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sccp on page 384

Copyright © 2016, Juniper Networks, Inc.

inactive-media-timeout (Security SIP)

Syntax inactive-media-timeout *value-in-seconds*;

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum amount of time that the temporary openings in the firewall (pinholes) remain open for media if no activity is detected.

Options ***value-in-seconds***—Maximum amount of time that the pinholes remain open.

Range: 0 through 2550 seconds

Default: 120 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

map-entry-timeout

Syntax map-entry-timeout *map-entry-timeout*;

Hierarchy Level [edit security alg msrpc]
[edit security alg sunrpc]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure the mapping entry timeout value. When the incoming traffic hits the mapping entry, the timeout value has been reset to configured value. The mapping entry is removed from the table when the timeout value expires. The lifetime of the mapping entry is global and applies to all entries in the table.

Options ***map-entry-timeout***—Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) or Sun RPC ALG mapping entry timeout value in hours.

Range: 1 through 72 hours.

Default: 32 hours.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- msrpc on page 370
- sunrpc on page 393

Copyright © 2016, Juniper Networks, Inc.

maximum-call-duration (Security)

Syntax maximum-call-duration *value-in-minutes*;

Hierarchy Level [edit security alg mgcp],
[edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the time at which the SIP call ends. The media session is released after the call has ended.

Options ***value-in-minutes***—Maximum amount of time at which the call ends and releases the media sessions.

Range: 3 through 720 minutes

Default: 720 minutes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369
- sip (Security) on page 387

maximum-message-length

Syntax maximum-message-length *number*;

Hierarchy Level [edit security alg dns]

Release Information Statement introduced in Junos OS Release 10.1.

Description Specify the maximum DNS message length.

Options **number**—Maximum length in bytes of a single DNS message.

Range: 512 through 8192 bytes.

Default: 512 bytes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- dns on page 351

Copyright © 2016, Juniper Networks, Inc.

media-source-port-any

Syntax media-source-port-any;

Hierarchy Level [edit security alg h323]

Release Information Statement introduced in Junos OS Release 8.5.

Description Allow media traffic from any port number. By default, this feature is disabled, which allows a temporary opening in the firewall (pinhole) for media traffic to be opened.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- h323 on page 360

message-flood (Security H323)

Syntax
```
message-flood {
  gatekeeper {
    threshold rate;
  }
}
```

Hierarchy Level [edit security alg h323 application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Limit the rate per second at which remote access server (RAS) requests to the gatekeeper are processed. Messages exceeding the threshold are dropped. This feature is disabled by default.

Options
gatekeeper threshold *rate*—Maximum number of RAS connection requests per second allowed per gateway.

Range: 1 through 65,535

Default: 1000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- h323 on page 360

Copyright © 2016, Juniper Networks, Inc.

mgcp

Syntax

```
mgcp {
  application-screen {
    connection-flood {
      threshold rate;
    }
    message-flood {
      threshold rate;
    }
    unknown-message {
      permit-nat-applied;
      permit-routed;
    }
  }
  disable;
  dscp-rewrite {
    code-point string;
  }
  inactive-media-timeout value-in-seconds;
  maximum-call-duration value-in-minutes;
  traceoptions {
    flag flag <extensive>;
  }
  transaction-timeout value-in-seconds;
}
```

Hierarchy Level [edit security alg]

Release Information Statement modified in Junos OS Release 9.2.

Description Specify the Media Gateway Control Protocol (MGCP) ALG on the device. MGCP is a text-based Application Layer protocol that can be used for call setup and call control.

Options **disable**—Disable the MGCP ALG. By default, the MGCP ALG is enabled.

NOTE: By default, the MGCP ALG is disabled for SRX Series devices.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the MGCP ALG on page 183

Copyright © 2016, Juniper Networks, Inc.

msrpc

Syntax

```
msrpc {
  disable;
  map-entry-timeout map-entry-timeout;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the Microsoft remote procedure call Application Layer Gateway (MS-RPC ALG) on the device.

Options

- **disable**—Disable the MS-RPC ALG. By default, the MS-RPC ALG is enabled.
- **map-entry-timeout** ***map-entry-timeout***—Specify the MS-RPC ALG mapping entry timeout value in hours.

  Range: 1 through 72 hours.

  Default: 32 hours.
- **traceoptions**—Configure MS-RPC ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Microsoft RPC ALGs on page 67

Copyright © 2016, Juniper Networks, Inc.

nat

Syntax

```
nat {
  destination {
    pool pool-name {
      address ip-address {
        (port port-number | to ip-address);
      }
      description text;
      routing-instance routing-instance-name;
    }
    rule-set rule-set-name {
      description text;
      from {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
      rule rule-name {
        description text;
        match {
          (destination-address <ip-address> | destination-address-name <address-name>);
          destination-port port-number;
          protocol [protocol-name-or-number];
          source-address [ip-address];
          source-address-name [address-name];
        }
        then {
          destination-nat (off | pool pool-name);
        }
      }
    }
  }
  proxy-arp {
    interface interface-name {
      address ip-address {
        to ip-address;
      }
    }
  }
  proxy-ndp {
    interface interface-name {
      address ip-address {
        to ip-address;
      }
    }
  }
  source {
    address-persistent;
    interface {
      port-overloading {
        off;
      }
    }
```

```
pool pool-name {
  address ip-address {
    to ip-address;
  }
  description text;
  host-address-base ip-address;
  overflow-pool (interface | pool-name);
  port {
    (no-translation | port-overloading-factor number | range port-low <to port-high>);
  }
  routing-instance routing-instance-name;
}
pool-default-port-range lower-port-range to upper-port-range;
pool-utilization-alarm {
  clear-threshold value;
  raise-threshold value;
}
port-randomization {
  disable;
}
port-round-robin {
  disable;
}
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      (destination-address <ip-address> | destination-address-name <address-name>);
      destination-port port-number;
      protocol [protocol-name-or-number];
      source-address [ip-address];
      source-address-name [address-name];
    }
    then {
      source-nat {
        interface {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number value;
            permit (any-remote-host | target-host | target-host-port);
          }
        }
        off;
        pool {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number number;
            permit (any-remote-host | target-host | target-host-port);
```


Copyright © 2016, Juniper Networks, Inc.

```
                    }
                    pool-name;
                  }
                }
              }
            }
            to {
              interface [interface-name];
              routing-instance [routing-instance-name];
              zone [zone-name];
            }
          }
        }
        static {
          rule-set rule-set-name {
            description text;
            from {
              interface [interface-name];
              routing-instance [routing-instance-name];
              zone [zone-name];
            }
            rule rule-name {
              description text;
              match {
                (destination-address ip-address | destination-address-name address-name);
              }
              then {
                static-nat {
                  inet {
                    routing-instance (default | routing-instance-name);
                  }
                  prefix {
                    address-prefix;
                    routing-instance (default | routing-instance-name);
                  }
                  prefix-name {
                    address-prefix-name;
                    routing-instance (default | routing-instance-name);
                  }
                }
              }
            }
          }
        }
        traceoptions {
          file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
        }
      }
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1.

Description Configure Network Address Translation (NAT) for SRX Series devices.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839
- Understanding Logical System Network Address Translation on page 3355

Copyright © 2016, Juniper Networks, Inc.

nat-pat-address

Syntax

```
nat-pat-address {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT with port address translation (PAT) configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Network Address Translation on page 3355
- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

policy (Security Policies)

Syntax

```
policy policy-name {
  description description;
  match {
    application {
      [application];
      any;
    }
    destination-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
```


Copyright © 2016, Juniper Networks, Inc.

```
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              web-redirect;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. The **services-offload** option added in Junos OS Release 11.4. Statement updated with the **source-identity** option and the **description** option added in Junos OS Release 12.1. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Define a security policy.

Copyright © 2016, Juniper Networks, Inc.

Options ***policy-name***—Name of the security policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320
- *Building Blocks Feature Guide for Security Devices*

pptp

Syntax

```
pptp {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Point-to-Point Tunneling Protocol (PPTP) ALG on the device.

Options
- **disable**—Disable the PPTP ALG. By default, the PPTP ALG is enabled.
- **traceoptions**—Configure PPTP ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- *Building Blocks Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

protect

Syntax

```
protect {
  deny {
    all {
      timeout value-in-seconds;
    }
    destination-ip address;
    timeout value-in-seconds;
  }
}
```

Hierarchy Level [edit security alg sip application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure options to protect servers against INVITE attacks.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

Copyright © 2016, Juniper Networks, Inc.

protocols (Security Zones Interfaces)

Syntax

```
protocols protocol-name {
    except;
}
```

Hierarchy Level [edit security zones security-zone *zone-name* interfaces *interface-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of routing protocol traffic that can reach the device on a per-interface basis.

Options

- ***protocol-name*** —Protocol for which traffic is allowed. The following protocols are supported:
  - **all**—Enable traffic from all possible protocols available.
  - **bfd**—Enable incoming Bidirectional Forwarding Detection (BFD) Protocol traffic.
  - **bgp**—Enable incoming BGP traffic.
  - **dvmrp**—Enable incoming Distance Vector Multicast Routing Protocol (DVMRP) traffic.
  - **igmp**—Enable incoming Internet Group Management Protocol (IGMP) traffic.
  - **ldp**—Enable incoming Label Distribution Protocol (LDP) traffic (UDP and TCP port 646).
  - **msdp**—Enable incoming Multicast Source Discovery Protocol (MSDP) traffic.
  - **nhrp**—Enable incoming Next Hop Resolution Protocol (NHRP) traffic.
  - **ospf**—Enable incoming OSPF traffic.
  - **ospf3**—Enable incoming OSPF version 3 traffic.
  - **pgm**—Enable incoming Pragmatic General Multicast (PGM) protocol traffic (IP protocol number 113).
  - **pim**—Enable incoming Protocol Independent Multicast (PIM) traffic.
  - **rip**—Enable incoming RIP traffic.
  - **ripng**—Enable incoming RIP next generation traffic.
  - **router-discovery**—Enable incoming router discovery traffic.
  - **rsvp**—Enable incoming Resource Resolution Protocol (RSVP) traffic (IP protocol number 46).
  - **sap**— Enable incoming Session Announcement Protocol (SAP) traffic. SAP always listens on **224.2.127.254:9875**.
  - **vrrp**—Enable incoming Virtual Router Redundancy Protocol (VRRP) traffic.

except—(Optional) except can only be used if all has been defined.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050

retain-hold-resource

Syntax
retain-hold-resource;

Hierarchy Level
[edit security alg sip]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Enable the device to not free media resources for a Session Initiation Protocol (SIP) Application Layer Gateway (ALG), even when a media stream is placed on hold. By default, media stream resources are released when the media stream is held.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- sip (Security) on page 387

rsh

Syntax

```
rsh {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the remote shell (RSH) ALG on the device.

Options

- **disable**—Disable the RSH ALG. By default, the RSH ALG is disabled.
- **traceoptions**—Configure RSH ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the RSH ALG on page 75

Copyright © 2016, Juniper Networks, Inc.

rtsp

Syntax

```
rtsp {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Real-Time Streaming Protocol (RTSP) ALG on the device.

Options

- **disable**—Disable the RTSP ALG. By default, the RTSP ALG is enabled.
- **traceoptions**—Configure RTSP ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the RTSP ALG on page 89

sccp

Syntax

```
sccp {
  application-screen {
    call-flood {
      threshold rate;
    }
    unknown-message {
      permit-nat-applied;
      permit-routed;
    }
  }
  disable;
  dscp-rewrite {
    code-point string;
  }
  inactive-media-timeout value-in-seconds;
  traceoptions {
    flag flag <extensive>;
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Skinny Client Control Protocol (SCCP) ALG on the device.

Options **disable**—Disable the SCCP ALG. By default, the SCCP ALG is enabled.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding SCCP ALGs on page 217

Copyright © 2016, Juniper Networks, Inc.

security-zone

Syntax

```
security-zone zone-name {
  address-book {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
  }
  application-tracking;
  description text;
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
  interfaces interface-name {
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
  }
  screen screen-name;
  tcp-rst;
}
```

Hierarchy Level [edit security zones]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description Define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Copyright © 2016, Juniper Networks, Inc.

Options

zone-name —Name of the security zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [edit security zones] Hierarchy Level on page 324
- Security Zones and Interfaces Overview on page 1049
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

Copyright © 2016, Juniper Networks, Inc.

sip (Security)

Syntax

```
sip {
  application-screen {
    protect {
      deny {
        all {
          timeout value-in-seconds;
        }
        destination-ip address;
        timeout value-in-seconds;
      }
    }
    unknown-message {
      permit-nat-applied;
      permit-routed;
    }
  }
  c-timeout value-in-minutes;
  disable;
  dscp-rewrite {
    code-point string;
  }
  inactive-media-timeout value-in-seconds;
  maximum-call-duration value-in-minutes;
  retain-hold-resource;
  t1-interval value-in-milliseconds;
  t4-interval value-in-seconds;
  traceoptions {
    flag flag <detail | extensive | terse>;
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement modified in Junos OS Release 9.2.

Description Specify the Session Initiation Protocol (SIP) ALG on the device.

NOTE: IPv6 is supported on the SIP ALG along with Network Address Translation Protocol Translation (NAT-PT) mode and NAT64 address translation.

The type of the <destination-ip-address> is changed from IPv4 address to IP prefix to support all kinds of IP addresses, and correspondingly a prefix is supported to allow multiple IP addresses.

Options **disable**—Disable the SIP ALG. Use the **set security alg *alg-name* disable** CLI command to disable the SIP ALG.

Copyright © 2016, Juniper Networks, Inc.

The remaining statements are explained separately. See CLI Explorer.

NOTE: By default, the SIP ALG is disabled for SRX5600 and SRX5800 devices. Use the CLI command **delete security alg *alg-name*** disable to enable the SIP ALG.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding the SIP ALG on page 241

source-address (Security Destination NAT)

Syntax
source-address [*ip-address*];

Hierarchy Level
[edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information
Statement modified in Junos OS Release 9.6.

Description
Specify source address to match the rule. You can configure multiple addresses or subnets.

Options
ip-address —Source address or a subnet.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-address (Security Policies)

Syntax

```
source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options

address—IP addresses, address sets, or wildcard addresses (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

source-nat

Syntax

```
source-nat {
  interface {
    persistent-nat {
      address-mapping;
      inactivity-timeout seconds;
      max-session-number value;
      permit (any-remote-host | target-host | target-host-port);
    }
  }
  off;
  pool <pool-name>;
    persistent-nat {
      address-mapping;
      inactivity-timeout seconds;
      max-session-number number;
      permit (any-remote-host | target-host | target-host-port);
  }
  rule-session-count-alarm (clear-threshold value | raise-threshold value);
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Specify the action of the source NAT rule.

Options

- **off**—Do not perform the source NAT operation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security — To view this statement in the configuration.
security-control— To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

sql

Syntax

```
sql {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Oracle SQL ALG on the device.

Options

- **disable**—Disable the SQL ALG. By default, the SQL ALG is enabled.
- **traceoptions**—Configure SQL ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the SQLNET ALG on page 103

Copyright © 2016, Juniper Networks, Inc.

static-nat

Syntax

```
static-nat {
  inet {
    routing-instance (default | routing-instance-name);
  }
  nptv6-prefix {
    address-prefix;
      routing-instance routing-instance-name;
  nptv6-prefix-name {
    address-prefix-name;
      routing-instance routing-instance-name;
  [
  prefix {
    address-prefix;
    routing-instance (default | routing-instance-name);
  }
  prefix-name {
    address-prefix-name;
    routing-instance (default | routing-instance-name);
  }
  rule-session-count-alarm (clear threshold value | raise threshold value);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

Description Specify the translated address of the static NAT rule.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

sunrpc

Syntax

```
sunrpc {
  disable;
  map-entry-timeout map-entry-timeout;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the Sun Microsystems remote procedure call (RPC) ALG on the device.

Options

- **disable**—Disable the Sun RPC ALG. By default, the Sun RPC ALG is enabled.
- **map-entry-timeout** ***map-entry-timeout***—Specify the Sun RPC ALG mapping entry timeout value in hours.

  Range: 1 through 72 hours.

  Default: 32 hours.
- **traceoptions**—Configure the Sun RPC ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Sun RPC ALGs on page 62

Copyright © 2016, Juniper Networks, Inc.

support-lib

Syntax

```
alg-support-lib {
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the Application Layer Gateway (ALG) support library.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- alg on page 327

Copyright © 2016, Juniper Networks, Inc.

system-services (Security Zones Host Inbound Traffic)

Syntax

```
system-services {
  (service-name | all <service-name except>);
}
```

Hierarchy Level [edit security zones security-zone *zone-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of incoming system service traffic that can reach the device for all interfaces in a zone. You can do this in one of several ways:

- You can enable traffic from each system service individually.
- You can enable traffic from all system services.
- You can enable traffic from all but some system services.

Options

- ***service-name*** —System-service for which traffic is allowed. The following system services are supported:
  - **all**—Enable traffic from the defined system services available on the Routing Engine (RE). Use the *except* option to disallow specific system services.
  - **any-service**—Enable all system services on entire port range including the system services that are not defined.
  - **bootp**—Enable traffic destined to BOOTP and DHCP relay agents.
  - **dhcp**—Enable incoming DHCP requests.
  - **dhcpv6**—Enable incoming DHCP requests for IPv6.
  - **dns**—Enable incoming DNS services.
  - **finger**—Enable incoming finger traffic.
  - **ftp**—Enable incoming FTP traffic.
  - **http**—Enable incoming J-Web or clear-text Web authentication traffic.
  - **https**—Enable incoming J-Web or Web authentication traffic over Secure Sockets Layer (SSL).
  - **ident-reset**—Enable the access that has been blocked by an unacknowledged identification request.
  - **ike**—Enable Internet Key Exchange traffic.
  - **lsping**—Enable label switched path ping service.
  - **netconf**—Enable incoming NETCONF service.
  - **ntp**—Enable incoming Network Time Protocol (NTP) traffic.
  - **ping**—Allow the device to respond to ICMP echo requests.
  - **r2cp**—Enable incoming Radio Router Control Protocol traffic.

Copyright © 2016, Juniper Networks, Inc.

- **reverse-ssh**—Reverse SSH traffic.
- **reverse-telnet**—Reverse Telnet traffic.
- **rlogin**—Enable incoming **rlogin** (remote login) traffic.
- **rpm**—Enable incoming Real-time performance monitoring (RPM) traffic.
- **rsh**—Enable incoming Remote Shell (**rsh**) traffic.
- **snmp**—Enable incoming SNMP traffic (UDP port 161).
- **snmp-trap**—Enable incoming SNMP traps (UDP port 162).
- **ssh**—Enable incoming SSH traffic.
- **telnet**—Enable incoming Telnet traffic.
- **tftp**—Enable TFTP services.
- **traceroute**—Enable incoming traceroute traffic (UDP port 33434).
- **xnm-clear-text**—Enable incoming Junos XML protocol traffic for all specified interfaces.
- **xnm-ssl**— Enable incoming Junos XML protocol-over-SSL traffic for all specified interfaces.

- **except**—(Optional) Enable specific incoming system service traffic but only when the *all* option has been defined . For example, to enable all but FTP and HTTP system service traffic:

  ```
  set security zones security-zone trust host-inbound-traffic system-services all
  set security zones security-zone trust host-inbound-traffic system-services ftp except
  set security zones security-zone trust host-inbound-traffic system-services http except
  ```

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

t1-interval

Syntax t1-interval *value-in-milliseconds*;

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum round-trip time (RTT) (in milliseconds) allowed for Session Initiation Protocol (SIP) transactions.

Options ***value-in-milliseconds***—Maximum round-trip time (RTT) allowed measured in milliseconds.

Range: 500 through 5000 milliseconds

Default: 500 milliseconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- sip (Security) on page 387

t4-interval

Syntax t4-interval *value-in-seconds*;

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum length of time (in seconds) that the network can take to clear messages between client and server Session Initiation Protocol (SIP) transactions.

Options ***value-in-seconds***—Maximum number of seconds that the network can take to clear messages between client and server transactions.

Range: 5 through 10 seconds

Default: 5 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- sip (Security) on page 387

Copyright © 2016, Juniper Networks, Inc.

talk

Syntax

```
talk {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the TALK program ALG on the device.

Options

- **disable**—Disable the TALK program ALG. By default, the TALK program ALG is enabled.
- **traceoptions**—Configure TALK program ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the TALK ALG on page 117

Copyright © 2016, Juniper Networks, Inc.

term (Applications)

Syntax

```
term term-name {
    alg application;
    destination-port port-identifier;
    icmp-code value;
    icmp-type value;
    icmp6-code value;
    icmp6-type value;
    inactivity-timeout (seconds | never);
    protocol number;
    rpc-program-number number;
    source-port port-number;
    uuid hex-value;
}
```

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define individual application protocols.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Security Policy Applications Overview on page 1171

Copyright © 2016, Juniper Networks, Inc.

tftp (Security ALG)

Syntax

```
tftp {
  disable;
  traceoptions {
    flag {
      all <extensive>;
    }
  }
}
```

Hierarchy Level [edit security alg]

Release Information Statement modified in Junos OS Release 9.2.

Description Configure the Trivial File Transfer Protocol (TFTP) ALG on the device.

Options

- **disable**—Disable the TFTP ALG. By default, the TFTP ALG is enabled.

  NOTE: By default, the TFTP ALG is disabled for SRX Series devices.

- **traceoptions**—Configure TFTP ALG tracing options.
  - **flag**—Trace operation to perform.
    - **all**—Trace all events.
    - ***extensive***—Display extensive amount of data.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding the TFTP ALG on page 131

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security ALG)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  level (brief | detail | extensive | verbose);
  no-remote-trace;
}
```

Hierarchy Level [edit security alg]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure ALG tracing options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **level**—Set the level of debugging the output option.
  - **brief**—Match brief messages
  - **detail**—Match detail messages.
  - **extensive**—Match extensive messages.
  - **verbose**—Match verbose messages.
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- alg on page 327

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security H323 ALG)

Syntax

```
traceoptions {
  flag flag <detail | extensive | terse>;
}
```

Hierarchy Level [edit security alg h323]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure H.323 tracing options.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled.
  - **chassis-cluster**—Trace chassis cluster information.
  - **h225-asn1**—Trace H.225 ASN.1 processing activity.
  - **h245**—Trace H.245 processing activity.
  - **h245-asn1**—Trace H.245 ASN.1 processing activity.
  - **q931**—Trace Q.931 processing activity.
  - **ras**—Trace remote access server (RAS) processing activity.
  - **ras-asn1**—Trace RAS ASN.1 processing activity.
- **detail**—Display moderate amount of data in trace.
- **extensive**—Display extensive amount of data in trace.
- **terse**—Display minimum amount of data in trace.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- h323 on page 360

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security MGCP ALG)

Syntax

```
traceoptions {
  flag flag <extensive>;
}
```

Hierarchy Level [edit security alg mgcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure Media Gateway Control Protocol (MGCP) tracing options.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled.
  - **call**—Trace call processing activity.
  - **chassis-cluster**—Trace chassis cluster information.
  - **decode**—Trace decoder operations activity.
  - **error**—Trace processing errors activity.
  - **nat**—Trace Network Address Translation (NAT) processing activity.
  - **packet**—Trace MGCP protocol packet processing activity.
  - **rm**—Trace MGCP Resource Management (Resmgr) functions activity.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security SCCP ALG)

Syntax

```
traceoptions {
  flag flag <extensive>;
}
```

Hierarchy Level [edit security alg sccp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure Skinny Client Control Protocol (SCCP) tracing options.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled.
  - **call**—Trace call processing activity.
  - **chassis-cluster**—Trace chassis cluster information.
  - **cli**—Trace CLI configuration activity and command changes.
  - **decode**—Trace decoder operations activity.
  - **error**—Trace processing errors activity.
  - **init**—Enable tracing for SCCP initialization errors
  - **nat**—Enable tracing for SCCP NAT processing
  - **rm**——Trace SCCP Resource Management (Resmgr) functions activity.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- sccp on page 384

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security SIP ALG)

Syntax

```
traceoptions {
  flag flag <detail | extensive | terse>;
}
```

Hierarchy Level [edit security alg sip]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure Session Initiation Protocol (SIP) tracing options.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled.
  - **call**—Trace call processing activity.
  - **chassis-cluster**—Trace chassis cluster information.
  - **nat**—Trace SIP Network Address Translation (NAT) processing activity.
  - **parser**—Trace SIP parser operations.
  - **rm**—Trace SIP Resource Management (Resmgr) functions activity.
- detail—Display moderate amount of data in trace.
- extensive—Display extensive amount of data in trace.
- terse—Display minimum amount of data in trace.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

Copyright © 2016, Juniper Networks, Inc.

traceoptions (System Services DNS)

Syntax

```
traceoptions {
  category {
    category-type;
  }
  file;
}
```

Hierarchy Level [edit system services dns]

Release Information Statement introduced in Junos OS Release 10.2.

Description Defines tracing options for DNS services.

Options

category—Specifies the logging category. See Table 15 for the different logging categories and their descriptions.

file—Trace file information.

Copyright © 2016, Juniper Networks, Inc.

Table 15: Category Names

Category Name	Description
client	Processing of client requests
config	Configuration file parsing and processing
database	Messages relating to the databases
default	Categories for which there is no specific configuration
delegation-only	Delegation only
dispatch	Dispatching of incoming packets to the server
dnssec	DNSSEC and TSIG protocol processing
edns-disabled	Log query using plain DNS
general	General information
lame-servers	Lame servers
network	Network options
notify	NOTIFY protocol
queries	DNS query resolver
resolver	DNS resolution security
security	Approval and denial of requests
unmatched	Unable to determine the class for messages named
update	Dynamic updates
update-security	Approval and denial of update requests
xfer-in	Zone transfers that the server is receiving xfer-out
xfer-out	Zone transfers that the server is sending

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

transaction-timeout

Syntax transaction-timeout *value-in-seconds*;

Hierarchy Level [edit security alg mgcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Timeout value for Media Gateway Control Protocol (MGCP) transactions. If the timeout value exceeds transaction will be removed by MGCP transactions ager out processing.

Options ***value-in-seconds***—Timeout value.

Range: 3 through 50 seconds

Default: 30 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369

Copyright © 2016, Juniper Networks, Inc.

unknown-message (Security H323 ALG)

Syntax

```
unknown-message {
  permit-nat-applied;
  permit-routed;
}
```

Hierarchy Level [edit security alg h323 application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how unidentified H.323 messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown H.323 (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped.

This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.

Options

- **permit-nat-applied**—Specify that unknown messages be allowed to pass if the session is in NAT mode.
- **permit-routed**—Specify that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- h323 on page 360

 Copyright © 2016, Juniper Networks, Inc.

unknown-message (Security MGCP ALG)

Syntax

```
unknown-message {
  permit-nat-applied;
  permit-routed;
}
```

Hierarchy Level [edit security alg mgcp application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how unidentified Media Gateway Control Protocol (MGCP) messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown MGCP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped.

This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.

Options

- **permit-nat-applied**—Specify that unknown messages be allowed to pass if the session is in NAT mode.
- **permit-routed**— Specify that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mgcp on page 369

Copyright © 2016, Juniper Networks, Inc.

unknown-message (Security SCCP ALG)

Syntax

```
unknown-message {
  permit-nat-applied;
  permit-routed;
}
```

Hierarchy Level [edit security alg sccp application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how unidentified Skinny Client Control Protocol (SCCP) messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown SCCP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped.

This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.

Options

- **permit-nat-applied**—Specify that unknown messages be allowed to pass if the session is in NAT mode.
- **permit-routed**— Specify that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sccp on page 384

Copyright © 2016, Juniper Networks, Inc.

unknown-message (Security SIP ALG)

Syntax

```
unknown-message {
  permit-nat-applied;
  permit-routed;
}
```

Hierarchy Level [edit security alg sip application-screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how unidentified Session Initiation Protocol (SIP) messages are handled by the device. The default is to drop unknown (unsupported) messages. Permitting unknown messages can compromise security and is not recommended. However, in a secure test or production environment, this statement can be useful for resolving interoperability issues with disparate vendor equipment. By permitting unknown SIP (unsupported) messages, you can get your network operational and later analyze your VoIP traffic to determine why some messages were being dropped.

This statement applies only to received packets identified as supported VoIP packets. If a packet cannot be identified, it is always dropped. If a packet is identified as a supported protocol, the message is forwarded without processing.

Options

- **permit-nat-applied**—Specify that unknown messages be allowed to pass if the session is in NAT mode.
- **permit-routed**— Specify that unknown messages be allowed to pass if the session is in Route mode. (Sessions in Transparent mode are treated as Route mode.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- sip (Security) on page 387

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 20

Operational Commands

- clear security alg h323 counters
- clear security alg ike-esp-nat
- clear security alg sccp calls
- clear security alg sccp counters
- clear security alg sip calls
- clear security alg sip counters
- clear security flow session application
- show chassis cluster data-plane statistics
- show chassis cluster statistics
- show security alg h323 counters
- show security alg ike-esp-nat summary
- show security alg msrpc
- show security alg sccp calls
- show security alg sccp counters
- show security alg sip calls
- show security alg sip counters
- show security alg sip rate
- show security alg status
- show security flow gate
- show security flow session
- show security flow session application
- show security flow session resource-manager
- show security idp policy-templates
- show security resource-manager group active
- show security resource-manager resource active
- show security resource-manager summary
- show security zones
- show security zones type

Copyright © 2016, Juniper Networks, Inc.

clear security alg h323 counters

Syntax clear security alg h323 counters
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Clear information about H.323 Application Layer Gateway (ALG) counters.

Options

- none—Clear H.323 ALG counters.
- **node**—(Optional) For chassis cluster configurations, clear H.323 counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- h323 on page 360
- show security alg h323 counters on page 430

List of Sample Output clear security alg h323 counters on page 416

Sample Output

clear security alg h323 counters

```
user@host> clear security alg h323 counters
```


Copyright © 2016, Juniper Networks, Inc.

clear security alg ike-esp-nat

Syntax clear security alg ike-esp-nat

Release Information Command introduced in Junos OS Release 10.2.

Description Clear state information about Application Layer Gateway (ALG) for IKE and ESP.

Required Privilege Level clear

Related Documentation
- show security alg ike-esp-nat summary on page 432

List of Sample Output clear security alg ike-esp-nat on page 417

Sample Output

clear security alg ike-esp-nat

```
user@host> clear security alg ike-esp-nat
10 active IKE-ESP alg state cleared
```

clear security alg sccp calls

Syntax

clear security alg sccp calls
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Clear Skinny Client Protocol (SCCP) Application Layer Gateway (ALG) call information.

Options

- none—Clear all SCCP ALG calls.
- **node**—(Optional) For chassis cluster configurations, clear SCCP calls on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level

clear

Related Documentation

- sccp on page 384
- show security alg sccp calls on page 435

Output Fields

This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security alg sccp counters

Syntax clear security alg sccp counters
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Clear Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) counters.

Options

- none—Clear all SCCP ALG counters.
- **node**—(Optional) For chassis cluster configurations, clear SCCP counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- sccp on page 384
- show security alg sccp counters on page 437

List of Sample Output clear security alg sccp counters on page 419

Sample Output

clear security alg sccp counters

```
user@host> clear security alg sccp counters
```

clear security alg sip calls

Syntax clear security alg sip calls
<node (*node-id* | all | local | primary)>

Release Information Command modified in Junos OS Release 9.2; **node** options added in Junos OS Release 9.0.

Description Clear Session Initiation Protocol (SIP) Application Layer Gateway (ALG) call information.

Options

- none—Clear all SIP ALG calls.
- **node**—(Optional) For chassis cluster configuration, clear SIP calls on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- sip (Security) on page 387
- show security alg sip calls on page 439

List of Sample Output clear security alg sip calls on page 420

Sample Output

clear security alg sip calls

```
user@host> clear security alg sip calls
```


Copyright © 2016, Juniper Networks, Inc.

clear security alg sip counters

Syntax clear security alg sip counters
<node (*node-id* | all | local | primary)>

Release Information Command modified in Junos OS Release 9,2; **node** options added in Junos OS Release 9.0.

Description Clear Session Initiation Protocol (SIP) Application Layer Gateway (ALG) counters.

Options

- none—Clear all SIP ALG counters.
- **node**—(Optional) For chassis cluster configurations, clear SIP counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- sip (Security) on page 387
- show security alg sip counters on page 441

List of Sample Output clear security alg sip counters on page 421

Sample Output

clear security alg sip counters

```
user@host> clear security alg sip counters
```

clear security flow session application

Syntax clear security flow session application
application-name
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear currently active sessions for application types or application sets.

Options

- *application-name* —Name of the specified application type or application set.
  - **dns**—Domain Name System
  - **ftp**—File Transfer Protocol
  - **ignore**—Ignore application type
  - **mgcp-ca**—Media Gateway Control Protocol with Call Agent
  - **mgcp-ua**—MGCP with User Agent
  - **ms-rpc**—Microsoft RPC
  - **pptp**—Point-to-Point Tunneling Protocol
  - **q931**—ISDN connection control protocol
  - **ras**—RAS
  - **realaudio**—RealAudio
  - **rsh**—UNIX remote shell services
  - **rtsp**—Real-Time Streaming Protocol
  - **sccp**—Skinny Client Control Protocol
  - **sip**—Session Initiation Protocol
  - **sqlnet-v2**—Oracle SQLNET
  - **sun-rpc**—Sun Microsystems RPC
  - **talk**—TALK program
  - **tftp**—Trivial File Transfer Protocol
- **node**—(Optional) For chassis cluster configurations, clear sessions for applications on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.

Copyright © 2016, Juniper Networks, Inc.

- **local** —Clear the local node.
- **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation
- show security flow session application on page 459

List of Sample Output
clear security flow session application dns on page 423
clear security flow session application dns node 0 on page 423

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session application dns

```
user@host> clear security flow session application dns
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session application dns node 0

```
user@host> clear security flow session application dns node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

show chassis cluster data-plane statistics

Syntax **show chassis cluster data-plane statistics**

Release Information Command introduced in Junos OS Release 9.3.

Description Display information about chassis cluster data plane statistics.

Required Privilege Level view

Related Documentation
- *clear chassis cluster data-plane statistics*

List of Sample Output show chassis cluster data-plane statistics on page 425

Output Fields Table 16 lists the output fields for the **show chassis cluster data-plane statistics** command. Output fields are listed in the approximate order in which they appear.

Table 16: show chassis cluster data-plane statistics Output Fields

Field Name	Field Description
Services Synchronized	• **Service name**—Name of the service. • **Rtos sent**—Number of runtime objects (RTOs) sent. • **Rtos received**—Number of RTOs received. • **Translation context**—Messages synchronizing Network Address Translation (NAT) translation context. • **Incoming NAT**—Messages synchronizing incoming Network Address Translation (NAT) service. • **Resource manager**—Messages synchronizing resource manager groups and resources. • **Session create**—Messages synchronizing session creation. • **Session close**—Messages synchronizing session close. • **Session change**—Messages synchronizing session change. • **Gate create**—Messages synchronizing creation of pinholes (temporary openings in the firewall). • **Session ageout refresh request**—Messages synchronizing request session after age-out. • **Session ageout refresh reply**—Messages synchronizing reply session after age-out. • **IPsec VPN**—Messages synchronizing VPN session. • **Firewall user authentication**—Messages synchronizing firewall user authentication session. • **MGCP ALG**—Messages synchronizing MGCP ALG sessions. • **H323 ALG**—Messages synchronizing H.323 ALG sessions. • **SIP ALG**—Messages synchronizing SIP ALG sessions. • **SCCP ALG**—Messages synchronizing SCCP ALG sessions. • **PPTP ALG**—Messages synchronizing PPTP ALG sessions. • **RTSP ALG**—Messages synchronizing RTSP ALG sessions.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show chassis cluster data-plane statistics

```
user@host> show chassis cluster data-plane statistics
Services Synchronized:
    Service name                              RTOs sent    RTOs received
    Translation context                       0            0
    Incoming NAT                              0            0
    Resource manager                          0            0
    Session create                            0            0
    Session close                             0            0
    Session change                            0            0
    Gate create                               0            0
    Session ageout refresh requests           0            0
    Session ageout refresh replies            0            0
    IPsec VPN                                 0            0
    Firewall user authentication              0            0
    MGCP ALG                                  0            0
    H323 ALG                                  0            0
    SIP ALG                                   0            0
    SCCP ALG                                  0            0
    PPTP ALG                                  0            0
    RTSP ALG                                  0            0
```

Copyright © 2016, Juniper Networks, Inc.

show chassis cluster statistics

Syntax **show chassis cluster statistics**

Release Information Command modified in Junos OS Release 9.0. Output changed to support dual control ports in Junos OS Release 10.0.

Description Display information about chassis cluster services and interfaces.

Required Privilege Level view

Related Documentation
- *clear chassis cluster statistics*

List of Sample Output
show chassis cluster statistics on page 427
show chassis cluster statistics (SRX5000 line devices) on page 428
show chassis cluster statistics (SRX5000 line devices) on page 429

Output Fields Table 17 lists the output fields for the **show chassis cluster statistics** command. Output fields are listed in the approximate order in which they appear.

Table 17: show chassis cluster statistics Output Fields

Field Name	Field Description
Control link statistics	Statistics of the control link used by chassis cluster traffic. Statistics for **Control link 1** are displayed when you use dual control links (SRX5000 lines only). Note that the output for the SRX5000 lines will always show **Control link 0** and **Control link 1** statistics, even though only one control link is active or working. • **Heartbeat packets sent**—Number of heartbeat messages sent on the control link. • **Heartbeat packets received**—Number of heartbeat messages received on the control link. • **Heartbeat packet errors**—Number of heartbeat packets received with errors on the control link.
Fabric link statistics	Statistics of the fabric link used by chassis cluster traffic. Statistics for **Child Link 1** are displayed when you use dual fabric links. • **Probes sent**—Number of probes sent on the fabric link. • **Probes received**—Number of probes received on the fabric link.

Copyright © 2016, Juniper Networks, Inc.

Table 17: show chassis cluster statistics Output Fields *(continued)*

Field Name	Field Description
Services Synchronized	• **Service name**—Name of the service. • **Rtos sent**—Number of runtime objects (RTOs) sent. • **Rtos received**—Number of RTOs received. • **Translation context**—Messages synchronizing Network Address Translation (NAT) translation context. • **Incoming NAT**—Messages synchronizing incoming Network Address Translation (NAT) service. • **Resource manager**—Messages synchronizing resource manager groups and resources. • **Session create**—Messages synchronizing session creation. • **Session close**—Messages synchronizing session close. • **Session change**—Messages synchronizing session change. • **Gate create**—Messages synchronizing creation of pinholes (temporary openings in the firewall). • **Session ageout refresh request**—Messages synchronizing request session after age-out. • **Session ageout refresh reply**—Messages synchronizing reply session after age-out. • **IPsec VPN**—Messages synchronizing VPN session. • **Firewall user authentication**—Messages synchronizing firewall user authentication session. • **MGCP ALG**—Messages synchronizing MGCP ALG sessions. • **H323 ALG**—Messages synchronizing H.323 ALG sessions. • **SIP ALG**—Messages synchronizing SIP ALG sessions. • **SCCP ALG**—Messages synchronizing SCCP ALG sessions. • **PPTP ALG**—Messages synchronizing PPTP ALG sessions. • **RTSP ALG**—Messages synchronizing RTSP ALG sessions. • **MAC address learning**—Messages synchronizing MAC address learning.

Sample Output

show chassis cluster statistics

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 798
        Heartbeat packets received: 784
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 793
        Probes received: 0
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          0            0
    Session close                           0            0
    Session change                          0            0
    Gate create                             0            0
```

Copyright © 2016, Juniper Networks, Inc.

```
    Session ageout refresh requests          0            0
    Session ageout refresh replies           0            0
    IPsec VPN                                0            0
    Firewall user authentication             0            0
    MGCP ALG                                 0            0
    H323 ALG                                 0            0
    SIP ALG                                  0            0
    SCCP ALG                                 0            0
    PPTP ALG                                 0            0
    RTSP ALG                                 0            0
    MAC address learning                     0            0
```

Sample Output

show chassis cluster statistics (SRX5000 line devices)

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
        Heartbeat packets errors: 0
    Control link 1:
        Heartbeat packets sent: 258689
        Heartbeat packets received: 258684
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 258681
        Probes received: 258681
    Child link 1
        Probes sent: 258501
        Probes received: 258501
Services Synchronized:
    Service name                             RTOs sent    RTOs received
    Translation context                      0            0
    Incoming NAT                             0            0
    Resource manager                         0            0
    Session create                           1            0
    Session close                            1            0
    Session change                           0            0
    Gate create                              0            0
    Session ageout refresh requests          0            0
    Session ageout refresh replies           0            0
    IPSec VPN                                0            0
    Firewall user authentication             0            0
    MGCP ALG                                 0            0
    H323 ALG                                 0            0
    SIP ALG                                  0            0
    SCCP ALG                                 0            0
    PPTP ALG                                 0            0
    RPC ALG                                  0            0
    RTSP ALG                                 0            0
    RAS ALG                                  0            0
    MAC address learning                     0            0
    GPRS GTP                                 0            0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show chassis cluster statistics (SRX5000 line devices)

```
user@host> show chassis cluster statistics
Control link statistics:
    Control link 0:
        Heartbeat packets sent: 82371
        Heartbeat packets received: 82321
        Heartbeat packets errors: 0
    Control link 1:
        Heartbeat packets sent: 0
        Heartbeat packets received: 0
        Heartbeat packets errors: 0
Fabric link statistics:
    Child link 0
        Probes sent: 258681
        Probes received: 258681
    Child link 1
        Probes sent: 258501
        Probes received: 258501
Services Synchronized:
    Service name                            RTOs sent    RTOs received
    Translation context                     0            0
    Incoming NAT                            0            0
    Resource manager                        0            0
    Session create                          1            0
    Session close                           1            0
    Session change                          0            0
    Gate create                             0            0
    Session ageout refresh requests         0            0
    Session ageout refresh replies          0            0
    IPSec VPN                               0            0
    Firewall user authentication            0            0
    MGCP ALG                                0            0
    H323 ALG                                0            0
    SIP ALG                                 0            0
    SCCP ALG                                0            0
    PPTP ALG                                0            0
    RPC ALG                                 0            0
    RTSP ALG                                0            0
    RAS ALG                                 0            0
    MAC address learning                    0            0
    GPRS GTP                                0            0
```

Copyright © 2016, Juniper Networks, Inc.

show security alg h323 counters

Syntax **show security alg h323 counters**
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Display H.323 Application Layer Gateway (ALG) counters information.

Options

- none—Display H.323 ALG counters. information.
- **node**—(Optional) For chassis cluster configurations, display H.323 counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- h323 on page 360
- clear security alg h323 counters on page 416

List of Sample Output show security alg h323 counters on page 431

Output Fields Table 18 lists the output fields for the **show security alg h323 counters** command. Output fields are listed in the approximate order in which they appear.

Table 18: show security alg h323 counters Output Fields

Field Name	Field Description
Packets received	Number of H.323 ALG packets received.
Packets dropped	Number of H.323 ALG packets dropped.
RAS message received	Number of incoming RAS (Endpoint Registration, Admission, and Status) messages per second per gatekeeper received and processed.
Q.931 message received	Counter for Q.931 message received.
H.245 message received	Counter for H.245 message received.

Copyright © 2016, Juniper Networks, Inc.

Table 18: show security alg h323 counters Output Fields *(continued)*

Field Name	Field Description
Number of calls	Total number of H.323 ALG calls. NOTE: This counter displays the number of call legs and may not display the exact number of voice calls that are active. For instance, for a single active voice call between two endpoints, this counter might display a value of **2**.
Number of active calls	Number of active H.323 ALG calls.
Decoding errors	Number of decoding errors.
Message flood dropped	Error counter for message flood dropped.
NAT errors	H.323 ALG Network Address Translation (NAT) errors.
Resource manager errors	H.323 ALG resource manager errors.

Sample Output

show security alg h323 counters

```
user@host> show security alg h323 counters
H.323 counters summary:
Packets received           :4060
Packets dropped            :24
RAS message received       :3690
Q.931 message received       :202
H.245 message received       :145
Number of calls        :25
Number of active calls       :0
H.323 Error Counters:
Decoding errors            :24
Message flood dropped       :0
NAT errors           :0
Resource manager errors       :0
H.323 Message Counters:
RRQ      : 431   RCF      :  49 ARQ  : 60  ACF      : 33
URQ      : 34   UCF      : 25 DRQ  : 55  DCF      : 44
oth RAS  : 2942  Setup    : 28 Alert    : 9 Connect    : 25
CallPrcd  : 18  Info    : 0  RelCmpl  : 39  Facility   : 14
Progress  : 0  Empty    : 65  OLC      : 20 OLC-ACK  :   20
oth H245  : 16
```

Copyright © 2016, Juniper Networks, Inc.

show security alg ike-esp-nat summary

Syntax **show security alg ike-esp-nat summary**

Release Information Command introduced in Junos OS Release 10.2.

Description Display Application Layer Gateway (ALG) for IKE and ESP information summary.

Required Privilege Level view

Related Documentation
- clear security alg ike-esp-nat on page 417

List of Sample Output show security alg ike-esp-nat summary on page 432

Sample Output

show security alg ike-esp-nat summary

```
user@host> security alg ike-esp-nat summary
Initiator cookie: d5732d9b4114de1a
Responder cookie: 4776fe31164ef
Session-ID:       13
ALG state :       1
Timeout:          6292
Used IKE cookies: 0
Maximum IKE cookies: 2400
```


Copyright © 2016, Juniper Networks, Inc.

show security alg msrpc

Syntax

show security alg msrpc
<object-id-map> <node (*node-id* | all | local | primary)>

Release Information

Command modified in Junos OS Release 10.1.

Description

Display Microsoft (MS) remote procedure call (RPC) Application Layer Gateway (ALG) information.

Options

- none—Display all MS-RPC ID (UUID) to object ID (OID) table information.
- **object-id-map**—(Optional) Display information from the MS-RPC ID (UUID) to object ID (OID) table.
- **node**—(Optional) For chassis cluster configurations, display MS-RPC UUID-to-object-ID mapping information on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- msrpc on page 370

List of Sample Output

show security alg msrpc object-id-map on page 434

Output Fields

Copyright © 2016, Juniper Networks, Inc.

Table 19 lists the output fields for the **show security alg msrpc** command. Output fields are listed in the approximate order in which they appear.

Table 19: show security alg msrpc Output Fields

Field Name	Field Description
UUID	MS-RPC ID.
OID	MS-RPC object ID.

Sample Output

show security alg msrpc object-id-map

```
user@host> show security alg msrpc object-id-map
UUID                                    OID
1be617c0-31a5-11cf-a7d8-00805f48a135    0x80000020
e3514235-4b06-11d1-ab04-00c04fc2dcd2    0x80000002
67df7c70-0f04-11ce-b13f-00aa003bac6c    0x80000014
```


Copyright © 2016, Juniper Networks, Inc.

show security alg sccp calls

Syntax

show security alg sccp calls
< **brief** | **detail** | node (*node-id* | all | local | primary) >

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Display information about Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) calls.

Options

- none | **brief**—Display brief call information.
- **detail**—(Optional) Display detailed call information.
- **node**—(Optional) For chassis cluster configurations, display SCCP calls on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- sccp on page 384
- clear security alg sccp calls on page 418

List of Sample Output

show security alg sccp calls on page 436
show security alg sccp calls detail on page 436

Output Fields

Table 20 lists the output fields for the **show security alg sccp calls** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 20: show security alg sccp calls Output Fields

Field Name	Field Description
Client IP address	IP address of the client.
Client zone	Client zone ID.
Call manager	IP address of the call manager.
Resource manager group	Resource manager group ID.

Sample Output

show security alg sccp calls

```
user@host> show security alg sccp calls
Client IP          Zone      Call Manager       Conference ID  RM group
11.0.102.91        7         13.0.99.226        16789504       2047
12.0.102.96        8         13.0.99.226        16789505       2048
```

Sample Output

show security alg sccp calls detail

```
user@host> show security alg sccp calls detail
Client IP address: 11.0.102.91
Client zone: 7
Call Manager IP: 13.0.99.226
Conference ID: 16789504
Resource manager group: 2048
SCCP channel information:
    Media transmit channel address (IP address/Port): 0.0.0.0:0
    Media transmit channel translated address (IP address/Port): 0.0.0.0:0
    Media transmit channel pass-through party ID (PPID): 0
    Media transmit channel resource ID: 0
    Media receive channel address (IP address/Port): 11.0.102.91:20060
    Media receive channel translated address (IP address/Port): 25.0.0.1:1032
    Media receive channel pass-through party ID (PPID): 16934451
    Media receive channel resource ID: 8185
    Multimedia transmit channel address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia transmit channel pass-through party ID (PPID): 0
    Multimedia transmit channel resource ID: 0
    Multimedia receive channel address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel translated address (IP address/Port): 0.0.0.0:0
    Multimedia receive channel pass-through party ID (PPID): 0
    Multimedia receive channel resource ID: 0
Total number of calls = 1
```


Copyright © 2016, Juniper Networks, Inc.

show security alg sccp counters

Syntax **show security alg sccp counters**
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Display information about Skinny Client Control Protocol (SCCP) Application Layer Gateway (ALG) counters.

Options
- none—Display all SCCP ALG counters.
- **node**—(Optional) For chassis cluster configurations, display SCCP counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- sccp on page 384
- clear security alg sccp counters on page 419

List of Sample Output show security alg sccp counters on page 438

Output Fields Table 21 lists the output fields for the **show security alg sccp counters** command. Output fields are listed in the approximate order in which they appear.

Table 21: show security alg sccp counters Output Fields

Field Name	Field Description
Active client sessions	Number of active SCCP ALG client sessions.
Active calls	Number of active SCCP ALG calls.
Total calls	Total number of SCCP ALG calls.
Packets received	Number of SCCP ALG packets received.
PDUs processed	Number of SCCP ALG protocol data units (PDU) processed.
Current call rate	Number of calls per second.
Packets dropped	Number of packets dropped by the SCCP ALG.

Copyright © 2016, Juniper Networks, Inc.

Table 21: show security alg sccp counters Output Fields *(continued)*

Field Name	Field Description
Decode errors	Number of decoding errors.
Protocol errors	Number of protocol errors.
Address translation errors	Number of NAT errors.
Policy lookup errors	Number of errors occurring during policy lookups.
Unknown PDUs	Number of unknown protocol data units (PDUs).
Maximum calls exceed	Number of times the maximum number of calls was exceeded.
Maximum call rate exceed	Number of times the maximum call rate was exceeded.
Initialization errors	Number of call initialization errors.
Internal errors	Number of internal errors.
Nonspecific error	Number of nonspecific errors.

Sample Output

show security alg sccp counters

```
user@host> show security alg sccp counters
SCCP call statistics:
  Active client sessions         : 4
  Active calls                   : 2
  Total calls                    : 3
  Packets received               : 232
  PDUs processed                 : 232
  Current call rate              : 0
Error counters:
  Packets dropped                : 0
  Decode errors                  : 0
  Protocol errors                : 0
  Address translation errors     : 0
  Policy lookup errors           : 0
  Unknown PDUs                   : 0
  Maximum calls exceeded         : 0
  Maximum call rate exceeded     : 0
  Initialization errors          : 0
  Internal errors                : 0
  Nonspecific error              : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security alg sip calls

Syntax
show security alg sip calls
<brief | detail | node (*node-id* | all | local | primary) >

Release Information
Command modified in Junos OS Release 9.2; **node** options added in Junos OS Release 9.0.

Description
Display information about Session Initiation Protocols (SIP) Application Layer Gateway (ALG) calls.

Options
- none | **brief**—Display brief call information.
- **detail**—(Optional) Display detailed information about SIP ALG calls.
- **node**—(Optional) For chassis cluster configurations, display SIP calls on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level
view

Related Documentation
- sip (Security) on page 387
- clear security alg sip calls on page 420

List of Sample Output
show security alg sip calls on page 439
show security alg sip calls detail on page 440

Output Fields
Table 22 lists the output fields for the **show security alg sip calls** command. Output fields are listed in the approximate order in which they appear.

Table 22: show security alg sip calls Output Fields

Field Name	Field Description
UAS callid	Call Identifier for the SIP ALG user agent server.
State	State of the SIP ALG user agent server.

Sample Output

show security alg sip calls

```
user@host> show security alg sip calls
```

```
Total number of calls: 1
     Call ID  : 1-2468@20.0.0.174
   Method   : INVITE
```

Sample Output

show security alg sip calls detail

```
user@host> show security alg sip calls detail
Total number of calls: 1
       Call ID  : 1-2468@20.0.0.174
     Method   : INVITE
     State    : SETUP
     Group ID : 307
```


Copyright © 2016, Juniper Networks, Inc.

show security alg sip counters

Syntax **show security alg sip counters**
<node (*node-id* | all | local | primary)>

Release Information Command modified in Junos OS Release 9.2; **node** options added in Junos OS Release 9.0.

Description Display information about Session Initiation Protocol (SIP) Application Layer Gateway (ALG) counters.

Options

- none—Display all SIP ALG counters.
- **node**—(Optional) For chassis cluster configurations, display SIP counters on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- sip (Security) on page 387
- clear security alg sip counters on page 421

List of Sample Output show security alg sip counters on page 443

Output Fields Table 23 lists the output fields for the **show security alg sip counters** command. Output fields are listed in the approximate order in which they appear.

Table 23: show security alg sip counters Output Fields

Field Name	Field Description
INVITE	Number of INVITE requests sent. An INVITE request is sent to invite another user to participate in a session.
CANCEL	Number of CANCEL requests sent. A user can send a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE had sent a final response for the INVITE before it received the CANCEL
ACK	Number of ACK requests sent. The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE request.
BYE	Number of BYE requests sent. A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session.
RR header exceeded max	Number of times the SIP ALG RR (Record-Route) headers exceeded the maximum limit.

Copyright © 2016, Juniper Networks, Inc.

Table 23: show security alg sip counters Output Fields *(continued)*

Field Name	Field Description
REGISTER	Number of REGISTER requests sent. A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user.
OPTIONS	Number of OPTIONS requests sent. An OPTION request is used by the User Agent (UA) to obtain information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports.
INFO	Number of INFO requests sent. AN INFO message is used to communicate mid-session signaling information along the signaling path for the call.
MESSAGE	Number of MESSAGE requests sent. SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call).
NOTIFY	Number of NOTIFY requests sent. NOTIFY requests are sent to inform subscribers of changes in state to which the subscriber has a subscription.
PRACK	Number of PRACK requests sent. The PRACK request plays the same role as ACK, but for provisional responses.
PUBLISH	Number of PUBLISH requests sent. The PUBLISH request used for publishing event state. PUBLISH is similar to REGISTER in that it allows a user to create, modify, and remove state in another entity which manages this state on behalf of the user.
REFER	Number of REFER requests sent. A REFER request is used to refer the recipient (identified by the Request-URI) to a third party by the contact information provided in the request.
SUBSCRIBE	Number of SUBSCRIBE requests sent. A SUBCRIBE request is used to request current state and state updates from a remote node.
UPDATE	Number of UPDATE requests sent. AN UPDATE request is used to create a temporary opening in the firewall (pinhole) for new or updates Session Description Protocol (SDP) information. The following fields are modified: Via, From, To, Call-ID, Contact, Route, and Record-Route.
BENOTIFY	Number of BENOTIFY requests sent. A BENOTIFY request is used to reduce the unnecessary SIP signaling traffic on application servers. Applications that do not need a response for a NOTIFY request can enhance performance by enabling BENOTIFY.
SERVICE	Number of SERVICE requests sent. The SERVICE method used by a SIP client to request a service of a SIP server. It is a standard SIP message and will be forwarded until it reaches the server or end user which is performing the service.
OTHER	Number of OTHER requests sent.
Total Pkt-in	Number of SIP ALG total packets received.

Copyright © 2016, Juniper Networks, Inc.

Table 23: show security alg sip counters Output Fields *(continued)*

Field Name	Field Description
Total Pkt dropped on error	Number of SIP ALG total packets dropped while transmission and retransmission of messages.
Call error	Number of SIP ALG call errors.
IP resolve error	Number of SIP ALG IP address resolution errors.
NAT error	Number of SIP ALG NAT errors.
Resource manager error	Number of SIP ALG resource manager errors.
Contact header exceeded max	Number of times the SIP ALG contact headers exceeded the maximum limit.
Invite Dropped due to call limit	Number of SIP ALG invite dropped due to call limits.
SIP msg not processed by stack	Number of SIP ALG stack errors.
SIP msg not processed by alg	Number of SIP ALG messages not processed by ALGs.
SIP unknown method dropped	Number of SIP ALG unknown method errors.
Decoding error	Number of SIP ALG decoding errors.
Request for disconnected call	Number of SIP ALG calls disconnected.
Request out of state	Number of SIP ALG messages out of state errors.

Sample Output

show security alg sip counters

```
user@host> show security alg sip counters
SIP message counters(T:Transmit, RT:Retransmit):
   Method      T    1xx    2xx    3xx   4xx    5xx    6xx
               RT   RT     RT     RT    RT     RT     RT
   INVITE      4    4      3      0     0      0      0
               0    0      0      0     0      0      0
   CANCEL      0    0      0      0     0      0      0
               0    0      0      0     0      0      0
      ACK      3    0      0      0     0      0      0
               0    0      0      0     0      0      0
      BYE      3    0      3      0     0      0      0
               0    0      0      0     0      0      0
 REGISTER      7    0      7      0     0      0      0
               0    0      0      0     0      0      0
  OPTIONS      0    0      0      0     0      0      0
               0    0      0      0     0      0      0
     INFO      0    0      0      0     0      0      0
               0    0      0      0     0      0      0
  MESSAGE      0    0      0      0     0      0      0
```

```
                0    0     0     0    0     0     0
    NOTIFY      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     PRACK      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   PUBLISH      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     REFER      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
 SUBSCRIBE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
    UPDATE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
  BENOTIFY      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
   SERVICE      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
     OTHER      0    0     0     0    0     0     0
                0    0     0     0    0     0     0
SIP Error Counters:
Total Pkt-in                        :34
Total Pkt dropped on error          :0
Call error                          :0
IP resolve error                    :0
NAT error                           :0
Resource manager error              :0
RR header exceeded max              :0
Contact header exceeded max         :0
Call Dropped due to limit                  : 0
SIP stack error                                : 0
SIP decode error                             : 0
SIP unknown method error               : 0
RTO message sent                           : 0
RTO message received                     : 0
RTO buffer allocation failure              : 0
RTO buffer transmit failure                : 0
RTO send processing error                : 0
RTO receive processing error             : 0
RTO receive invalid length                 : 0
RTO receive call process error           : 0
RTO receive call allocation error        : 0
RTO receive call register error            : 0
RTO receive invalid status error         : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security alg sip rate

Syntax **show security alg sip rate**
<node (***node-id*** | all | local | primary)>

Release Information Command modified in Junos OS Release 9.2; **node** options added in Junos OS Release 9.0.

Description Display rate information for Session Initiation Protocol (SIP) Application Layer Protocol (ALG) messages.

Options
- none—Display all SIP ALG rate information.
- **node**—(Optional) For chassis cluster configurations, display SIP rate on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- sip (Security) on page 387

List of Sample Output show security alg sip rate on page 445

Output Fields Table 24 lists the output fields for the **show security alg sip rate** command. Output fields are listed in the approximate order in which they appear.

Table 24: show security alg sip rate Output Fields

Field Name	Field Description
CPU ticks	SIP ALG CPU ticks per microsecond.
Time taken	Time, in microseconds, that the last SIP ALG message needed to transit the network.
Total time	Total time, in microseconds, during an interval of less than 10 minutes for the specified number of SIP ALG messages to transit the network.
Rate	Number of SIP ALG messages per second transiting the network.

Sample Output

show security alg sip rate

```
user@host> show security alg sip rate
```

```
CPU ticks per us is 166
Time taken for the last message is 1103 us
Total time taken for 3124 messages is 6221482 us (in less than 10 minutes)
Rate: 502 messages/second
```


Copyright © 2016, Juniper Networks, Inc.

show security alg status

Syntax **show security alg status**

Release Information Command modified in Junos OS Release 9.5.

Description Display the status (enabled/disabled) of the supported Application Layer Gateway (ALG) transactions.

The following list describes the default status on each of these devices:

- SRX1500, SRX5400, SRX5600, and SRX5800 devices—FTP, TFTP, DNS, MS-RPC, PPTP, SUNRPC, SQL, and TALK ALGs are enabled by default. All other ALGs are disabled.
- SRX300, SRX320, SRX340, and SRX1500 devices—All supported ALGs are enabled by default.
- On all SRX Series devices — The RSH ALG is disabled by default.

Options none—Display status of all supported ALGs.

Required Privilege Level view

Related Documentation
- alg on page 327

Output Fields The following list describes the output fields for the **show security alg status** command. Output fields are listed in the approximate order in which they appear.

- **DNS**—Domain Name Server
- **FTP**—File Transfer Protocol
- **H323**—H.323 protocol
- **MGCP**—Media Gateway Control Protocol
- **MSRPC**—Microsoft remote procedure call
- **PPTP**—Point-to-Point Tunneling Protocol
- **RSH**—UNIX remote shell services
- **RTSP**—Real-Time Streaming Protocol
- **SCCP**—Skinny Client Control Protocol
- **SIP**—Session Initiation Protocol
- **SQL**—Oracle SQL
- **SUNRPC**—Sun Microsystems remote procedure call
- **TALK**—TALK program

Copyright © 2016, Juniper Networks, Inc.

- **TFTP—**Trivial File Transfer Protocol
- **IKE-ESP—**Internet Key Exchange and Encapsulating Security Payload

Sample Output

show security alg status

```
user@host> show security alg status
ALG Status :

  DNS      : Enabled

  FTP      : Enabled

  H323     : Disabled

  MGCP     : Disabled

  MSRPC    : Enabled

  PPTP     : Enabled

  RSH      : Disabled

  RTSP     : Disabled

  SCCP     : Disabled

  SIP      : Disabled

  SQL      : Enabled

  SUNRPC   : Enabled

  TALK     : Enabled

  TFTP     : Enabled

  IKE-ESP  : Disabled
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate

Syntax

show security flow gate
[<filter>] [brief | summary}

Release Information

Command introduced in Junos OS Release 8.5; Filter and display options added in Junos OS Release 10.2.

Description

Display information about temporary openings known as pinholes or gates in the security firewall.

Pinholes are used by applications that commonly have both control and data sessions and must create openings in the firewall for the data sessions based on information from the parent sessions.

Options

- destination-port—Destination port
- destination-prefix—Destination IP prefix or address
- protocol—IP protocol number
- source-port—Source port
- source-prefix—Source IP prefix or address
- brief | summary—Display the specified level of output.

Required Privilege Level

view

Related Documentation

- show security flow gate brief node on page 2059
- show security flow gate destination-port on page 2065
- show security flow gate destination-prefix on page 2068
- show security flow gate protocol on page 2071
- show security flow gate summary node on page 2074

List of Sample Output

show security flow gate on page 450
show security flow gate brief on page 451
show security flow gate summary on page 452

Output Fields

Table 25 lists the output fields for the **show security flow gate** command. Output fields are listed in the approximate order in which they appear.

Table 25: show security flow gate Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 25: show security flow gate Output Fields *(continued)*

Field Name	Field Description
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Maximum gates	Number of maximum gates

Sample Output

show security flow gate

```
user@host> show security flow gate
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.198.51.100.252/64515-64515
Translated: 0.0.0.0/0->10.0.31.161/25415
Protocol: udp
Application: none/0
Age: 101 seconds
Flags: 0xe001
Zone: untrust
Reference count: 1
Resource: 5-1024-8185
ole: 0.0.0.0-0.0.0.0/0-198.51.100.252/1046-1046
Translated: 198.51.100.252/36039-> 203.0.113.1/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
```


Copyright © 2016, Juniper Networks, Inc.

```
Flags: 0xe200
Zone: untrust
Reference count: 1
Resource: 5-1024-8189
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/24101-24101
Translated: 0.0.0.0/0-> 198.51.100.252/24101
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
Zone: trust
Reference count: 1
Resource: 5-1024-8188
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.5-198.51.100.252/24100-24100
Translated: 0.0.0.0/0->198.51.100.252/24100
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
Zone: trust
Reference count: 1
Resource: 5-1024-8191
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/5060-5060
Translated: 0.0.0.0/0->198.51.100.252/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
Flags: 0xe200
Zone: trust
Reference count: 1
Resource: 5-1024-8190
```

show security flow gate brief

```
root> show security flow gate brief
Flow Gates on FPC4 PIC1:

Hole: 192.0.2.1-192.0.2.1/0->192.0.2.100-192.0.2.100/38143-38143
  Translated: 192.0.2.1->192.0.2.100/38143
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65532 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate summary

```
root> show security flow gate summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session

Syntax **show security flow session**
[*filter*] **[brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 8.5. Support for filter and view options added in Junos OS Release 10.2. Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2. Policy ID filter added in Junos OS Release 12.3X48-D10. Support for connection tag added in Junos OS Release 15.1X49-D40.

Description Display information about all currently active security sessions on the device.

NOTE: For the normal flow sessions, the **show security flow session** command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. Because of this, the output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific **show** command for examples of the filtered output.

  application—Predefined application name

  application-firewall—Application firewall enabled

  application-firewall-rule-set—Application firewall enabled with the specified rule set

  application-traffic-control—Application traffic control session

  application-traffic-control-rule-set—Application traffic control rule set name and rule name

  conn-tag—Session connection tag

  destination-port—Destination port

  destination-prefix—Destination IP prefix or address

  dynamic-application—Dynamic application

  dynamic-application-group—Dynamic application

  encrypted—Encrypted traffic

  extensive—Display detailed output

Copyright © 2016, Juniper Networks, Inc.

family—Display session by family

idp—IDP enabled sessions

interface—Name of incoming or outgoing interface

logical-system (all | *logical-system-name*)—Name of a specific logical system or **all** to display all logical systems

nat—Display sessions with network address translation

policy-id—Display session information based on policy ID; the range is 1 through 4,294,967,295

protocol—IP protocol number

resource-manager—Resource manager

root-logical-system—Display root logical system as default

security-intelligence—Display security intelligence sessions

services-offload—Display services offload sessions

session-identifier—Display session with specified session identifier

source-port—Source port

source-prefix—Source IP prefix

summary—Display output summary

tunnel—Tunnel sessions

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session all on page 1902

List of Sample Output
show security flow session on page 456
show security flow session brief on page 457
show security flow session extensive on page 457
show security flow session summary on page 457

Output Fields Table 26 lists the output fields for the **show security flow session** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 26: show security flow session Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Application traffic control rule-set	AppQoS rule set for this session.
Rule	AppQoS rule for this session.
Forwarding class	The AppQoS forwarding class name for this session that distinguishes the transmission priority
DSCP code point	Differentiated Services (DiffServ) code point (DSCP) value remarked by the matching rule for this session.

Copyright © 2016, Juniper Networks, Inc.

Table 26: show security flow session Output Fields *(continued)*

Field Name	Field Description
Loss priority	One of four priority levels set by the matching rule to control discarding a packet during periods of congestion. A high loss priority means a high probability that the packet could be dropped during a period of congestion.
Rate limiter client to server	The rate-limiter profile assigned to the client-to-server traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Rate limiter server to client	The rate-limiter profile assigned to the server-to-client traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Maximum number of sessions permitted.

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
 Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/10000/2000 --> 203.0.113.11/10000/1000;udp, Conn Tag: 0x0, If:
 reth0.0, Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session application

Syntax

show security flow session application
application-name **[brief | extensive | summary]**

Release Information

Command introduced in Junos OS Release 8.5. Filter and view options added in Junos OS Release 10.2.

Description

Display information about each session of the specified application type.

Options

- ***application-name***—Type of application about which to display sessions information. Possible values are:
  - dns—Domain Name System
  - ftp—File Transfer Protocol
  - ignore—Ignore application type
  - mgcp-ca—Media Gateway Control Protocol with Call Agent
  - mgcp-ua—MGCP with User Agent
  - pptp—Point-to-Point Tunneling Protocol
  - q931—ISDN connection control protocol
  - ras—Remote Access Server
  - realaudio—RealAudio
  - rsh—UNIX remote shell services
  - rtsp—Real-Time Streaming Protocol
  - sccp—Skinny Client Control Protocol
  - sip—Session Initiation Protocol
  - sqlnet-v2—Oracle SQLNET
  - talk—TALK program
  - tftp—Trivial File Transfer Protocol
- brief | extensive | summary—Display the specified level of output.

Required Privilege Level

view

Related Documentation

- clear security flow session application on page 422

List of Sample Output

show security flow session application telnet on page 461
show security flow session application telnet brief on page 461
show security flow session application telnet extensive on page 461
show security flow session application telnet summary on page 462

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 27 lists the output fields for the **show security flow session application** command. Output fields are listed in the approximate order in which they appear.

Table 27: show security flow session application Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 27: show security flow session application Output Fields *(continued)*

Field Name	Field Description
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session application telnet

```
root> show security flow session application telnet
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes:
610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes:
602
Total sessions: 1
```

show security flow session application telnet brief

```
root> show security flow session application telnet brief
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes:
610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes:
602
Total sessions: 1
```

show security flow session application telnet extensive

```
root>  show security flow session application telnet extensive
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Status: Normal
Flag: 0x40
Policy name: default-policy/2
Source NAT pool: Null, Application: junos-telnet/10
Maximum timeout: 1800, Current timeout: 1788
Session State: Valid
Start time: 670184, Duration: 33
   In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp,
    Interface: ge-0/0/2.0,
    Session token: 0x180, Flag: 0x0x21
    Route: 0x60010, Gateway: 203.0.113.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 10, Bytes: 610
   Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp,
    Interface: ge-0/0/1.0,
    Session token: 0x1c0, Flag: 0x0x20
    Route: 0x70010, Gateway: 192.0.2.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 602
Total sessions: 1
```

show security flow session application telnet summary

```
root> show security flow session application telnet summary
Flow Sessions on FPC4 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC5 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager

Syntax **show security flow session resource-manager**
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

Description Display information about sessions created by the resource manager.

Options none—Display all resource manager sessions.

brief | extensive | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session resource-manager on page 1912

List of Sample Output
show security flow session resource-manager on page 464
show security flow session resource-manager brief on page 465
show security flow session resource-manager extensive on page 465
show security flow session resource-manager summary on page 466

Output Fields Table 28 lists the output fields for the **show security flow session resource-manager** command. Output fields are listed in the approximate order in which they appear.

Table 28: show security flow session resource-manager Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
Resource information	Information about the session particular to the resource manager, including the name of the ALG, the group ID. and the resource ID.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 28: show security flow session resource-manager Output Fields *(continued)*

Field Name	Field Description
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.

Sample Output

show security flow session resource-manager

```
root> show security flow session resource-manager
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1734, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586,
 CP Session ID: 410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes:
803, CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager brief

```
root> show security flow session resource-manager brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1704, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586,
 CP Session ID: 410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes:
803, CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session resource-manager extensive

```
root> show security flow session resource-manager extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Status: Normal
Flags: 0x42/0x0/0x2010103
Policy name: p1/4
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1682
Session State: Valid
Start time: 160496, Duration: 153
Client: FTP ALG, Group: 1, Resource: 0
   In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0002621
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 586
    CP Session ID: 410001274
   Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0002620
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 803
    CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager summary

```
root> show security flow session resource-manager summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security idp policy-templates

Syntax show security idp policy-templates

Release Information Command introduced in Junos OS Release 10.1.

Description Display the list of available policy templates.

Required Privilege Level view

Related Documentation
- *show security idp active-policy*

Output Fields **user@host> show security idp policy-templates**

Sample Output

```
DMZ_Services
DNS_Service
File_Server
Getting_Started
IDP_Default
Recommended
Web_Server
```

show security resource-manager group active

Syntax

show security resource-manager group active
<group-number >
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Display security information about active groups created through the resource manager.

Options

- none—Display resource manager group service information for all active groups.
- ***group-number*** —(Optional) Display resource manager group service information for a specific group identification number.
- **node**—(Optional) For chassis cluster configurations, display active resource manager group service information on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output

show security resource-manager group active on page 469
show security resource-manager group active 2048 on page 469
show security resource-manager group active node primary on page 469
show security resource-manager group active node all on page 469
show security resource-manager group active 1024 node all on page 469

Output Fields

Table 29 lists the output fields for the **show security resource-manager group** command. Output fields are listed in the approximate order in which they appear.

Table 29: show security resource-manager group Output Fields

Field Name	Field Description
Total groups	Total number of groups in the system.
active groups	Number of active groups.
Group ID	Identification number whose group information is displayed.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security resource-manager group active

```
user@host> show security resource-manager group active
Total groups 32, active groups 0
```

Sample Output

show security resource-manager group active 2048

```
user@host> show security resource-manager group active 2048
Total groups 2048, active groups 1
Group ID 2048: state - Active
          : Virtual System      - root
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8190
               Resource ID - 8188
               Resource ID - 8187
```

Sample Output

show security resource-manager group active node primary

```
user@host> show security resource-manager group active node primary
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active node all

```
user@host> show security resource-manager group active node all
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
node1:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active 1024 node all

```
user@host> show security resource-manager group active 1024 node all
node0:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8192
               Resource ID - 8188
               Resource ID - 8187
```

Copyright © 2016, Juniper Networks, Inc.

```
node1:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8187
               Resource ID - 8186
               Resource ID - 8190
```


Copyright © 2016, Juniper Networks, Inc.

show security resource-manager resource active

Syntax

show security resource-manager resource **active**
<resource-id >
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Display security information about active resources created through the resource manager.

Options

- none—Display information for all active resources.
- ***resource-id*** —(Optional) Display information for a resource with a specific identification number.
- **node**—(Optional) For chassis cluster configurations, display active resource manager information on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output

show security resource-manager resource active on page 472
show security resource-manager resource active 5 on page 472
show security resource-manager resource active node local on page 472
show security resource-manager resource active node primary on page 472

Output Fields

Table 30 lists the output fields for the **show security resource-manager resource** command. Output fields are listed in the approximate order in which they appear.

Table 30: show security resource-manager resource Output Fields

Field Name	Field Description
Total resources	Total number of resources in the system.
active resources	Number of active resources.
Resource ID	Identification number whose resource information is displayed.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security resource-manager resource active

```
user@host> show security resource-manager resource active
    Resource ID 7: Group ID - 2, Application - JSF_sip

    Resource ID 6: Group ID - 2, Application - JSF_sip

    Resource ID 5: Group ID - 2, Application - JSF_sip

    Resource ID 4: Group ID - 2, Application - JSF_sip

    Resource ID 3: Group ID - 2, Application - JSF_sip

    Resource ID 1: Group ID - 2, Application - JSF_sip

    Resource ID 2: Group ID - 2, Application - JSF_sip
Total Resources 4326, active resources 7
```

Sample Output

show security resource-manager resource active 5

```
user@host> show security resource-manager resource active 5
Resource ID 5: state - Active
    Application         - asl_client
    Parent group        - 2
    Policy              - 5
    From zone           - untrust
    To zone             - trust
    Resource timeout    - 0
    Number of sessions  - 0
    Number of Holes     - 1
        Source IP range        - {0.0.0.0, 0.0.0.0}
        Source port range      - {0, 0}
        Destination IP range   - {33.1.0.200, 33.1.0.200}
        Destination port range - {5060, 5060}
        Translated             - {0.0.0.0/0 -> 33.1.0.200/5060}
        Protocol               - 17
        Reference count        - 1
```

Sample Output

show security resource-manager resource active node local

```
user@host> show security resource-manager resource active node local
node0:
--------------------------------------------------------------------------
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Sample Output

show security resource-manager resource active node primary

```
user@host> show security resource-manager resource active node primary
node0:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

```
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Copyright © 2016, Juniper Networks, Inc.

show security resource-manager summary

Syntax show security resource-manager summary

Release Information Command introduced in Junos OS Release 11.4.

Description Display summary information about active resources, clients, groups, and sessions created through the resource manager.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output show security resource-manager summary on page 474

Output Fields Table 31 lists the output fields for the **show security resource-manager summary** command. Output fields are listed in the approximate order in which they appear.

Table 31: show security resource-manager summary Output Fields

Field Name	Field Description
Active resource-manager clients	Number of active resource manager clients.
Active resource-manager groups	Number of active resource manager groups.
Active resource-manager resources	Number of active resource manager resources.
Active resource-manager sessions	Number of active resource manager sessions.

Sample Output

show security resource-manager summary

```
user@host> show security resource-manager summary

Active resource-manager clients   : 15
Active resource-manager groups    : 1
Active resource-manager resources : 1
Active resource-manager sessions  : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security zones

Syntax
show security zones
<detail | terse>
< *zone-name* >

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones.

Options
- none—Display information about all zones.
- **detail | terse**—(Optional) Display the specified level of output.
- ***zone-name*** —(Optional) Display information about the specified zone.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones on page 476
show security zones abc on page 476
show security zones abc detail on page 476
show security zones terse on page 477

Output Fields
Table 32 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 32: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                     Type
my-internal              Security
my-external              Security
dmz                      Security
```

show security zones type

Syntax
show security zones type
(functional | security)
<detail | terse>

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones of the specified type.

Options
- **functional**—Display functional zones.
- **security**—Display security zones.
- **detail | terse**—(Optional) Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones type functional on page 479
show security zones type security on page 479
show security zones type security terse on page 479
show security zones type security detail on page 479

Output Fields
Table 33 lists the output fields for the **show security zones type** command. Output fields are listed in the approximate order in which they appear.

Table 33: show security zones type Output Fields

Field Name	Field Description
Security zone	Zone name.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones type functional

```
user@host> show security zones type functional
Functional zone: management
  Description: management zone
  Policy configurable: No
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security

```
user@host> show security zones type security
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security terse

```
user@host> show security zones type security terse
Zone                    Type
trust                   Security
untrust                 Security
junos-host              Security
```

Sample Output

show security zones type security detail

```
user@host> show security zones type security detail
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```


Copyright © 2016, Juniper Networks, Inc.

AppSecure Services Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 21

Overview

- Understanding AppSecure Services on page 483

Understanding AppSecure Services

An individual can connect to the network using multiple devices simultaneously, making it impractical to identify a user, an application, or a device by a group of statically allocated IP addresses and port numbers. Junos OS application identification recognizes traffic at different network layers using characteristics other than port number.

Once the application is determined, AppSecure service modules can be configured to monitor and control traffic for tracking, prioritization, access control, detection, and prevention based on the application ID of the traffic.

- AppTrack—Tracks and reports applications passing through the device.
- AppFW—Implements an application firewall using application-based rules.
- AppQoS—Provides quality of service prioritization based on application awareness.
- Intrusion Detection and Prevention (IDP)—Applies appropriate attack objects to applications running on nonstandard ports. Application identification improves IDP performance by narrowing the scope of attack signatures for applications without decoders.

Related Documentation

- Understanding Application Identification Techniques on page 485

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 22

Understanding Application Identification

- Understanding Application Identification Techniques on page 485
- Understanding the Junos OS Application Identification Database on page 488

Understanding Application Identification Techniques

Historically, firewalls have used the IP address and port numbers as a way of enforcing policies. That strategy is based on the assumption that users connect to the network from fixed locations and access particular resources using specific port numbers.

Today, wireless networking and mobile devices require a different strategy. The way in which devices connect to the network changes rapidly. An individual can connect to the network using multiple devices simultaneously. It is no longer practical to identify a user, application, or device by a group of statically allocated IP addresses and port numbers.

- Junos OS Next-Generation Application Identification on page 485
- Application Signature Mapping on page 486
- Application Identification Match Sequence on page 486

Junos OS Next-Generation Application Identification

Next-generation application identification builds on the legacy application identification functionality and provides more effective detection capabilities for evasive applications such as Skype, BitTorrent, and Tor.

Junos OS application identification recognizes Web-based and other applications and protocols at different network layers using characteristics other than port number. Applications are identified by using a protocol bundle containing application signatures and parsing information. The identification is based on protocol parsing and decoding and session management.

The detection mechanism has its own data feed and constructs to identify applications.

The following features are supported in application identification:

- Support for protocols and applications, including video streaming, peer-to-peer communication, social networking, and messaging
- Identification of services within applications

Copyright © 2016, Juniper Networks, Inc.

- Ability to distinguish actions launched within an application (such as login, browse, chat, and file transfer)
- Support for all versions of protocols and application decoders and dynamic updates of decoders
- Support for encrypted and compressed traffic and most complex tunneling protocols
- Ability to identify all protocols from Layer 3 to Layer 7 and above Layer 7.

Application Signature Mapping

Application signature mapping is a precise method of identifying the application that issued traffic on the network. Signature mapping operates at Layer 7 and inspects the actual content of the payload.

Applications are identified by using a downloadable protocol bundle. Application signatures and parsing information of the first few packets are compared to the content of the database. If the payload contains the same information as an entry in the database, the application of the traffic is identified as the application mapped to that database entry.

Juniper Networks provides a predefined application identification database that contains entries for a comprehensive set of known applications, such as FTP and DNS, and applications that operate over the HTTP protocol, such as Facebook, Kazaa, and many instant messaging programs. A signature subscription allows you to download the database from Juniper Networks and regularly update the content as new predefined signatures are added.

Application Identification Match Sequence

Figure 28 shows the sequence in which mapping techniques are applied and how the application is determined.

Copyright © 2016, Juniper Networks, Inc.

Figure 28: Mapping Sequence

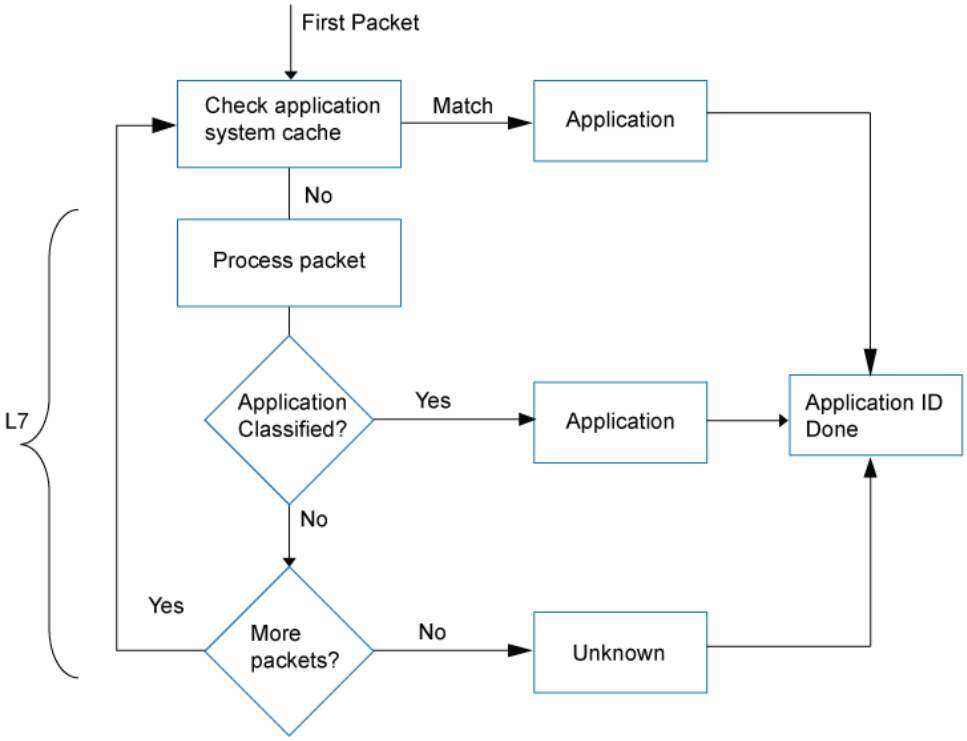

In application identification, every packet in the flow passes through the application identification engine for processing until the application is identified. Application bindings are saved in the application system cache (ASC) to expedite future identification process.

Application signatures identify an application based on protocol grammar analysis in the first few packets of a session. If the application identification engine has not yet identified the application, it passes the packets and waits for more data.

The application identification module matches applications for both client-to-server and server-to-client sessions.

Once the application is determined, AppSecure service modules can be configured to monitor and control traffic for tracking, prioritization, access control, detection, and prevention based on the application ID of the traffic.

- AppTrack—Tracks and reports applications passing through the device.
- Intrusion Detection and Prevention (IDP)—Applies appropriate attack objects to applications running on nonstandard ports. Application identification improves IDP performance by narrowing the scope of attack signatures for applications without decoders.
- AppFW—Implements an application firewall using application-based rules.
- AppQoS—Provides quality of service prioritization based on application awareness.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding AppTrack on page 573
- Application Firewall Overview on page 555
- *IDP Policies Overview*
- Understanding Application QoS (AppQoS) on page 583

Understanding the Junos OS Application Identification Database

A predefined signature database is available on the Juniper Networks Security Engineering website. This database includes a library of application signatures.

The predefined signature package provides identification criteria for known application signatures and is updated periodically.

Whenever new applications are added, the protocol bundle is updated and generated for all relevant platforms. It is packaged together with other application signature files. This package will be available for download through the security download website.

A subscription service allows you to regularly download the latest signatures for up-to-date coverage without having to create entries for your own use.

Application identification is enabled by default and is automatically turned on when you configure Intrusion Detection and Prevention (IDP), AppFW, AppQoS, or AppTrack.

NOTE: **Updates to the Junos OS predefined application signature package are authorized by a separately licensed subscription service. You must install the application identification application signature update license key on your device to download and install the signature database updates provided by Juniper Networks. When your license key expires, you can continue to use the locally stored application signature package contents but you cannot update the package.**

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- *Understanding IDP Application Identification*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 23

Installing Application Signature Package

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package on page 496
- Example: Scheduling the Application Signature Package Updates on page 499
- Scheduling the Application Signature Package Updates As Part of the IDP Security Package on page 501
- Example: Downloading and Installing the Application Identification Package in Chassis Cluster Mode on page 503
- Verifying the Junos OS Application Identification Extracted Application Package on page 506
- Uninstalling the Junos OS Application Identification Application Package on page 507
- Disabling and Reenabling Junos OS Application Identification on page 507

Understanding the Junos OS Application Package Installation

Juniper Networks regularly updates the predefined application signature package database and makes it available to subscribers on the Juniper Networks website. This package includes signature definitions of known application objects that can be used to identify applications for tracking, firewall policies, quality of service prioritization, and Intrusion Detection and Prevention (IDP). The database contains application objects such as FTP, DNS, Facebook, Kazaa, and many instant messenger programs.

You need to download and install the application signature package before configuring application services. The application signature package is included in the IDP installation directly and does not need to be downloaded separately.

- If you have IDP enabled and plan to use application identification, you can continue to run the IDP signature database download. To download the IDP signature database, run the following command: **request security idp security-package download**. The application package download can be performed manually or automatically. See

Copyright © 2016, Juniper Networks, Inc.

"Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package" on page 496.

NOTE: If you have an IDP-enabled device and plan to use application identification, we recommend that you download only the IDP signature database. This will avoid having two versions of the application database, which could become out of sync.

- If you do not have IDP enabled and plan to use application identification, you can run the following commands: **request services application-identification download** and **request services application-identification install**. These commands will download the application signature database and install it on the device.

  You can perform the download manually or automatically. When you download the extracted package manually, you can change the download URL.

  After downloading and installing the application signature package, use CLI commands to download and install database updates, and view summary and detailed application information.

  See "Downloading and Installing the Junos OS Application Signature Package Manually" on page 493 or "Example: Scheduling the Application Signature Package Updates" on page 499.

NOTE: The Junos OS application signature package update is a separately licensed subscription service. You must install the application signature package update license key on your device to download and install the signature database updates provided by Juniper Networks. If your license key expires, you can continue to use the locally stored application signature package content but you cannot update the data.

NOTE: On all SRX Series devices, J-Web pages for AppSecure Services are preliminary. We recommend using the CLI for configuration of AppSecure features.

Upgrading to Next-Generation Application Identification

You must install Junos OS Release 12.1X47-D10 to migrate from existing, or legacy, application identification to next-generation application identification.

SRX Series devices installed with Junos OS builds with legacy application identification include legacy application identification security packages. When you upgrade these devices with Junos OS Release 12.1X47-D10, the next-generation application identification security package is installed along with the default protocol bundle. The device is automatically upgraded to next-generation application identification.

Copyright © 2016, Juniper Networks, Inc.

NOTE:

- The next-generation application identification security package introduces incremental updates to the legacy application identification package. You are not required to remove or uninstall any existing applications.
- Applications supported in previous releases (Junos OS Release 12.1X46 or prior) might have new aliases or alternative names in the new version. So existing configurations using such application works in Junos OS Release 12.1X47; however, related logs and other information will use the new name. You can use the **show services application-identification application detail** *new-application-name* command to get the details of the applications.
- When you upgrade Junos OS, you can include the **validate** or **no-validate** options with the **request system software add** command. Because the existing features, which are not part of next-generation application identification are deprecated, however, incompatibility issues are not seen.
- Next-generation application identification eliminates the generation of new nested applications and treats existing nested applications as normal applications. In addition, next-generation application identification does not support custom applications or custom application groups. Existing configurations involving any nested applications, custom applications, or custom application groups are ignored with warning messages.

Related Documentation

- Understanding the Junos OS Application Identification Database on page 488
- *Understanding the IDP Signature Database*.
- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package on page 496
- Example: Scheduling the Application Signature Package Updates on page 499

Installing and Verifying Licenses for an Application Signature Package

The Junos OS application signature package update is a separately licensed subscription service. You must install the application signature package update license key on your device to download and install the signature database updates provided by Juniper Networks. If your license key expires, you can continue to use the locally stored application signature package content.

Licensing is usually ordered when the device is purchased, and this information is bound to the chassis serial number. These instructions assume that you already have the license. If you did not order the license during the purchase of the device, contact your account team or Juniper customer care for assistance. For more information, refer to the Knowledge Base article KB9731 at http://kb.juniper.net/InfoCenter/index?page=home.

Copyright © 2016, Juniper Networks, Inc.

NOTE: On SRX1500 devices, AppSecure is part of Juniper Networks Secure Edge software (a default shipping software package on the SRX1500). A separate license key is not required on your device to download and install the AppID signature database updates, or to use other AppSecure features such as AppFW, AppQoS, and AppTrack.

NOTE: On SRX300, SRX320, SRX340, and SRX345 devices, AppSecure is part of Juniper Networks Secure Edge software or IPS subscription license. A separate license key is not required on your device to download and install the AppID signature database updates, or to use other AppSecure features such as AppFW, AppQoS, and AppTrack.

You can install the license on the SRX Series device using either the automatic method or manual method as follows:

- Install your license automatically on the device.

  To install or update your license automatically, your device must be connected to the Internet .

  user@host> **request system license update**

  ```
  Trying to update license keys from https://ae1.juniper.net, use 'show system
  license' to check status.
  ```

- Install the licenses manually on the device.

  user@host> **request system license add terminal**

  ```
  [Type ^D at a new line to end input,
   enter blank line between each license key]
  ```

  Paste the license key and press Enter to continue.

- Verify the license is installed on your device.

  Use the **show system license command** command to view license usage, as shown in the following example:

  ```
  License usage:
                                  Licenses     Licenses    Licenses    Expiry
    Feature name                      used    installed      needed
    logical-system                       4            1           3    permanent


   License identifier: JUNOSXXXXXX
    License version: 2
    Valid for device: AA4XXXX005
    Features:
      appid-sig        - APPID Signature
        date-based, 2014-02-17 08:00:00 GMT-8 - 2015-02-11 08:00:00 GMT-8
  ```

  The output sample is truncated to display only license usage details.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package on page 496

Downloading and Installing the Junos OS Application Signature Package Manually

This example shows how to download the application signature package, create a policy, and identify it as the active policy.

- Requirements on page 493
- Overview on page 493
- Configuration on page 493
- Verification on page 495

Requirements

Before you begin:

- Ensure that your SRX Series device has a connection to the Internet to download security package updates.

  NOTE: **DNS must be set up because you need to resolve the name of the update server.**

- Ensure that you have installed the application identification feature license. See "Installing and Verifying Licenses for an Application Signature Package" on page 491.

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 12.1X47-D10

Overview

Juniper Networks regularly updates the predefined application signature package database and makes it available on the Juniper Networks website. This package includes application objects that can be used in Intrusion Detection and Prevention (IDP), application firewall policy, and AppTrack to match traffic.

Configuration

CLI Quick Configuration

CLI quick configuration is not available for this example because manual intervention is required during the configuration.

Copyright © 2016, Juniper Networks, Inc.

Downloading and Installing Application Identification

Step-by-Step Procedure

1. Download the application package.

   user@host> **request services application-identification download**

   ```
   Please use command "request services application-identification download
   status" to check status
   ```

   Download retrieves the application package from the Juniper Networks security website https://signatures.juniper.net/cgi-bin/index.cgi.

   You can also download a specific version of the application package or download the application package from the specific location by using the following options:

   - To download a specific version of the application package:

     user@host>**request services application-identification download version** ***version-number***

   - To change the download URL for the application package from configuration mode:

     [edit]
     user@host# **set services application-identification download url** ***URL or File Path***

   NOTE: If you change the download URL and you want to keep that change, make sure you commit the configuration.

2. Check the download status.

   user@host>**request services application-identification download status**

   ```
   Application package 2345 is downloaded successfully
   ```

   NOTE: You can also use the system log to view the result of the download.

3. Install the application package.

   user@host>**request services application-identification install**

   ```
   Please use command "request services application-identification install
   status" to check status and use command "request services
   application-identification proto-bundle-status" to check protocol bundle
   status
   ```

   The application package is installed in the application signature database on the device.

4. Check the installation status of the application package.

   The command output displays information about the downloaded and installed versions of the application package and protocol bundle.

Copyright © 2016, Juniper Networks, Inc.

- To view the installation status:

  user@host>**request services application-identification install status**

  ```
  Install application package 2345 succeed
  ```

- To view the protocol bundle status:

  user@host>**request services application-identification proto-bundle-status**

  ```
  Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and
  application secpack version (2345) is loaded and activated.
  ```

NOTE: It is possible that an application signature was removed from the newer version of an application signature database. If this signature is used in an existing application firewall policy on your device, the installation of the new database will fail. An installation status message identifies the signature that is no longer valid. To update the database successfully, remove all references to the deleted signature from your existing policies and groups, and rerun the install command.

Verification

Confirm that the configuration is working properly.

Verifying the Application Identification Status

Purpose Verify that the application identification configuration is working properly.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show services application-identification status** command.

```
pic: 1/0

Application Identification
 Status                          Enabled
 Sessions under app detection    0
 Engine Version                  4.18.1-20 (build date Jan 25 2014)
 Max TCP session packet memory   30000
 Max C2S bytes                   1024
 Max S2C bytes                   0
 Force packet plugin             Disabled
 Force stream plugin             Disabled
 Statistics collection interval  1 (in minutes)

Application System Cache
 Status                          Enabled
 Negative cache status           Disabled
 Max Number of entries in cache  131072
 Cache timeout in seconds        3600

Protocol Bundle
 Download Server              https://services.netscreen.com/cgi-bin/index.cgi

 AutoUpdate                      Enabled
Slot 1:
 Status                          Active
 Version                         1.30.4-22.005 (build date Jan 17 2014)
 Sessions                        0
Slot 2
 Status                          Free
```

Meaning The **Status: Enabled** field shows that application identification is enabled on the device.

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Example: Scheduling the Application Signature Package Updates on page 499
- Verifying the Junos OS Application Identification Extracted Application Package on page 506
- Uninstalling the Junos OS Application Identification Application Package on page 507

Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package

You can download and install application signatures through intrusion detection and prevention (IDP) security packages.

Copyright © 2016, Juniper Networks, Inc.

This example shows how to enhance security by downloading and installing the IDP signatures and application signature package. In this case, both IDP signature pack and application signature pack are downloaded with a single command.

- Requirements on page 497
- Overview on page 497
- Configuration on page 497
- Verification on page 498

Requirements

Before you begin:

- Ensure that your SRX Series device has a connection to the Internet to download security package updates.

  NOTE: **DNS must be set up because you need to resolve the name of the update server.**

- Ensure that you have installed the application identification feature license. See "Installing and Verifying Licenses for an Application Signature Package" on page 491.

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 12.1X47-D10

Overview

In this example, you download and install the signature database from the Juniper Networks website.

Configuration

Downloading and Installing the Signature Database

CLI Quick Configuration

CLI quick configuration is not available for this example because manual intervention is required during the configuration.

Step-by-Step Procedure

To download and install application signatures:

1. Download the signature database.

   ```
   [edit]
   user@host# run request security idp security-package download

   Will be processed in async mode. Check the status using the status checking
    CLI
   ```

Copyright © 2016, Juniper Networks, Inc.

NOTE: Downloading the database might take some time depending on the database size and the speed of your Internet connection.

2. Check the security package download status.

```
[edit]
user@host# run request security idp security-package download status
```

```
Done;Successfully downloaded
from(https://services.netscreen.com/cgi-bin/index.cgi).
Version info:2230(Mon Feb  4 19:40:13 2013 GMT-8, Detector=12.6.160121210)
```

3. Install the attack database.

```
[edit]
user@host# run request security idp security-package install
```

```
Will be processed in async mode. Check the status using the status checking
 CLI
```

NOTE: Installing the attack database might take some time depending on the security database size.

4. Check the attack database install status. The command output displays information about the downloaded and installed versions of the attack database.

```
[edit]
user@host# run request security idp security-package install status
```

```
Done;Attack DB update : successful - [UpdateNumber=2230,ExportDate=Mon Feb
 4 19:40:13 2013 GMT-8,Detector=12.6.160121210]
Updating control-plane with new detector : successful
Updating data-plane with new attack or detector : successful
```

5. Confirm your IDP security package version.

```
[edit]
user@host# run show security idp security-package-version
```

```
Attack database version:2230(Mon Feb  4 19:40:13 2013 GMT-8)
  Detector version :12.6.160121210
  Policy template version :2230
```

6. Confirm your application identification package version.

```
[edit]
user@host# run show services application-identification version
```

```
Application package version: 1884
```

Verification

Confirm that the application signature package is being updated properly.

Copyright © 2016, Juniper Networks, Inc.

Verifying application signature package

Purpose Verify services application identification version

Action From operational mode, enter the **show services application-identification version** command.

```
user@host> show services application-identification version
Application package version: 1884
```

Meaning The sample output shows that, the services application identification version is 1884.

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Verifying the Junos OS Application Identification Extracted Application Package on page 506
- Uninstalling the Junos OS Application Identification Application Package on page 507

Example: Scheduling the Application Signature Package Updates

This example shows how to set up automatic updates of the predefined application signature package.

- Requirements on page 499
- Overview on page 499
- Configuration on page 500
- Verification on page 500

Requirements

Before you begin:

- Ensure that your SRX Series device has a connection to the Internet to download security package updates.

  NOTE: DNS must be set up because you need to resolve the name of the update server.

- Ensure that you have installed the application identification feature license. See "Installing and Verifying Licenses for an Application Signature Package" on page 491.

Overview

In this example, you want to download the current version of the application signature package periodically. The download should start at 11:59 PM on December 10. To maintain

Copyright © 2016, Juniper Networks, Inc.

the most current information, you want to update the package automatically every 2 days from your company's intranet site.

Configuration

GUI Step-by-Step Procedure

To set up the automatic download and periodic update with the J-Web interface:

1. Enter **Configure>Security>AppSecure Settings** to display the Applications Signature page.
2. Click **Global Settings**.
3. Click the **Download Scheduler** tab, and modify the following fields:
   - URL: **https://signatures.juniper.net/cgi-bin/index.cgi**
   - Enable Schedule Update: Select the check box.
   - Interval: **48**
4. Click **Reset Setting** to clear the existing start time, enter the new start time in MM-DD.hh:mm format, and click **OK**.
   - Start Time: **12-10.23:59**
5. Click **Commit Options>Commit** to commit your changes.
6. Click **Check Status** to monitor the progress of an active download or update, or to check the outcome of the latest update.

Step-by-Step Procedure

To use the CLI to automatically update the Junos OS application signature package:

1. Specify the URL for the security package. The security package includes the detector and the latest attack objects and groups. The following statement specifies https://signatures.juniper.net/cgi-bin/index.cgi as the URL for downloading signature database updates:

   ```
   [edit]
   user@host# set services application-identification download url
     https://signatures.juniper.net/cgi-bin/index.cgi
   ```

2. Specify the time and interval for download. The following statement sets the interval as 48 hours and the start time as 11:59 pm on December 10:

   ```
   [edit]
   user@host# set services application-identification download automatic interval 48
     start-time 12-10.23:59
   ```

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

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

Verification

To verify that the application signature package is being updated properly, enter the **show services application-identification version** command. Review the version number and details for the latest update.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Verifying the Junos OS Application Identification Extracted Application Package on page 506

Scheduling the Application Signature Package Updates As Part of the IDP Security Package

The configuration instructions in this example describe how to setup automatic updates of application identification signature package (part of IDP security package) at a specified date and time.

- Requirements on page 501
- Overview on page 501
- Configuration on page 502
- Verification on page 502

Requirements

Before you begin:

- Ensure that your SRX Series device has a connection to the Internet to download security package updates.

  NOTE: **DNS must be set up because you need to resolve the name of the update server.**

- Ensure that you have installed the application identification feature license. See "Installing and Verifying Licenses for an Application Signature Package" on page 491.

Overview

In this example, you want to download the current version of the application signature package periodically. The download should start at 11:59 PM on December 10. To maintain the most current information, you want to update the package automatically every 2 days from your company's intranet site.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To set up the automatic download and periodic update with the J-Web interface:

1. Enter **Configure>Security>IDP>Signature Updates** to display the Security IDP Signature Configuration page.
2. Click **Download Settings** and modify the URL: **https://signatures.juniper.net/cgi-bin/index.cgi**
3. Click the **Auto Download Settings** tab, and modify the following fields:
   - Interval: **48**
   - Start Time: **2013-12-10.23:59:55**
   - Enable Schedule Update: Select the check box.
4. Click **Reset Setting** to clear the existing fields, enter the new values. Click **OK**.
5. Click **Commit Options>Commit** to commit your changes.
6. Click **Check Status** to monitor the progress of an active download or update, or to check the outcome of the latest update.

Step-by-Step Procedure

To use the CLI to automatically update the Junos OS application signature package:

1. Specify the URL for the security package. The security package includes the detector and the latest attack objects and groups. The following statement specifies https://signatures.juniper.net/cgi-bin/index.cgi as the URL for downloading signature database updates:

   ```
   [edit]
   user@host# set security idp security-package url
     https://signatures.juniper.net/cgi-bin/index.cgi
   ```

2. Specify the time and interval for download. The following statement sets the interval as 48 hours and the start time as 11:55 pm on December 10, 2013:

   ```
   [edit]
   user@host# set security idp security-package automatic interval 48 start-time
     2013-12-10.23:55:55
   ```

3. Enable an automatic download and update of the security package.

   ```
   [edit]
   user@host# set security idp security-package automatic enable
   ```

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

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

Verification

Confirm that the application signature package is being updated properly.

Copyright © 2016, Juniper Networks, Inc.

Verifying application signature package

Purpose Verify services application identification version

Action From operational mode, enter the **show services application-identification version** command.

```
user@host> show services application-identification version
Application package version: 1884
```

Meaning The sample output shows that, the services application identification version is 1884.

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Verifying the Junos OS Application Identification Extracted Application Package on page 506

Example: Downloading and Installing the Application Identification Package in Chassis Cluster Mode

This example shows how to download and install the application signature package database to a device operating in chassis cluster mode.

- Requirements on page 503
- Overview on page 504
- Downloading and Installing the Application Identification Package on page 504

Requirements

Before you begin:

- Set the chassis cluster node ID and cluster ID. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*.
- Ensure that your SRX Series device has a connection to the Internet to download security package updates.

NOTE: DNS must be set up because you need to resolve the name of the update server.

- Ensure that you have installed application identification feature license. See "Installing and Verifying Licenses for an Application Signature Package" on page 491.

Copyright © 2016, Juniper Networks, Inc.

Overview

If you use application identification, you can download the predefined application signature package database. Juniper Networks regularly updates the database and makes it available on the Juniper Networks website. This package includes application objects that can be used to match traffic in IDP, application firewall policies, and application tracking. For more details, see "Understanding the Junos OS Application Package Installation" on page 489.

When you download the application identification security package on a device operating in chassis cluster mode, the security package is downloaded to the primary node and then synchronized to the secondary node.

Downloading and Installing the Application Identification Package

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

To download and install an application package:

1. Download the application package on the primary node.

   {primary:node0}[edit]

   user@host> request services application-identification download

   ```
   Please use command "request services application-identification download status"
    to check status
   ```

2. Check the application package download status.

   {primary:node0}[edit]

   user@host> request services application-identification download status

   On a successful download, the following message is displayed

   ```
   Application package 2345 is downloaded successfully
   ```

   The application package is installed in the application signature database on the primary node, and application identification files are synchronized on the primary and secondary nodes.

3. Update the application package using **install** command.

   {primary:node0}[edit]

   user@host> request services application-identification install

   ```
   node0:
   --------------------------------------------------------------------------
   Please use command "request services application-identification install status"
    to check status and use command "request services application-identification
    proto-bundle-status" to check protocol bundle status

   node1:
   --------------------------------------------------------------------------
   Please use command "request services application-identification install status"
    to check status and use command "request services application-identification
   ```


Copyright © 2016, Juniper Networks, Inc.

```
proto-bundle-status" to check protocol bundle status
```

4. Check the application package update status. The command output displays information about the downloaded and installed versions of the application package.

   {primary:node0}[edit]

   user@host> request services application-identification install status

   ```
   node0:
   --------------------------------------------------------------------------
   Install application package 2345 succeed
   node1:
   --------------------------------------------------------------------------
   Install application package 2345 succeed
   ```

   NOTE: It is possible that an application signature is removed from the new version of an application signature database. If this signature is used in an existing application firewall policy on your device, the installation of the new database will fail. An installation status message identifies the signature that is no longer valid. To update the database successfully, remove all references to the deleted signature from your existing policies and groups, and rerun the install command.

To uninstall an application package:

1. Uninstall the application package using **uninstall** command.

   {primary:node0}[edit]

   user@host> request services application-identification uninstall

   ```
   node0:
   --------------------------------------------------------------------------
   Please use command "request services application-identification uninstall
   status" to check status and use command "request services
   application-identification proto-bundle-status" to check protocol bundle status
   node1:
   --------------------------------------------------------------------------
   Please use command "request services application-identification uninstall
   status" to check status and use command "request services
   application-identification proto-bundle-status" to check protocol bundle status
   ```

2. Check the uninstall status of the application package.

   {primary:node0}[edit]

   user@host> request services application-identification uninstall status

   ```
   node0:
   --------------------------------------------------------------------------
   Uninstall application package 2345 succeed
   node1:
   --------------------------------------------------------------------------
   Uninstall application package 2345 succeed
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Check the uninstall status of protocol bundle:

   user@host>**request services application-identification proto-bundle-status**

```
Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and application
 secpack version (2345) is unloaded and deactivated
```

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Installing and Verifying Licenses for an Application Signature Package on page 491
- Verifying the Junos OS Application Identification Extracted Application Package on page 506

Verifying the Junos OS Application Identification Extracted Application Package

Purpose After successful download and installation of the application package, use the following commands to view the predefined application signature package content.

Action

- View the current version of the application package:

  show services application-identification version

```
Application package version: 1608
```

- View the current status of the application package:

  show services application-identification status

```
pic: 1/0

Application Identification
 Status                          Enabled
 Sessions under app detection    0
 Engine Version                  4.18.1-20 (build date Jan 25 2014)
 Max TCP session packet memory   30000
 Max C2S bytes                   1024
 Max S2C bytes                   0
 Force packet plugin             Disabled
 Force stream plugin             Disabled
 Statistics collection interval  1 (in minutes)

Application System Cache
 Status                          Enabled
 Negative cache status           Disabled
 Max Number of entries in cache  131072
 Cache timeout in seconds        3600

Protocol Bundle
 Download Server
https://services.netscreen.com/cgi-bin/index.cgi
 AutoUpdate                      Enabled
Slot 1:
 Status                          Active
 Version                         1.30.4-22.005 (build date Jan 17 2014)
 Sessions                        0
Slot 2
 Status                          Free
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding the Junos OS Application Package Installation on page 489
- Downloading and Installing the Junos OS Application Signature Package Manually on page 493

Uninstalling the Junos OS Application Identification Application Package

You can uninstall the predefined application package. The uninstall operation will fail if there are any active security policies referenced in the predefined application signatures in the Junos OS configuration

To uninstall application package:

1. Uninstall the application package:

   user@host> **request services application-identification uninstall**

   ```
   Please use command "request services application-identification uninstall
   status" to check status and use command "request services
   application-identification proto-bundle-status" to check protocol bundle status
   ```

2. Check the uninstall operation status of the application package. The command output displays information about the uninstall status of the application package and protocol bundle.

   - Check the uninstall status:

     user@host>**request services application-identification uninstall status**

     ```
     Uninstall application package 2345 succeed
     ```

   - Check the uninstall status of protocol bundle:

     user@host>**request services application-identification proto-bundle-status**

     ```
     Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and
     application secpack version (2345) is unloaded and deactivated
     ```

The application package and protocol bundle are uninstalled on the device. To reinstall application identification, you need to download application package and reinstall it again.

Related Documentation

- Downloading and Installing the Junos OS Application Signature Package Manually on page 493
- Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package on page 496
- Verifying the Junos OS Application Identification Extracted Application Package on page 506

Disabling and Reenabling Junos OS Application Identification

Application identification is enabled by default. You can disable application identification with the CLI.

Copyright © 2016, Juniper Networks, Inc.

To disable application identification:

```
user@host# set services application-identification no-application-identification
```

If you want to reenable application identification, delete the configuration statement that specifies disabling of application identification:

```
user@host# delete services application-identification no-application-identification
```

If you are finished configuring the device, commit the configuration.

To verify the configuration, enter the **show services application-identification** command. See "Verifying the Junos OS Application Identification Extracted Application Package" on page 506.

Related Documentation

- Understanding Application Identification Techniques on page 485
- Understanding the Junos OS Application Identification Database on page 488

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 24

Custom Application Signatures

- Understanding Junos OS Application Identification Custom Application Signatures on page 509
- Example: Configuring Junos OS Application Identification Custom Application Signatures on page 512

Understanding Junos OS Application Identification Custom Application Signatures

Application identification supports user-defined custom application signatures and signature groups. Custom application signatures are unique to your environment and are not part of the predefined application package. You must install application signature package on your device to use custom signatures. When the custom signatures are configured, you cannot uninstall the application signature package.

Custom application signatures are required:

- To control traffic particular to an environment
- To bring visibility for unknown or unclassified applications by developing custom applications.
- To identify applications over Layer 7 and transiting or temporary applications, and to achieve further granularity of known applications
- To perform QoS for your specific application

You can create custom application signatures using CLI by specifying a name, protocol, port where the application runs, and match criteria. For more details, see "Example: Configuring Junos OS Application Identification Custom Application Signatures" on page 512.

CAUTION: We recommend that only advanced Junos OS users attempt to customize application signatures.

You can view application signatures and application signature groups by using the **show services application-identification application** and **show services application-identification group** commands.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The following features are not supported:

- Copying a predefined application signature to use as a model by entering the **request services application-identification application copy** command
- Prioritizing custom signatures over a specific predefined custom signature
- Complete Perl Compatible Regular Expressions (PCRE)-based character set, and unicode-based characters
- Enforcing of order among members in Layer 7-based signatures
- The wildcard address for address-based signatures (Layer 3 and Layer 4)

Unlike predefined signatures and groups, custom application signatures and groups are saved in the configuration hierarchy, not in the predefined application signature database. Custom application signatures and signature groups are located in the **[services application-identification]** hierarchy.

SRX Series devices support the following types of custom signatures:

- ICMP-Based Mapping on page 510
- Address-Based Mapping on page 510
- IP Protocol-Based Mapping on page 511
- Layer 7-Based Signatures on page 511

ICMP-Based Mapping

The ICMP mapping technique maps standard ICMP message types and optional codes to a unique application name. This mapping technique lets you differentiate between various types of ICMP messages.

NOTE: IDP works only with TCP or UDP traffic. ICMP mapping, therefore, does not apply to IDP and cannot support IDP features such as custom attacks.

NOTE: The ICMP mapping technique used for mapping standard ICMP message types and optional codes are not supported for ICMPv6 traffic.

Address-Based Mapping

Layer 3 and Layer 4 address mapping defines an application by the IP address and optional port range of the traffic.

To ensure adequate security, use address mapping when the configuration of your private network predicts application traffic to or from trusted servers. Address mapping provides efficiency and accuracy in handling traffic from a known application.

Copyright © 2016, Juniper Networks, Inc.

Layer 3 and Layer 4 address-based custom applications, you can match the IP address and/or port range to either source IP address or destination IP address and/or port range. When both IP address and port range are configured, both must match either source or destination tuples (IP address and port range) of the packet.

Consider a Session Initiation Protocol (SIP) server that initiates sessions from its known port 5060. Because all traffic from this IP address and port is generated by only the SIP application, the SIP application can be mapped to the server's IP address and port 5060 for application identification. In this way, all traffic with this IP address and port is identified as SIP application traffic.

NOTE: When you configure an address-based application and a TCP/UDP stream-based application, and a session matches both applications, TCP/UDP stream-based application is reported as application and address-based application is reported as extended application.

IP Protocol-Based Mapping

Standard IP protocol numbers can map an application to IP traffic. As with address mapping, to ensure adequate security, use IP protocol mapping only in your private network for trusted servers.

NOTE: IDP works only with TCP or UDP traffic. IP protocol mapping, therefore, does not apply to IDP and cannot support IDP features such as custom attacks.

Layer 7-Based Signatures

Layer 7 custom signatures define an application running over TCP or UDP or Layer 7 applications. Layer 7-based custom application signatures are required for the identification of multiple applications running on the same Layer 7 protocols. For example, applications such as Facebook and Yahoo Messenger can both run over HTTP, but there is a need to identify them as two different applications running on the same Layer 7 protocol.

Layer 7-based custom application signatures detect applications based on the patterns in HTTP contexts. However, some HTTP sessions are encrypted in SSL, also called Transport Layer Security (TLS). Application identification can also extract the server name information or the server certification from the TLS or SSL sessions. It can also detect patterns in TCP or UDP payload in Layer 7 applications.

Related Documentation

- Understanding Application Identification Techniques on page 485
- Understanding the Junos OS Application Package Installation on page 489
- Understanding the Junos OS Application Identification Database on page 488

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Junos OS Application Identification Custom Application Signatures

This example shows how to configure custom application signatures for Junos OS application identification.

CAUTION: We recommend that only advanced Junos OS users attempt to customize application signatures.

- Requirements on page 512
- Overview on page 512
- Configuration on page 512
- Verification on page 516

Requirements

Before you begin:

- SRX Series device with application signature package installed. See "Downloading and Installing the Junos OS Application Signature Package Manually" on page 493 or "Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package" on page 496.
- SRX Series device must be running Junos OS Release 15.1X49-D40 or later.

Overview

Application identification supports custom application signatures to detect applications as they pass through the device. When you configure custom signatures, make sure that your signatures are unique.

In this example, you create custom application signatures for applications based on ICMP, IP protocol, IP address, and Layer 7.

Configuration

CLI Quick Configuration

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

HTTP Context-Based Custom Signatures

```
set services application-identification application mycustom-http over HTTP signature
    s1 member m01 context http-header-host
set services application-identification application mycustom-http over HTTP signature
    s1 member m01 pattern .*agent1.*
set services application-identification application mycustom-http over HTTP signature
    s1 member m01 direction any
```


Copyright © 2016, Juniper Networks, Inc.

SSL Context-Based Custom Signatures

```
set services application-identification application mycustom-ssl over SSL signature s1
  member m01 context ssl-server-name
set services application-identification application mycustom-ssl over SSL signature s1
  member m01 pattern "example\.com"
set services application-identification application mycustom-ssl over SSL signature s1
  member m01 direction any
```

TCP Stream-Based Custom Signatures

```
set services application-identification application mycustom-tcp over TCP signature s1
  member m01 context stream
set services application-identification application mycustom-tcp over TCP signature s1
  member m01 pattern "12345678901234567890123456789O"
set services application-identification application mycustom-tcp over TCP signature s1
  member m01 direction client-to-server
```

ICMP-Based

```
set services application-identification application MY-ICMP icmp-mapping type 100
set services application-identification application MY-ICMP icmp-mapping code 1
```

Layer 3/Layer 4 Address-Based

```
set services application-identification application My-ADDRESS address-mapping
  ADDR-SAMPLE filter ip 5.0.0.1/32
set services application-identification application My-ADDRESS address-mapping
  ADDR-SAMPLE filter port-range udp 5000-6000
```

IP Protocol-Based

```
set services application-identification application MY-IGMP ip-protocol-mapping protocol
  2
```

Step-by-Step Procedure

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

To configure HTTP context-based custom signatures:

1. Configure an application based on HTTP context. Define an application signature to match the pattern, a unique application signature identifier, application signature member identifier, and set the context to be matched.

   ```
   [edit services application-identification]
   user@host# set application mycustom-http over HTTP signature s1 member m01
     context http-header-host
   ```

2. Configure a pattern to match the context.

   ```
   [edit services application-identification]
   user@host# set application mycustom-http over HTTP signature s1 member m01
     pattern .*agent1.*
   ```

3. Configure the connection direction of the packets to apply pattern matching.

   ```
   [edit services application-identification]
   user@host# set application mycustom-http over HTTP signature s1 member m01
     direction any
   ```

Step-by-Step Procedure

To configure SSL context-based custom signatures:

1. Configure an application based on SSL. Define an application signature to match the pattern, a unique application signature identifier, application signature member identifier, and set the context to be matched.

Copyright © 2016, Juniper Networks, Inc.

```
[edit services application-identification]
user@host# set application mycustom-ssl over SSL signature s1 member m01
  context ssl-server-name
```

2. Configure a pattern to match the context.

```
[edit services application-identification]
user@host# set application mycustom-ssl over SSL signature s1 member m01
  pattern "example\.com"
```

3. Configure the connection direction of the packets to apply pattern matching.

```
[edit services application-identification]
user@host# set application mycustom-ssl over SSL signature s1 member m01
  direction any
```

Step-by-Step Procedure

To configure TCP stream-based custom signatures:

1. Configure an application based on TCP. Define an application signature to match the pattern, a unique application signature identifier, application signature member identifier, and set the context to be matched.

```
[edit services application-identification]
user@host# set application mycustom-tcp over TCP signature s1 member m01
  context stream
```

2. Configure a pattern to match the context.

```
[edit services application-identification]
user@host# set application mycustom-tcp over TCP signature s1 member m01
  pattern ""12345678901234567890123456789O"
```

3. Configure the connection direction of the packets to apply pattern matching.

```
[edit services application-identification]
user@host# set application mycustom-tcp over TCP signature s1 member m01
  direction client-to-server
```

Step-by-Step Procedure

To configure ICMP-based custom applications signatures:

1. Define the type of ICMP mapping. The type field identifies the ICMP message.

```
[edit services application-identification]
user@host# set application MY-ICMP icmp-mapping type 100
```

2. Define the code for ICMP mapping. The code field provides further information about the associated type field.

```
[edit services application-identification]
user@host# set application MY-ICMP icmp-mapping code 1
```

Step-by-Step Procedure

To configure Layer 3 or Layer 4 address-based custom applications signatures:

1. Configure the application to match the specified IP address.

```
[edit services application-identification]
user@host# set application My-ADDRESS address-mapping ADDR-SAMPLE filter
  ip 5.0.0.1/32
```

2. Configure the port range for TCP or UDP.

Copyright © 2016, Juniper Networks, Inc.

```
[edit services application-identification]
user@host# set application My-ADDRESS address-mapping ADDR-SAMPLE filter
  port-range udp 5000-6000
```

Step-by-Step Procedure

To configure IP protocol mapping-based custom application signatures:

- Specify the IP protocol value for an application to match.

```
[edit services application-identification]
user@host# set application MY-IGMP ip-protocol-mapping protocol 2
```

Results

From configuration mode, confirm your configuration by entering the **show services application-identification** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services application-identification

download {
  url https://services.netscreen.com/cgi-bin/index.cgi;
}
application MY-ICMP {
  icmp-mapping {
    type 100;
    code 1;
  }
}
application MY-IGMP {
  ip-protocol-mapping {
    protocol 2;
  }
}
application My-ADDRESS {
  address-mapping ADDR-SAMPLE {
    filter {
      ip 5.0.0.1/32;
      port-range {
        udp 5000-6000;
      }
    }
  }
}
application mycustom-http {
  over HTTP {
    signature s1 {
      member m01 {
        context http-header-host;
        pattern ".*agent1.*;";
        direction any;
      }
    }
  }
}
application mycustom-ssl {
  over SSL {
```

Copyright © 2016, Juniper Networks, Inc.

```
        signature s1 {
          member m01 {
            context ssl-server-name;
            pattern "example\.com";
            direction any;
          }
        }
      }
    }
    application mycustom-tcp {
      over TCP {
        signature s1 {
          member m01 {
            context stream;
            pattern 12345678901234567890123901234567;
            direction client-to-server;
          }
        }
      }
    }
```

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

Verification

Verifying the Custom Application Definitions

Purpose Display predefined and custom application signatures and settings that are configured on your device. Note that predefined application signature names use the prefix “junos:”

Action From configuration mode, enter the **show services application-identification application detail *name*** command.

See show services application-identification application

Related Documentation

- Understanding Application Identification Techniques on page 485
- Understanding Junos OS Application Identification Custom Application Signatures on page 509
- Understanding the Junos OS Application Package Installation on page 489
- Understanding the Junos OS Application Identification Database on page 488

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 25

Configuring Application Groups

- Customizing Application Groups for Junos OS Application Identification on page 517
- Enabling Application Groups in Junos OS Application Identification on page 518
- Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management on page 518

Customizing Application Groups for Junos OS Application Identification

In Junos OS, application identification allows you to group applications in policies. Applications can be grouped under predefined and custom application groups. The entire predefined application group can be downloaded as part of the IDP or application identification security package. You can create custom application groups with a set of similar applications for consistent reuse when defining policies.

Application group support associates related applications under a single name for simplified, consistent reuse when using any application services.

The hierarchy of application groups resembles a tree structure with associated applications as the leaf nodes. The group *any* refers to the root node. The group *unassigned* is always situated one level from the root and initially contains all applications. When a group is defined, applications are assigned from the unassigned group to the new group. When a group is deleted, its applications are moved back to the unassigned group.

All predefined application groups have the prefix “junos“ in the application group name to prevent naming conflicts with custom application groups. You cannot modify the list of applications within a predefined application group. However, you can copy a predefined application group to use it as a template for creating a custom application group.

To customize a predefined application group, you must first disable the predefined group. Note that a disabled predefined application group remains disabled after an application database update. You can then use the operational command **request services application-identification group** to copy the disabled predefined application group. The copied group is placed in the configuration file, and the prefix “junos” is changed to “my”. At this point, you can modify the list of applications in “my” application group and rename the group with a unique name.

To reassign an application from one custom group to another, you must remove the application from its current custom application group, and then reassign it to the other.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Application Identification Techniques on page 485
- Enabling Application Groups in Junos OS Application Identification on page 518
- Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management on page 518

Enabling Application Groups in Junos OS Application Identification

All application groups are enabled by default. Predefined application groups are enabled at installation.

- For predefined application groups, you can disable and reenable a group using the **request services application-identification group** command. You cannot delete a predefined signature or signature group.
  - To disable a predefined application group:

    user@host> request services application-identification group disable *predefined-application-group-name*
  - To reenable a disabled predefined application group:

    user@host> request services application-identification group enable *predefined-application-group-name*

Related Documentation

- Understanding Application Identification Techniques on page 485
- Customizing Application Groups for Junos OS Application Identification on page 517
- Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management on page 518

Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management

This example shows how to configure custom application groups for Junos OS application identification for consistent reuse when defining policies.

- Requirements on page 518
- Overview on page 519
- Configuration on page 519

Requirements

Before you begin, install an entire signature database from an IDP or an application identification security package. See "Downloading and Installing the Junos OS Application Signature Package Manually" on page 493 or "Downloading and Installing the Junos OS Application Signature Package As Part of the IDP Security Package" on page 496.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you define applications for an application group, delete an application from an application group, and include an application group within another application group.

In Junos OS, application identification allows you to group applications in policies. Applications can be grouped under predefined and custom application groups. The entire predefined application group can be downloaded as part of the IDP or application identification security package. You can create custom application groups with a set of similar applications for consistent reuse when defining policies.

NOTE: You cannot modify the applications defined in a predefined application group. However, you can copy a predefined application group using the operational command **request services application-identification group** *group-name* **copy** to create a custom application group and modify the list of applications. For more information, see request services application-identification group.

Configuration

- Configuring Junos OS Application Identification User-Defined Application Groups on page 519
- Deleting an Application from a User-Defined Application Group on page 521
- Creating Child Application Groups for an Application Group on page 521

Configuring Junos OS Application Identification User-Defined Application Groups

CLI Quick Configuration

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

```
set services application-identification application-group my_web
set services application-identification application-group my_web applications junos:HTTP
set services application-identification application-group my_web applications junos:FTP
set services application-identification application-group my_web applications
    junos:GOPHER
set services application-identification application-group my_web applications
    junos:AMAZON
set services application-identification application-group my_peer
set services application-identification application-group my_peer applications
    junos:BITTORRENT
set services application-identification application-group my_peer applications
    junos:BITTORRENT-DHT
set services application-identification application-group my_peer applications
    junos:BITTORRENT-UDP
set services application-identification application-group my_peer applications
    junos:BITTRACKER
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure a custom application group for application identification:

1. Set the name of your custom application group.

```
[edit services application-identification]
user@host# set application-group my_web
```

2. Add the list of applications that you want to include in your custom application group.

```
[edit services application-identification]
user@host# set application-group my_web applications junos:HTTP
user@host# set application-group my_web applications junos:FTP
user@host# set application-group my_web applications junos:GOPHER
user@host# set application-group my_web applications junos:AMAZON
```

3. Set the name of a second custom application group.

```
[edit services application-identification]
user@host# set application-group my_peer
```

4. Add the list of applications that you want to include in the group.

```
[edit services application-identification]
user@host# set application-group my_peer applications junos:BITTORRENT
user@host# set application-group my_peer applications junos:BITTORRENT-DHT
user@host# set application-group my_peer applications junos:BITTORRENT-UDP
user@host# set application-group my_peer applications junos:BITTRACKER
```

Results

From configuration mode, confirm your configuration by entering the **show services application-identification group** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services application-identification application-group my_web
  applications {
    junos:HTTP;
    junos:FTP;
    junos:GOPHER;
    junos:AMAZON
  }
user@host# show services application-identification application-group my_peer
  applications {
    junos:BITTORRENT;
    junos:BITTORRENT-DHT;
    junos:BITTORRENT-UDP;
    junos:BITTRACKER;
  }
```

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

Copyright © 2016, Juniper Networks, Inc.

Deleting an Application from a User-Defined Application Group

CLI Quick Configuration

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

```
[edit]
delete services application-identification application-group my_web applications
  junos:AMAZON
```

Step-by-Step Procedure

To delete an application from a custom application group:

- Delete an application from a custom application group.

```
[edit services application-identification]
user@host# delete application-group my_web applications junos:AMAZON
```

Results

From configuration mode, confirm your configuration by entering the **show services application-identification application group detail** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services application-identification group detail
  application group my_web {
    junos:HTTP;
    junos:FTP;
    junos:GOPHER;
  }
```

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

Creating Child Application Groups for an Application Group

CLI Quick Configuration

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

```
set services application-identification application-group p2p
set services application-identification application-group p2p application-groups my_web
set services application-identification application-group p2p application-groups my_peer
```

Step-by-Step Procedure

To configure child application groups for a custom application group:

1. Set the name of the custom application group in which you are configuring the child application groups.

```
[edit services application-identification]
user@host# set application-group p2p
```

2. Add the child application groups.

Copyright © 2016, Juniper Networks, Inc.

```
[edit services application-identification]
user@host# set application-group p2p application-groups my_web
uer@host# set application-group p2p application-groups my_peer
```

Results From configuration mode, confirm your configuration by entering the **show services application-identification application-group** *application-group-name* command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services application-identification application-group p2p
  applications-groups {
    my_web;
    my_peer;
  }
```

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

Related Documentation

- Understanding Application Identification Techniques on page 485
- Customizing Application Groups for Junos OS Application Identification on page 517
- Enabling Application Groups in Junos OS Application Identification on page 518

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 26

Configuring Application System Cache

- Understanding the Application System Cache on page 523
- Deactivating Application System Cache Information for Application Identification (CLI Procedure) on page 524
- Verifying Application System Cache Statistics on page 524

Understanding the Application System Cache

Application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service.

Once an application is identified, its information is saved in the ASC so that only one matching entry is required for an application running on a particular system, thereby expediting the identification process.

By default, the ASC saves the mapping information for 3600 seconds. However, you can configure the cache timeout value by using the CLI.

To minimize the impact on performance, application system cache is refreshed only when Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) traffic triggers a cache lookup. Without a cache lookup, the entries in the ASC remain unchanged even after cache timeout.

NOTE: **When you configure a new custom application signature or modify an existing custom signature, all the existing application system cache entries for predefined and custom applications will be cleared.**

Related Documentation

- Understanding Application Identification Techniques on page 485
- Understanding the Junos OS Application Identification Database on page 488
- Verifying Application System Cache Statistics on page 524

Copyright © 2016, Juniper Networks, Inc.

Deactivating Application System Cache Information for Application Identification (CLI Procedure)

Application caching is turned on by default. You can manually turn this caching off using the CLI.

```
user@host# set services application-identification no-application-system-cache
```

When you use the **show** command in the CLI operation mode for the application system cache (ASC), application cache is listed as **off**. Note that if the cache contains data from the prior implementation, the cached data is also displayed.

```
user@host> show services application-identification application-system-cache

application-cache: off
  nested-application-cache: on
  cache-unknown-result: on
  cache-entry-timeout: 3600 seconds
```

Related Documentation

- Understanding Application Identification Techniques on page 485
- Verifying Application System Cache Statistics on page 524
- Understanding the Junos OS Application Identification Database on page 488

Verifying Application System Cache Statistics

Purpose Verify the application system cache (ASC) statistics.

NOTE: The application system cache will display the cache for application identification applications.

Action From CLI operation mode, enter the **show services application-identification application-system-cache** command.

Sample Output

```
user@host> show services application-identification application-system-cache
application-cache: on
  nested-application-cache: on
  cache-unknown-result: on
  cache-entry-timeout: 3600 seconds
```

Meaning The output shows a summary of the ASC statistics information. Verify the following information:

- IP address—Displays the destination address.
- Port—Displays the destination port on the server.

Copyright © 2016, Juniper Networks, Inc.

- Protocol—Displays the protocol type on the destination port.
- Application—Displays the name of the application identified on the destination port.

NOTE: On branch SRX Series devices, when there are a large number of ASC entries (10,000 or more), and the entries are to be listed in the output for the command **show services application-identification application-system-cache**, it results in CLI session timeout.

Related Documentation

- Understanding Application Identification Techniques on page 485
- Deactivating Application System Cache Information for Application Identification (CLI Procedure) on page 524

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 27

Controlling Application Identification Performance

- Onbox Application Identification Statistics on page 527
- Understanding Jumbo Frames Support for Junos OS Application Identification Services on page 528
- Improving the Application Traffic Throughput on page 528

Onbox Application Identification Statistics

Application Identification services provide statistical information per session. These statistics provide customers with an application usage profile. The Onbox Application Identification Statistics feature adds application-level statistics to the AppSecure suite. Application statistics allow an administrator to access cumulative statistics as well as statistics accumulated over user-defined intervals.

With this feature, the administrator can clear the statistics and configure the interval values while maintaining bytes and session count statistics. Because the statistics count occurs at session close event time, the byte and session counts are not updated until the session closes. SRX Series devices support a history of eight intervals that an administrator can use to display application session and byte counts.

If application grouping is supported in your configuration of Junos OS, then the Onbox Application Identification Statistic feature supports onbox per-group matching statistics. The statistics are maintained for predefined groups only.

Reinstalling an application signature package will not clear the application statistics. If the application is disabled, there will not be any traffic for that application, but the application is still maintained in the statistics. It does not matter if you are reinstalling a predefined application, because applications are tracked according to application type. For predefined group statistics, reinstalling a security package will not clear the statistics. However, any changes to group memberships are updated. For example, junos:web might have 50 applications in the current release and 60 applications following an upgrade. Applications that are deleted and application groups that are renamed are handled in the same way as applications that are added.

The Application Identification module maintains a 64-bit session counters for each application on each Services Processing Unit (SPU). The counter increments when a

Copyright © 2016, Juniper Networks, Inc.

session is identified as a particular application. Another set of 64-bit counters aggregates the total bytes per application on the SPU. Counters for unspecified applications are also maintained. Statistics from multiple SPUs for both sessions and bytes are aggregated on the Routing Engine and presented to the users.

Individual SPUs have interval timers to roll over statistics per *interval* time. To configure the interval for statistics collection, use the **set services application-identification application-statistics interval time** command. Whenever the Routing Engine queries for the required interval, the corresponding statistics are fetched from each SPU, aggregated in the Routing Engine and presented to the user.

Use the **clear-services-application-identification-application-statistics** command to reset the counters manually. Counters reset automatically when a device is upgraded or rebooted, when flowd restarts, or when there is a change in the interval timer.

Related Documentation

- Understanding Application Identification Techniques on page 485

Understanding Jumbo Frames Support for Junos OS Application Identification Services

Application identification support the larger jumbo frame size of 9192 bytes. Although jumbo frames are enabled by default, you can adjust the maximum transmission unit (MTU) size by using the [**set interfaces**] command. CPU overhead can be reduced while processing jumbo frames.

Related Documentation

- Understanding Application Identification Techniques on page 485
- Understanding the Junos OS Application Identification Database on page 488

Improving the Application Traffic Throughput

The application traffic throughput can be improved by setting the deep packet inspection (DPI) in performance mode with default packet inspection limit as two packets, including both client-to-server and server-to-client directions. By default, performance mode is disabled on SRX Series devices.

To improve the application traffic throughput:

1. Enable the DPI performance mode.

   ```
   [edit]
   user@host# set services application-identification enable-performance-mode
   ```

2. (Optional) You can set the maximum packet threshold for DPI performance mode, including both client-to-server and server-to-client directions.

   You can set the packet inspection limit from 1 through 100.

   ```
   [edit]
   user@host# set services application-identification enable-performance-mode
     max-packet-threshold value
   ```

3. Commit the configuration.

Copyright © 2016, Juniper Networks, Inc.

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

Use the **show services application-identification status** command to display detailed information about application identification status.

show services application-identification status (DPI Performance Mode Enabled)

```
user@host> show services application-identification status
pic: 2/1

Application Identification
Status                            Enabled
Sessions under app detection      0
Engine Version                    4.18.2-24.006 (build date Jul 30 2014)
Max TCP session packet memory     30000
Force packet plugin               Disabled
Force stream plugin               Disabled
DPI Performance mode:             Enabled
Statistics collection interval    1 (in minutes)

Application System Cache
Status                            Enabled
Negative cache status             Disabled
Max Number of entries in cache    262144
Cache timeout                     3600 (in seconds)

Protocol Bundle
Download Server                 https://services.netscreen.com/cgi-bin/index.cgi
AutoUpdate                        Disabled
Slot 1:
Application package version       2399
Status                            Active
Version                           1.40.0-26.006 (build date May 1 2014)
Sessions                          0
Slot 2
Application package version       0
Status                            Free
Version
Sessions                          0
```

The DPI Performance mode field displays whether the DPI performance mode is enabled or not. This field is displayed in the CLI command output only if the performance mode is enabled.

If you want to set DPI to default accuracy mode and disable the performance mode, delete the configuration statement that specifies enabling of the performance mode:

To disable the performance mode:

1. Delete the performance mode.

   ```
   [edit]
   user@host# delete services application-identification enable-performance-mode
   ```

2. Commit the configuration.

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

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- enable-performance-mode on page 681

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 28

Configuring Encrypted Files Using SSL Proxy

- SSL Proxy Overview on page 531
- Application Firewall, IDP, and Application Tracking with SSL Proxy Overview on page 540
- Configuring SSL Proxy on page 540
- Working with the Certificate Revocation Lists for SSL Proxy on page 550
- Enabling Debugging and Tracing for SSL Proxy on page 552

SSL Proxy Overview

Secure Sockets Layer (SSL) is an application-level protocol that provides encryption technology for the Internet. SSL, also called Transport Layer Security (TLS), ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private-public key exchange pairs for this level of security.

Server authentication guards against fraudulent transmissions by enabling a Web browser to validate the identity of a webserver. Confidentiality mechanisms ensure that communications are private. SSL enforces confidentiality by encrypting data to prevent unauthorized users from eavesdropping on electronic communications. Finally, message integrity ensures that the contents of a communication have not been tampered with.

SSL proxy is transparent; that is, it performs SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence. Existing features like SSL offload and SSL inspection require the servers to share their secret keys to be able to decrypt the SSL traffic. However, sharing server keys is sometimes not feasible or might not be available in certain circumstances, in which case the SSL traffic cannot be decrypted.

SSL proxy addresses this problem by ensuring that it has the keys to encrypt and decrypt the payload:

- For the server, SSL proxy acts as a client—Because SSL proxy generates the shared pre-master key, it determines the keys to encrypt and decrypt.
- For the client, SSL proxy acts as a server—SSL proxy first authenticates the original server and replaces the public key in the original server certificate with a key that is

Copyright © 2016, Juniper Networks, Inc.

known to it. It then generates a new certificate by replacing the original issuer of the certificate with its own identity and signs this new certificate with its own public key (provided as a part of the proxy profile configuration). When the client accepts such a certificate, it sends a shared pre-master key encrypted with the public key on the certificate. Because SSL proxy replaced the original key with its own key, it is able to receive the shared pre-master key. Decryption and encryption take place in each direction (client and server), and the keys are different for both encryption and decryption.

Figure 29 depicts how SSL inspection (on an existing SRX Series IDP module) is typically used to protect servers. SSL inspection requires access to the private keys used by the servers so that the SRX Series device can decrypt the encrypted traffic.

Figure 29: SSL Inspection on an Existing SRX Series IDP Module

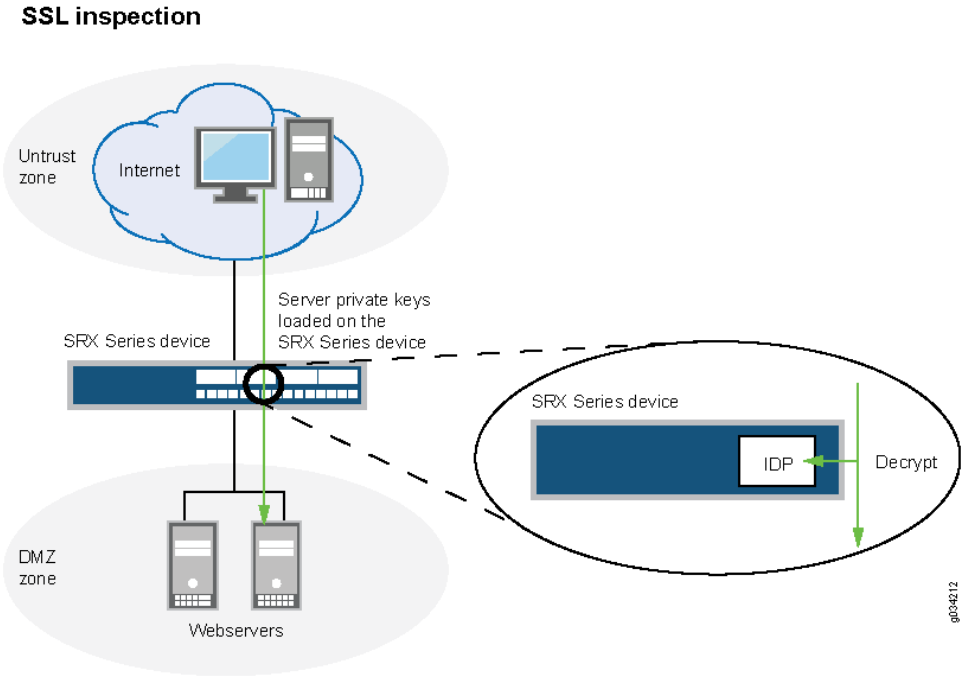

Figure 30 shows how SSL proxy works on an encrypted payload. When application firewall (AppFW), Intrusion Detection and Prevention (IDP), or application tracking (AppTrack) is configured, the SSL proxy acts as an SSL server by terminating the SSL session from the client and establishing a new SSL session to the server, the SRX Series device decrypts and then reencrypts all SSL proxy traffic. SSL proxy uses the following:

- SSL-T-SSL terminator on the client side.
- SSL-I-SSL initiator on the server side.
- Configured AppFW, IDP, or AppTrack services use the decrypted SSL sessions.

Copyright © 2016, Juniper Networks, Inc.

NOTE: If none of the services (AppFW, IDP, or AppTrack) are configured, then SSL proxy services are bypassed even if an SSL proxy profile is attached to a firewall policy.

Figure 30: SSL Proxy on an Encrypted Payload

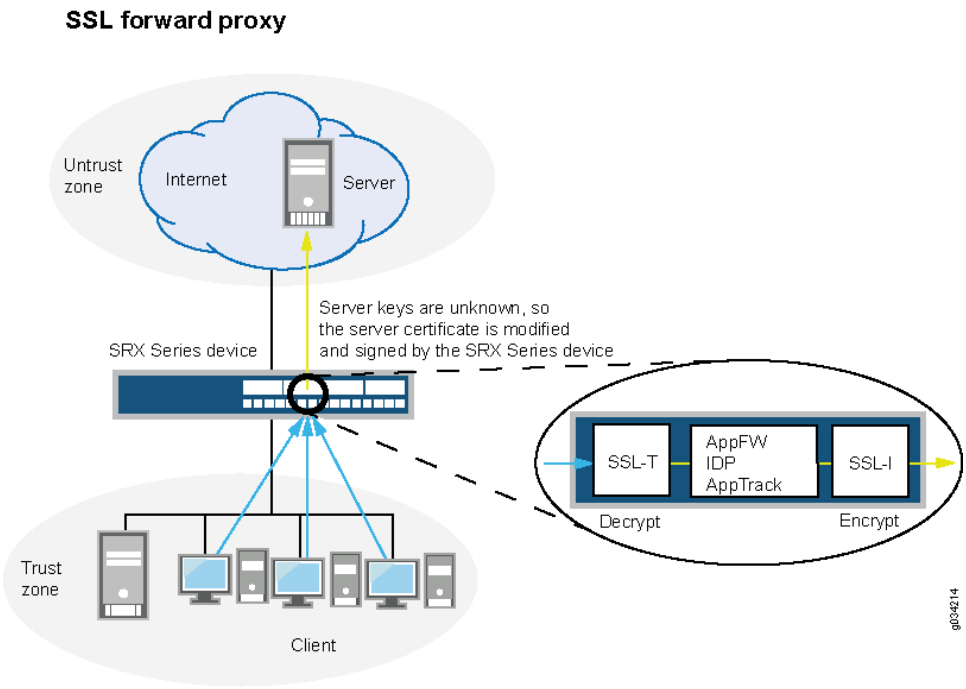

Supported Ciphers in Proxy Mode

An SSL cipher comprises encryption ciphers, authentication method, and compression. Table 34 displays a list of supported ciphers. NULL ciphers are excluded.

The following SSL protocols are supported on SRX Series devices:

- TLS version 1.0—Provides secure communication over networks by providing privacy and data integrity between communicating applications
- TLS version 1.1—This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.
- TLS version 1.2 — This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.

Table 34: Supported SSL Cipher List

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_RC4_128_MD5	RSA key exchange	128-bit RC4	Message Digest 5 (MD5) hash

Copyright © 2016, Juniper Networks, Inc.

Table 34: Supported SSL Cipher List *(continued)*

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_RC4_128_SHA	RSA key exchange	128-bit RC4	Secure Hash Algorithm (SHA) hash
RSA_WITH_DES_CBC_SHA	RSA key exchange	DES CBC	SHA hash
RSA_WITH_3DES_EDE_CBC_SHA	RSA key exchange	3DES EDE/CBC	SHA hash
RSA_WITH_AES_128_CBC_SHA	RSA key exchange	128-bit AES/CBC	SHA hash
RSA_WITH_AES_256_CBC_SHA	RSA key exchange	256-bit AES/CBC	SHA hash
RSA_EXPORT_WITH_RC4_40_MD5	RSA-export	40-bit RC4	MD5 hash
RSA_EXPORT_WITH_DES40_CBC_SHA	RSA-export	40-bit DES/CBC	SHA hash
RSA_EXPORT1024_WITH_DES_CBC_SHA	RSA 1024 bit export	DES/CBC	SHA hash
RSA_EXPORT1024_WITH_RC4_56_MD5	RSA 1024 bit export	56-bit RC4	MD5 hash
RSA_EXPORT1024_WITH_RC4_56_SHA	RSA 1024 bit export	56-bit RC4	SHA hash

NOTE: Cipher suites that have “export” in the title are intended for use outside of the United States and might have encryption algorithms with limited key sizes.

Export ciphers are not enabled by default. You need to either configure the export ciphers to enable or install a domestic package.

NOTE: Supported SSL ciphers for HTTPS firewall authentication are RSA_WITH_3DES_EDE_CBC_SHA , RSA_WITH_AES_128_CBC_SHA, and RSA_WITH_AES_256_CBC_SHA.

Server Authentication

Implicit trust between the client and the device (because the client accepts the certificate generated by the device) is an important aspect of SSL proxy. It is extremely important that server authentication is not compromised; however, in reality, self-signed certificates and certificates with anomalies are in abundance. Anomalies can include expired certificates, instances of common name not matching a domain name, and so forth. Server authentication is governed by setting the **ignore-server-auth-failure** option in the SSL proxy.

- By default, the **ignore-server-auth-failure** option is not defined as an action in the SSL proxy profile, and the following occurs:

Copyright © 2016, Juniper Networks, Inc.

  - If authentication succeeds, a new certificate is generated by replacing the keys and changing the issuer name to the issuer name that is configured in the root CA certificate in the proxy profile.
  - If authentication fails, the connection is dropped.

- If the **ignore-server-auth-failure** option is defined as an action in the SSL proxy profile, the following occurs:
  - If the certificate is self-signed, a new certificate is generated by replacing the keys only. The issuer name is not changed. This ensures that the client browser displays a warning that the certificate is not valid.
  - If the certificate has expired or if the common name does not match the domain name, a new certificate is generated by replacing the keys and changing the issuer name to SSL-PROXY: DUMMY_CERT:GENERATED DUE TO SRVR AUTH FAILURE. This ensures that the client browser displays a warning that the certificate is not valid.

Trusted CA List

SSL proxy ensures secure transmission of data between a client and a server. Before establishing a secure connection, SSL proxy checks CA certificates to verify signatures on server certificates. For this reason, a reasonable list of trusted CA certificates is required to effectively authenticate servers.

Junos OS provides the following options for trusted CA certificates:

- Loading the default trusted CA list—Junos OS provides a default list of certificates that contains well-known trusted CA certificates similar to the default certificates used by most common browsers. Without these default certificates, browsers would not be able to validate the identity of most websites and would mark them as untrusted sites.

  The Junos OS package contains the default CA certificates as a PEM file (for example, trusted_CA.pem). After you download the package and reboot your device, you can easily load the default certificates on your system using a CLI command.

  We recommend you load the default trusted CA list if you want to trust the same CA certificates as common browsers and avoid importing CA certificates manually.

- Importing the trusted CA list manually—You can import your own trusted CA certificates using the Public Key Infrastructure (PKI). The PKI helps verify and authenticate the validity of the trusted CA certificates. You create CA profile groups that include trusted CA certificates, then import the group on your device for server authentication.

- Ignoring server authentication—You can use the **ignore-server-auth-failure** option to ignore server authentication completely. In this case, SSL proxy ignores errors encountered during the server certificate verification process (such as CA signature verification failure, self-signed certificates, and certificate expiry).

  We do not recommend this option for authentication, because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause for dropped SSL sessions. See "Enabling Debugging and Tracing for SSL Proxy" on page 552.

Copyright © 2016, Juniper Networks, Inc.

Root CA

In a public key infrastructure (PKI) hierarchy, the root CA is at the top of the trust path. The root CA identifies the server certificate as a trusted certificate.

Client Authentication

Currently, client authentication is not supported in SSL proxy. If a server requests client authentication, a warning is issued that a certificate is not available. The warning lets the server determine whether to continue or to exit.

Whitelists

Because SSL encryption and decryption are complicated and expensive procedures, network administrators can selectively bypass SSL proxy processing for some sessions. Such sessions mostly include connections and transactions with trusted servers or domains with which network administrators are very familiar. There are also legal requirements to exempt financial and banking sites. Such exemptions are achieved by configuring the IP addresses or domain names of the servers under whitelists.

Dynamic Resolution of Domain Names

The IP addresses associated with domain names are dynamic and can change at any time. Whenever a domain IP address changes, it is propagated to the SSL proxy configuration (similar to what is done in the firewall policy configuration).

Session Resumption

An SSL session refers to the set of parameters and encryption keys created by performing a full handshake. A connection is the conversation or active data transfer that occurs within the session. The computational overhead of a complete SSL handshake and generation of master keys is considerable. In short-lived sessions, the time taken for the SSL handshake can be more than the time for data transfer.

To improve throughput and still maintain an appropriate level of security, SSL session resumption provides a session caching mechanism so that session information, such as the pre-master secret key and agreed-upon ciphers, can be cached for both the client and server. The cached information is identified by a session ID. In subsequent connections both parties agree to use the session ID to retrieve the information rather than create a new pre-master secret key. Session resumption shortens the handshake process and accelerates SSL transactions.

Session Renegotiation

After a session is created and SSL tunnel transport has been established, a change in SSL parameters requires renegotiation. SSL proxy supports both secure (RFC 5746) and nonsecure (TLS v1.0, TLS v1.1, and TLS v1.2) renegotiation. When session resumption is enabled, session renegotiation is useful in the following situations:

- Cipher keys need to be refreshed after a prolonged SSL session.
- Stronger ciphers need to be applied for a more secure connection.

Copyright © 2016, Juniper Networks, Inc.

A change in an SSL proxy profile that modifies a certificate, cipher strength, or trusted CA list flushes cache entries when the modified policy is committed. When a session is resumed, the SSL parameters associated with its session ID are retrieved from the cache. If the SSL proxy profile is not altered, cache entries corresponding to that profile are not flushed and the session continues. If the cache has been flushed, however, a full handshake must be performed to establish the new SSL parameters. (There is no impact to non-SSL sessions.)

SSL Proxy Logs

When logging is enabled in an SSL proxy profile, SSL proxy can generate the messages shown in Table 35.

Table 35: SSL Proxy Logs

Syslog Type	Description
SSL_PROXY_SSL_SESSION_DROP	Logs generated when a session is dropped by SSL proxy.
SSL_PROXY_SSL_SESSION_ALLOW	Logs generated when a session is processed by SSL proxy even after encountering some minor errors.
SSL_PROXY_SESSION_IGNORE	Logs generated if non-SSL sessions are initially mistaken as SSL sessions.
SSL_PROXY_SESSION_WHITELIST	Logs generated when a session is whitelisted.
SSL_PROXY_ERROR	Logs used for reporting errors.
SSL_PROXY_WARNING	Logs used for reporting warnings.
SSL_PROXY_INFO	Logs used for reporting general information.

All logs contain similar information as shown in the following example (actual order of appearance):

```
logical-system-name, session-id, source-ip-address, source-port,
destination-ip-address,destination-port,
nat-source-ip-address, nat-source-port, nat-destination-ip-address,
nat-destination-port, proxy profile name, source-zone-name,
source-interface-name, destination-zone-name,destination-interface-name, message
```

The **message** field contains the reason for the log generation. One of three prefixes shown in Table 36 identifies the source of the message. Other fields are descriptively labeled.

Table 36: SSL Proxy Log Prefixes

Prefix	Description
system	Logs generated due to errors related to the device or an action taken as part of the SSL proxy profile. Most logs fall into this category.
openssl error	Logs generated during the handshaking process if an error is detected by the openssl library.

Copyright © 2016, Juniper Networks, Inc.

Table 36: SSL Proxy Log Prefixes *(continued)*

Prefix	Description
certificate error	Logs generated during the handshaking process if an error is detected in the certificate (x509 related errors).

Sample logs:

```
Jun  1 05:11:13 4.0.0.254 junos-ssl-proxy: SSL_PROXY_SSL_SESSION_DROP: lsys:root
  23 <4.0.0.1/35090->5.0.0.1/443>  NAT:< 4.0.0.1/35090->5.0.0.1/443 >
ssl-inspect-profile  <untrust:ge-0/0/0.0->trust:ge-0/0/1.0> message:certificate
error: self signed certificate
```

NOTE: These logs capture sessions that are dropped by SSL proxy, not sessions that are marked by other modules that also use SSL proxy services.

For SSL_PROXY_SESSION_WHITELIST messages, an additional **host** field is included after the **session-id** and contains the IP address of the server or domain that has been whitelisted.

```
Jun  1 05:25:36 4.0.0.254 junos-ssl-proxy: SSL_PROXY_SESSION_WHITELIST: lsys:root
 24 host:5.0.0.1 <4.0.0.1/35090->5.0.0.1/443>  NAT:< 4.0.0.1/35090->5.0.0.1/443
> ssl-inspect-profile  <untrust:ge-0/0/0.0->trust:ge-0/0/1.0> message:system:
session whitelisted
```

Leveraging Dynamic Application Identification

SSL proxy uses application identification services to dynamically detect if a particular session is SSL encrypted. SSL proxies are allowed only if a session is SSL encrypted. The following rules apply for a session:

- Session is marked **Encrypted=Yes** in the application system cache. If the session is marked **Encrypted=Yes**, it indicates that the final match from application identification for that session is SSL encrypted, and SSL proxy transitions to a state where proxy functionality can be initiated.
- Session is marked **Encrypted=No** in the application system cache. If a non-SSL entry is found in the application system cache, it indicates that the final match from application identification for that session is non-SSL and SSL proxy ignores the session.
- An entry is not found in the application system cache. This can happen on the first session, or when the application system cache has been cleaned or has expired. In such a scenario, SSL proxy cannot wait for the final match (requires traffic in both directions). In SSL proxy, traffic in reverse direction happens only if SSL proxy has initiated an SSL handshake. Initially, for such a scenario SSL proxy tries to leverage prematch or aggressive match results from application identification , and if the results indicate SSL, SSL proxy will go ahead with the handshake.
- Application identification fails due to resource constraints and other errors. Whenever the result from application identification is not available, SSL proxy will assume static port binding and will try to initiate SSL handshake on the session. This will succeed for actual SSL sessions, but it will result in dropped sessions for non SSL sessions.

Copyright © 2016, Juniper Networks, Inc.

Logical Systems Support

It is possible to enable SSL proxy on firewall policies that are configured using logical systems; however, note the following limitations:

- The “services” category is currently not supported in logical systems configuration. Because SSL proxy is under “services,” you cannot configure SSL proxy profiles on a per-logical-system basis.
- Because proxy profiles configured at a global level (within “services ssl proxy”) are visible across logical system configurations, it is possible to configure proxy profiles at a global level and then attach them to the firewall policies of one or more logical systems.

Limitations

The following limitations currently apply to SSL proxy:

- Only TLSv1.0, TLSv1.1, and TLSv1.2 are supported.
- SSL v3.0 is not supported.
- Only X.509v3 certificate is supported
- Client authentication of SSL handshake is not supported.

NOTE: On all high-end SRX Series devices, for a particular session, the SSL proxy is only enabled if a relevant feature related to SSL traffic is also enabled. Features that are related to SSL traffic are IDP, application identification, application firewall, and application tracking. If none of the above listed features are active on a session, the SSL proxy bypasses the session and logs are not generated in this scenario

NOTE: On all SRX Series devices, the current SSL proxy implementation has the following connectivity limitations:

- The SSLv2 protocol is not supported. SSL sessions using SSLv2 are dropped
- SSL sessions where client certificate authentication is mandatory are dropped
- SSL sessions where renegotiation is requested are dropped

Related Documentation

- Understanding Address Books on page 1069
- Understanding Global Address Books on page 1071
- Understanding Self-Signed Certificates on page 6433
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

Application Firewall, IDP, and Application Tracking with SSL Proxy Overview

With the implementation of SSL proxy, AppID can identify applications encrypted in SSL. SSL proxy can be enabled as an application service in a regular firewall policy rule. Intrusion Detection and Prevention (IDP), application firewall (AppFW), and application tracking (AppTrack) services can use the decrypted content from SSL proxy. On the SSL payload, IDP can inspect attacks and anomalies; for example, HTTP chunk length overflow on HTTPS. On encrypted applications, such as Facebook, AppFW can enforce policies and AppTrack (when configured in the from and to zones) can report logging issues based on dynamic applications.

NOTE: If none of the services (AppFW, IDP, or AppTrack) are configured, then SSL proxy services are bypassed even if an SSL proxy is attached to a firewall policy.

NOTE: The IDP module will not perform an SSL inspection on a session if an SSL proxy is enabled for that session. That is, if both SSL inspection and SSL proxy are enabled on a session, SSL proxy will always take precedence.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Example: Configuring Application Firewall When SSL Proxy Is Enabled on page 568

Configuring SSL Proxy

SSL proxy works transparently between the client and the server. All requests from a client first go to the proxy server; the proxy server evaluates the request, and if the request is valid, forwards the request to the outbound side. Similarly, inbound requests are also evaluated by the proxy server. Both client and server interpret that they are communicating with each other; however, it is the SSL proxy that functions between the two.

SSL proxies provide encryption and decryption by residing between the server and the client. Because SSL proxies are hidden from both the server and the client, secret keys are shared between the two to decrypt the SSL traffic. Proxies are known as *forward proxies* because proxy servers are used to hide any detailed information from the servers.

Integrity, confidentiality, and authenticity of traffic are validated through PKI, which includes digital certificates issued by the CA, certificate validity and expiration dates, details about the certificate owner and issuer, and security policies.

- SSL Proxy Configuration Overview on page 541
- Configuring a Root CA Certificate on page 542
- Configuring a CA Profile Group on page 543

Copyright © 2016, Juniper Networks, Inc.

- Configuring a Trusted CA Profile on page 544
- Importing a Root CA Certificate into a Browser on page 545
- Applying an SSL Proxy Profile to a Security Policy on page 546
- Creating a Whitelist of Exempted Destinations on page 547
- Configuring SSL Proxy Logging on page 549
- Configuring Ciphers on page 549
- Exporting Certificates to a Specified Location on page 550
- Ignoring Server Authentication on page 550

SSL Proxy Configuration Overview

Figure 31 displays an overview of how SSL proxy is configured. It includes some required steps, such as configuring the root CA certificate, loading a CA profile group, and applying an SSL proxy profile to a security policy, and some optional steps, such as creating whitelists and SSL proxy logging.

Figure 31: SSL Proxy Configuration Overview

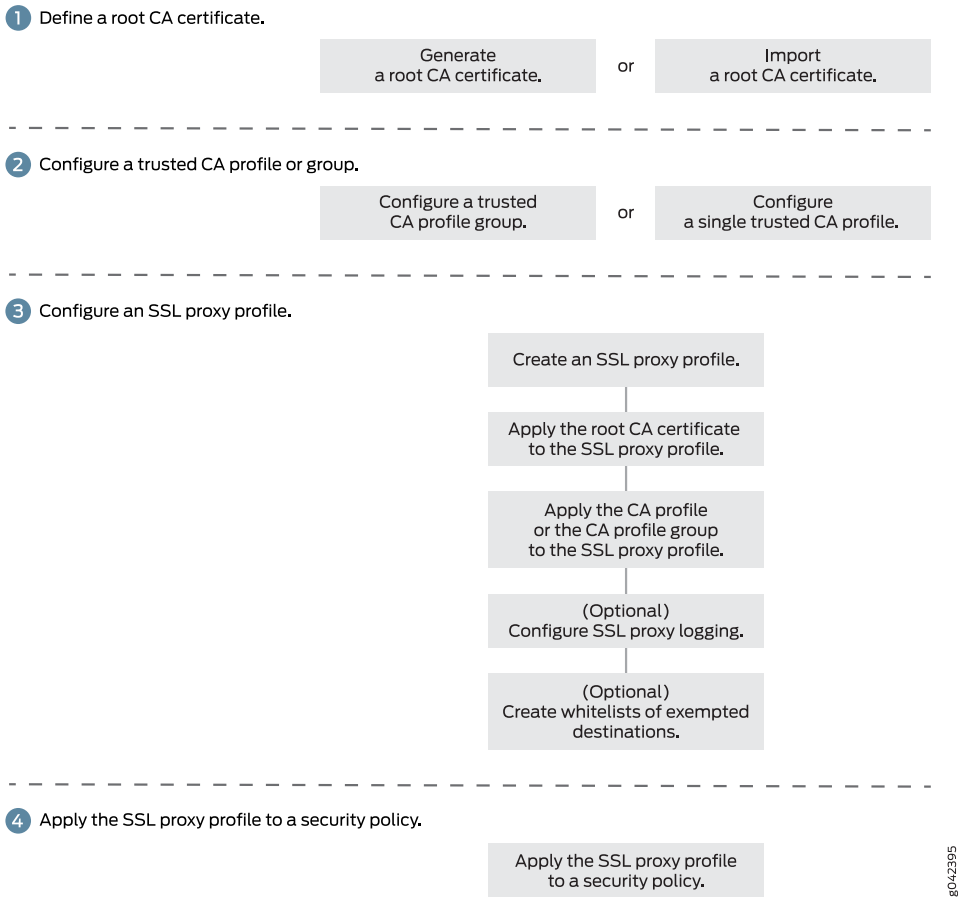

Configuring a Root CA Certificate

A CA can issue multiple certificates in the form of a tree structure. A root certificate is the topmost certificate of the tree, the private key of which is used to *sign* other certificates. All certificates immediately below the root certificate inherit the signature or trustworthiness of the root certificate. This is somewhat like the *notarizing* of an identity.

You can configure a root CA certificate by first obtaining a root CA certificate (by either generating a self-signed one or importing one) and then applying it to an SSL proxy profile. There are two ways you can obtain a root CA certificate—by using the Junos OS CLI on an SRX Series device or by using OpenSSL on a UNIX device.

To generate a root CA certificate using the Junos OS CLI, follow these steps on an SRX Series device:

1. From operational mode, generate a PKI public/private key pair for a local digital certificate.

   user@host>**request security pki generate-key-pair certificate-id** ***certificate-id*** **size** ***size*** **type** ***type***

2. From operational mode, define a self-signed certificate. Specify certificate details such as the certificate identifier (generated in the previous step), a fully qualified domain name (FQDN) for the certificate, and an e-mail address of the entity owning the certificate. You can also specify other information such as the common name and the organization involved. By configuring the **add-ca-constraint** option, you make sure that the certificate can be used for signing other certificates.

   user@host>**request security pki local-certificate generate-self-signed certificate-id** ***certificate-id*** **domain-name** ***domain-name*** **subject** ***subject*** **email** ***email-id*** **add-ca-constraint**

3. From configuration mode, apply the loaded certificate as root-ca in the SSL proxy profile.

   [edit]
   user@host# **set services ssl proxy profile** ***profile-name*** **root-ca** ***certificate-id***

4. Import the root CA as a trusted CA into client browsers. This is required for the client browsers to trust the certificates signed by the SRX Series device. See "Importing a Root CA Certificate into a Browser" on page 545.

To generate a root CA certificate using OpenSSL, follow these steps on a UNIX device:

1. Create folders **keys** and **certs**.

   mkdir /etc/pki/tls/keys
   mkdir /etc/pki/tls/certs

2. Change to the **openssl** directory.

   cd /etc/pki/tls

3. Create a CA certificate key. The following command creates an RSA key using the 3DES encryption named **ca.key** that is 2048 in length. You also need to enter a

Copyright © 2016, Juniper Networks, Inc.

password that is used to encrypt the private key. This is critical to security if the key is lost because it will still be encrypted.

```
% openssl genrsa -des3 -out keys/ssl-proxy-ca.key 2048
```

4. Create a CA certificate based on the CA private key (created in the previous step). The expiration date for this certificate is 3 years or 1095 days. However, you can set it to a different value. When creating the certificate, you need to enter the password and the certificate information that includes distinguished name (DN), country name, and so forth.

   ```
   % openssl req -new -x509 -days 1095 –key keys/ssl-proxy-ca.key -out
     certs/ssl-inspect-ca.cer
   ```

5. Import the CA private and public keys into the SRX Series device. Copy the **ca.key** and **ca.cer** keys to the **/var/tmp** directory on the SRX Series device. You can copy using SCP, or open the files and copy them into “vi” on the SRX Series device to create new files.

   ```
   user@host> request security pki local-certificate load certificate-id ssl-inspect-ca key
     /var/tmp/ssl-inspect-ca.key filename /var/tmp/ssl-inspect-ca.cer passphrase
     password
   ```

6. From configuration mode, apply the loaded certificate as root-ca in the SSL proxy profile.

   ```
   [edit]
   user@host# set services ssl proxy profile ssl-inspect-profile root-ca ssl-inspect-ca
   ```

7. Import the root CA as a trusted CA into client browsers. This is required for the client browsers to trust the certificates signed by the SRX Series device. See “Importing a Root CA Certificate into a Browser” on page 545.

Configuring a CA Profile Group

The CA profile defines the certificate information to be used for authentication. It includes the public key that SSL proxy uses when generating a new certificate. Junos OS allows you to create a group of CA profiles and load multiple certificates in one action, view information about all certificates in a group, and delete unwanted CA groups.

You can load a group of CA profiles by obtaining a list of trusted CA certificates, defining a CA group, and attaching the CA group to the SSL proxy profile.

1. Obtain a list of trusted CA certificates by following one of these methods:

   - Junos OS provides a default list of trusted CA certificates that you can load on your system using the **default** command option. The Junos OS package contains the default CA certificates as a PEM file (for example, **trusted_CA.pem**). After you download the Junos OS package and reboot your device, the default certificates are available on your system.

     From operational mode, load the default trusted CA certificates (the group name identifies the CA profile group):

     ```
     user@host> request security pki ca-certificate ca-profile-group load ca-group-name
       group-name filename default
     ```

Copyright © 2016, Juniper Networks, Inc.

- Alternatively, you can define your own list of trusted CA certificates and import them on your system. You get the list of trusted CAs in a single PEM file (for example **IE-all.pem**) and save the PEM file in a specific location (for example, **/var/tmp**). See Knowledge Base Article KB23144.

  From operational mode, load the trusted list to the device (the group name identifies the CA profile group):

  user@host> **request security pki ca-certificate ca-profile-group load ca-group-name** ***group-name*** **filename** */var/tmp/IE-all.pem*

2. From configuration mode, attach the CA profile group to the SSL proxy profile. You can attach one or all CA profile groups at a time:

   - To attach one CA profile group (the group name identifies the CA profile group):

     [edit]
     user@host# **set services ssl proxy profile** ***profile-name*** **trusted-ca** ***group-name***

   - To attach all CA profile groups:

     [edit]
     user@host# **set services ssl proxy profile** ***profile-name*** **trusted-ca all**

You can easily display information about all certificates in a CA profile group:

user@host> **show security pki ca-certificates ca-profile-group** *group-name*

You can delete a CA profile group. Remember that deleting a CA profile group deletes all certificates that belong to that group:

user@host> **clear security pki ca-certificates ca-profile-group** *group-name*

Configuring a Trusted CA Profile

Typically, you import a list of trusted CA certificates by creating a group of CA profiles. However, you can also configure a single CA profile (containing one or multiple certificates) and import it using PKI commands. This section shows you how to import a trusted CA certificate from your browser's certificate store into your SRX Series device. The certificate that is configured under the trusted CA is loaded using the PKI commands and is used for validating the server certificate chain.

1. From configuration mode, configure the CA profile used for loading the certificate.

   [edit]
   user@host# **set security pki ca-profile** ***profile-name*** **ca-identity** ***ca-identity***

2. From operational mode, load the certificate using PKI commands.

   user@host> **request security pki ca-certificate load ca-profile** ***profile-name*** **filename** ***filename***

3. From configuration mode, disable the revocation check.

NOTE: CRL checks are not supported; we recommend that you disable revocation checks.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security pki ca-profile profile-name ca-identity ca-identity
  revocation-check disable
```

4. From configuration mode, configure the loaded certificate as a trusted CA in the SSL proxy profile.

```
[edit]
user@host# set services ssl proxy profile ssl-proxy-profile-name trusted-ca
  ca-profile-name
```

NOTE: More than one trusted CA can be configured for a profile.

5. (Optional) If you have multiple trusted CA certificates, you do not have to specify each trusted CA separately. You can load *all* the trusted CA certificates using the following command from configuration mode.

```
[edit]
user@host# set services ssl proxy profile ssl-proxy-profile-name trusted-ca all
```

NOTE: Alternatively, you can import a set of trusted CAs from your browser into the SRX Series device. See Knowledge Base article KB23144.

Importing a Root CA Certificate into a Browser

In order to have your browser or system automatically trust all certificates signed by the root CA configured in the SSL proxy profile, you must instruct your platform or browser to trust the CA root certificate.

To import a root CA certificate:

1. Generate a PEM format file for the configured root CA.

```
request security pki local-certificate export certificate-id root-ca type pem filename
  path/file-name.pem
```

2. Import a root CA certificate into a browser.

   From Internet Explorer (version 8.0):

   a. From the Tools menu, choose **Internet Options**.

   b. On the Content tab, click **Certificates**.

Copyright © 2016, Juniper Networks, Inc.

c. Select the **Trusted Root Certification Authorities** tab and click **Import**.

d. In the Certificate Import Wizard, navigate to the required root CA certificate and select it.

From Firefox (version 39.0):

a. From the Tools menu, choose **Options**.

b. From the Advanced menu, select the **Certificates** tab and click **View Certificate**.

c. In the Certificate Manager window, select the **Authorities** tab and click **Import**.

d. Navigate to the required root CA certificate and select it.

From Google Chrome (45.0):

a. From the Settings menu, choose **Show Advanced Settings**.

b. From the Advanced menu, select the **Certificates** tab and click **View Certificate**.

c. Under HTTPS/SSL, click **Manage Certificates**.

d. In the Certificate window, select **Trusted Root Certification Authorities** and click **Import**.

e. In the Certificate Import Wizard, navigate to the required root CA certificate and select it.

Applying an SSL Proxy Profile to a Security Policy

SSL proxy is enabled as an application service within a security policy. In a security policy, you specify the traffic that you want the SSL proxy enabled on as match criteria and then specify the SSL proxy CA profile to be applied to the traffic. Figure 32 displays a graphical view of SSL proxy profile and security policy configuration.

Copyright © 2016, Juniper Networks, Inc.

Figure 32: Applying an SSL Proxy Profile to a Security Policy

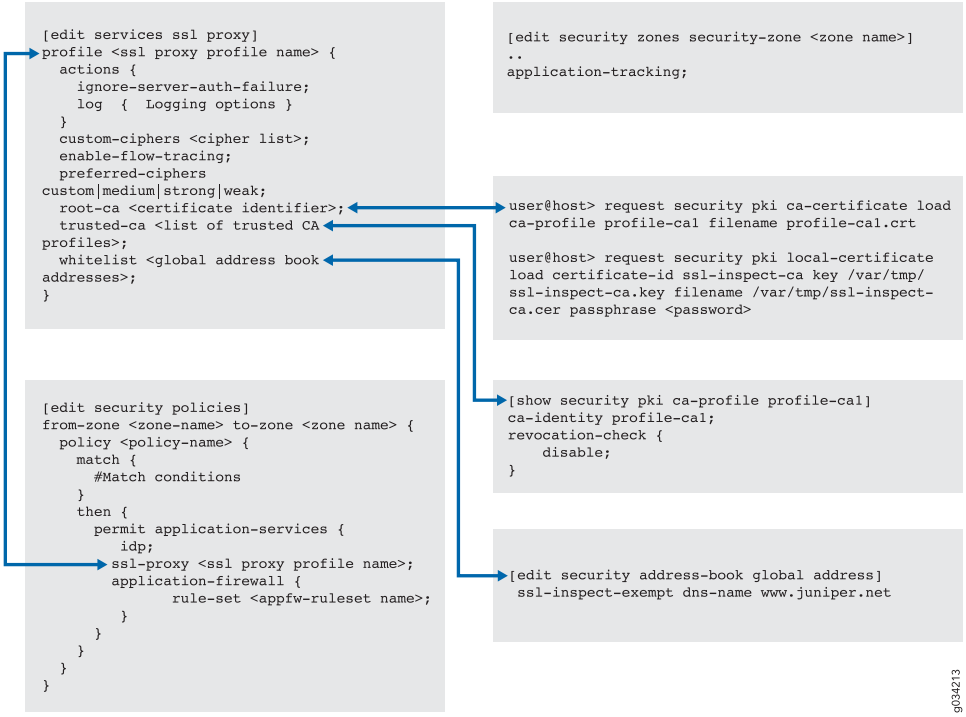

To enable SSL proxy in a security policy:

1. Create a security policy and specify the match criteria for the policy. As match criteria, specify the traffic for which you want to enable SSL proxy.

   [edit]
   user@host# **set security policies from-zone trust to-zone untrust policy *policy-name* match source-address *source-address***
   user@host# **set security policies from-zone trust to-zone untrust policy *policy-name* match destination-address *destination-address***
   user@host# **set security policies from-zone trust to-zone untrust policy *policy-name* match application *application***

2. Apply the SSL proxy profile to the security policy.

   [edit]
   user@host# **set security policies from-zone trust to-zone untrust policy *policy-name* then permit application-services ssl-proxy profile-name *profile-name***

Creating a Whitelist of Exempted Destinations

SSL encryption and decryption are complicated and expensive procedures. You can selectively bypass SSL proxy processing for some sessions by configuring a whitelist. Typically, you would configure the whitelist to include trusted servers or domains with which you are very familiar. You might also include financial and banking sites that you are legally required to include.

Copyright © 2016, Juniper Networks, Inc.

Whitelists include addresses that you want to exempt from undergoing SSL proxy processing. For example, if you want to exempt all sessions to www.example.net then you would include it in the whitelist. To configure the whiltelist, you specify the domain that you want to exempt in an address book and then configure the address in the SSL proxy profile.

1. Configure the domain in the address book.

   ```
   [edit]
   user@host# set security address-book global address address dns-name
     www.example.net
   ```

2. Specify the global address book address in the SSL proxy profile.

   ```
   [edit]
   user@host# set services ssl proxy profile profile-name whitelist address
   ```

Whitelist addresses and address sets are created under the global address book. The following type of addresses (from the global address book) are supported:

- IPv4 addresses (plain text). For example:

  ```
  [edit]
  user@host# set security address-book global address address-name ipv4-prefix
  ```

- IPv4 address range. For example:

  ```
  [edit]
  user@host# set security address-book global address address-name range-address
    range-low to range-high
  ```

- IPv4 wildcard. For example:

  ```
  [edit]
  user@host# set security address-book global address address-name wildcard-address
    a.d.d.r/netmask
  ```

  Noncontiguous netmasks are not supported. For example:

  - 5.0.0.1/255.255.0.0 is supported.
  - 5.0.0.1/255.255.0.255 is NOT supported.

- IPv6 address (plain text). For example:

  ```
  [edit]
  user@host# set security address-book global address address-name ipv6-prefix
  ```

- DNS name. For example:

  ```
  [edit]
  user@host# set security address-book global address address-name dns-name
    domain-name
  ```

- Translated IP addresses. Sessions are whitelisted based on the actual IP address and not on the translated IP address. Because of this, in the whitelist configuration of the SSL proxy profile, the actual IP address should be provided and not the translated IP addresses.

Copyright © 2016, Juniper Networks, Inc.

For example, consider a destination NAT rule that translates destination IP address **192.0.2.10/24** to **198.51.100.8/24** using the following commands:

```
[edit]
user@host# set security nat destination pool d1 address 198.51.100.8/24
user@host# set security nat destination rule-set dst-nat rule r1 match
  destination-address 192.0.2.10/24
user@host# set security nat destination rule-set dst-nat rule r1 then destination-nat
  pool d1
```

In this scenario, to exempt a session from SSL proxy inspection, the following IP address should be added to the whitelist:

```
[edit]
user@host# set security address-book global address ssl-proxy-exempted-addr
  192.0.2.10/24
user@host# set services ssl proxy profile ssl-inspect-profile whitelist
  ssl-proxy-exempted-addr
```

Configuring SSL Proxy Logging

When configuring SSL proxy, you can choose to set the option to receive some or all of the logs. SSL proxy logs contain the logical system name, SSL proxy whitelists, policy information, SSL proxy information, and other information that helps you troubleshoot when there is an error.

You can configure logging of *all* or specific events, such as error, warning, and information events. You can also configure logging of sessions that are whitelisted, dropped, ignored, or allowed after an error occurs.

```
[edit]
user@host# set services ssl proxy profile profile-name actions log all
user@host# set services ssl proxy profile profile-name actions log sessions-whitelisted
user@host# set services ssl proxy profile profile-name actions log sessions-allowed
user@host# set services ssl proxy profile profile-name actions log errors
```

You can use **enable-flow-tracing** option to enable debug tracing.

Configuring Ciphers

You can configure the following ciphers for an SSL proxy profile:

- **preferred-ciphers**–Preferred ciphers allow you to define an SSL cipher that can be used with acceptable key strength. Ciphers are divided in three categories depending on their key strength: strong, medium, or weak.
- **custom-ciphers**–Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories, you can select ciphers from each of the categories to form a custom cipher set. To configure custom ciphers, you must set **preferred-ciphers** to custom.

  The following example shows how to create a custom cipher. In this example, you set **prefered-cipher** to custom and add the cipher list (rsa-with-3des-ede-cbc-sha and rsa-with-aes-256-cbc-sha):

  ```
  set services ssl proxy profile profile-name preferred-ciphers custom
  ```

Copyright © 2016, Juniper Networks, Inc.

```
set services ssl proxy profile profile-name custom-ciphers rsa-with-3des-ede-cbc-sha
set services ssl proxy profile profile-name custom-ciphers rsa-with-aes-256-cbc-sha
```

Exporting Certificates to a Specified Location

When a self-signed certificate is generated using a PKI command, the newly generated certificate is stored in a predefined location (**var/db/certs/common/local**).

Use the following command to export the certificate to a specific location (within the device). You can specify the certificate ID, the filename, and the type of file format (DER/PEM):

```
user@host> request security pki local-certificate export certificate-id certificate-id
user@host> request security pki local-certificate export filename filename
user@host> request security pki local-certificate export type der
```

Ignoring Server Authentication

Junos OS allows you to configure an option to ignore server authentication completely. If you configure your system to ignore authentication, then any errors encountered during server certificate verification at the time of the SSL handshake are ignored. Commonly ignored errors include the inability to verify CA signature, incorrect certificate expiration dates, and so forth. If this option is not set, all the sessions where the server sends self-signed certificates are dropped when errors are encountered.

We do not recommend using this option for authentication because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause of dropped SSL sessions.

From configuration mode, specify to ignore server authentication:

```
[edit]
user@host# set services ssl proxy profile profile-name actions ignore-server-auth-failure
```

Related Documentation

- SSL Proxy Overview on page 531
- Enabling Debugging and Tracing for SSL Proxy on page 552
- Understanding Self-Signed Certificates on page 6433
- show services ssl proxy statistics on page 829
- clear services ssl proxy statistics on page 752

Working with the Certificate Revocation Lists for SSL Proxy

A certificate issued by a certificate authority (CA) is supposed to be valid untill the expiration of the validity period. In the normal course of business, a CA can revoke an issued certificate. A certificate is revoked if it is suspected that the certificate has been compromised. Some of the examples are:

- Unspecified (no particular reason is given).
- Private key associated with the certificate was compromised.

Copyright © 2016, Juniper Networks, Inc.

- Private key associated with the CA that issued the certificate was compromised.
- The owner of the certificate is no longer affiliated with the issuer of the certificate and does not have rights to access the certificate or does not require it any longer.
- Another certificate replaces the original certificate.
- The CA that issued the certificate has ceased to operate.
- The certificate is on hold pending further action. It is treated as revoked but might be accepted in the future.

Once the CA determines to revoke a certificate, it publishes the information by some means so that the enduser certificate can use the information to validate a certificate. The CA can publish this information using certificate revocation list (CRL).

The CRL contains the list of digital certificates that have been canceled before their expiration date. When a participating device uses a digital certificate, it checks the certificate signature and validity. It also acquires the most recently issued CRL and checks that the certificate serial number is not on that CRL. By default, CRL verification is enabled on SSL proxy profile.

CRL validation on SRX Series device involves checking for revoked certificates from servers. You can enable or disable the CRL validation to meet your specific security requirements.

- Disabling CRL Verification on page 551
- Allowing Sessions When CRL Information Is Not Available on page 551
- Allowing Sessions When CRL Status Is Unknown on page 552

Disabling CRL Verification

In order to enhance security, the certificate revocation checking feature has been enabled by default on SRX Series devices on any SSL proxy profile. You can enable or disable the CRL validation to meet your specific security requirements.

- To disable CRL verification:

```
[edit]
user@host# set services ssl proxy profile profile-name actions crl disable
```

You can reenable CRL validation by using the **delete services ssl proxy profile profile-name actions crl disable** command.

Allowing Sessions When CRL Information Is Not Available

Sometimes CRL information might not be available because of various reasons. For example:

- CRL download failed and the PKI daemon did not or could not fetch the CRL from the CA.
- The CRL path was not available from the configuration and it is not present in the root or intermediate certificate, or no URL was configured.

Copyright © 2016, Juniper Networks, Inc.

You can allow or drop the sessions when a CRL information is not available.

- To ensure that the sessions are not dropped for any reason when CRL information is not available:

  ```
  [edit]
  user@host# edit set services ssl proxy profile profile-name actions crl if-not-present
    allow
  ```

- To drop the sessions when CRL information is not available:

  ```
  [edit]
  user@host# edit set services ssl proxy profile profile-name actions crl if-not-present
    drop
  ```

Allowing Sessions When CRL Status Is Unknown

You can configure how an SRX Series device will respond when updated CRL information is not available, and the server certificate that is currently offered is not known to be revoked from a previous query. Certificates are presumed not to be revoked, by default, which means they are valid, and a temporary failure to obtain a CRL does not automatically result in an SSL handshake failure. By default, sessions are allowed if CRL status is unknown.

You can configure an SRX Series device to accept a certificate without a reliable confirmation available on the revocation status.

- To allow the sessions when a certificate is revoked and the revocation reason is on hold:

  ```
  [edit]
  user@host# edit set services ssl proxy profile profile-name actions crl
    ignore-hold-instruction-code
  ```

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540

Enabling Debugging and Tracing for SSL Proxy

Debug tracing on both Routing Engine and the Packet Forwarding Engine can be enabled for SSL proxy by setting the following configuration:

```
user@host# set services ssl traceoptions
```

Table 37 shows the supported levels for trace options.

Table 37: Trace Levels

Cause Type	Description
Brief	Only error traces on both the Routing Engine and the Packet Forwarding Engine.

Copyright © 2016, Juniper Networks, Inc.

Table 37: Trace Levels *(continued)*

Detail	Packet Forwarding Engine–Only event details up to the handshake should be traced. Routing Engine–Traces related to commit. No periodic traces on the Routing Engine will be available
Extensive	Packet Forwarding Engine–Data transfer summary available. Routing Engine–Traces related to commit (more extensive). No periodic traces on the Routing Engine will be available.
Verbose	All traces are available.

Table 38 shows the flags that are supported.

Table 38: Supported Flags in Trace

Cause Type	Description
cli-configuration	Configuration-related traces only.
initiation	Enable tracing on the SSL-I plug-in.
proxy	Enable tracing on the SSL-Proxy-Policy plug-in.
termination	Enable tracing on the SSL-T plug-in.
selected-profile	Enable tracing only for profiles that have **enable-flow-tracing** set.

You can enable logs in the SSL proxy profile to get to the root cause for the drop. The following errors are some of the most common:

- Server certification validation error. Check the trusted CA configuration to verify your configuration.
- System failures such as memory allocation failures.
- Ciphers do not match.
- SSL versions do not match.
- SSL options are not supported.
- Root CA has expired. You need to load a new root CA.

You can enable the **ignore-server-auth-failure** option in the SSL proxy profile to ensure that certificate validation, root CA expiration dates, and other such issues are ignored. If sessions are inspected after the **ignore-server-auth-failure** option is enabled, the problem is localized.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 29

Configuring Application Firewall

- Application Firewall Overview on page 555
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560
- Example: Configuring an Application Group for Application Firewall on page 564
- Example: Configuring Application Firewall When SSL Proxy Is Enabled on page 568

Application Firewall Overview

Traditionally, applications like HTTP, SMTP, and DNS use well-known standard ports and are easily controlled by a stateful firewall. However, it is possible to run these applications on any port as long as the client and server are using the same protocol as the well-known ports.

Evasive applications could remain undetected with a standard firewall that functions at Layer 3 or Layer 4 by transmitting other protocols over these well-known ports that are usually open by a firewall. AppFW enforces protocol and policy control at Layer 7. It inspects the actual content of the payload and ensures that it conforms to the policy, rather than identifying the application based on Layer 3 and Layer 4 information.

Additionally, with the growing popularity of Web applications and the shift from traditional full client-based applications to the Web, more and more traffic is being transmitted over HTTP. An application firewall identifies not only HTTP but also any application running on top of it, letting you properly enforce policies. For example, an application firewall rule could block HTTP traffic from Facebook but allow Web access to HTTP traffic from MS Outlook.

A security administrator implements an application firewall by performing the following tasks:

- Define one or more application firewall rule sets.
- Create rules for each rule set that permit, reject, or deny traffic based on the application ID.
- Configure a security policy to invoke the application firewall service and specify the rule set to be applied to permitted traffic.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Understanding Application Firewall Rule Sets on page 556
- Configuring an Application Firewall Within a Security Policy on page 557
- Application Group Support for Application Firewall on page 557
- Redirecting Users on page 558
- Session Logging for Application Firewalls on page 559
- Application Firewall Support in Chassis Cluster on page 559

Understanding Application Firewall Rule Sets

An application firewall permits, rejects, or denies traffic based on the application of the traffic. The firewall consists of one or more rule sets with rules that specify match criteria, including dynamic applications, and the action to be taken for matching traffic.

An application firewall rule set consists of:

- The name of the rule set
- One or more rules
- A single default rule

Each rule defines dynamic applications to permit, reject, or deny. Each rule consists of:

- The name of the rule
- A list of dynamic applications to be used as match criteria
- The action to take for any traffic that matches one of the specified applications
  - Reject—Notify the client, drop the traffic, close the session, and log the event.
  - Deny—Drop the traffic, close the session, and log the event.
  - Permit—Permit the traffic.

The default rule defines the action to be taken for any traffic that does not match one of the rules. An application firewall rule set must contain a default rule.

There is no limit to the number of dynamic applications in a rule or to the number of rules in a rule set. However, there is a limit to the overall number of rule sets and rules.

The junos:UNKNOWN keyword is reserved for unknown dynamic applications. In the following cases, the application ID is set to junos:UNKNOWN:

- The traffic does not match an application signature in the database.
- The system encounters an error when identifying the application.
- The session fails over to another device.

Traffic with an application ID of junos:UNKNOWN matches a rule with a dynamic application of junos:UNKNOWN. If there is no rule defined for junos:UNKNOWN, the default rule is applied.

Copyright © 2016, Juniper Networks, Inc.

Configuring an Application Firewall Within a Security Policy

An application firewall is invoked using the **then permit** statement of the security policy.

Any traffic denied or rejected by the security policy based on Layer 3 or Layer 4 criteria is dropped immediately. Traffic permitted by the security policy is further assessed by the application firewall at Layer 7 based on its application ID.

The following sample policy, outbound-traffic, permits matching HTTP traffic, and invokes application services and an application firewall. The rule set, unknown-traffic, permits, denies, or rejects, traffic based on its match criteria.

```
[edit security policies from-zone trust to-zone untrust outbound-traffic]
user@host# set match source-address 1.1.1.0
user@host# set match destination-address 2.2.2.0
user@host# set match application junos-http
user@host# set then permit application-services application-firewall rule-set
  unknown-traffic
```

Traffic is processed in the following sequence:

1. Match the zone pair specified in the policy.
2. When specified, match the source and destination IP addresses, ports, and application type.
3. Apply the security policy action to matching traffic.
   - Reject—Notify the client, drop the traffic, and log the event.
   - Deny—Drop the traffic, and log the event.
   - Permit—Open a session, log the event, and apply services as specified.
     - Invoke application services to retrieve the application ID for the traffic.
     - Apply the specified application firewall rule set.

NOTE: All IP fragmented packets received on the SRX Series device must be reassembled before forwarding.

Application Group Support for Application Firewall

Application group support associates related applications under a single name for simplified, consistent reuse when using any application services. As the predefined signature database changes, the content of a predefined application group can be modified to include new signatures without affecting existing firewall rules. When you define application firewall rules, you can specify dynamic application groups as match criteria.

Copyright © 2016, Juniper Networks, Inc.

NOTE: An application group can contain applications and groups simultaneously. It is possible to assign one application to multiple groups. There is no limit to the number of dynamic application groups contained in one rule.

For information on creating or listing application groups, see "Customizing Application Groups for Junos OS Application Identification" on page 517.

NOTE: On all SRX Series devices, when ALG is enabled, application identification includes the ALG result to identify the application of the control sessions. Application firewall permits ALG data sessions whenever control sessions are permitted. If the control session is denied, there will be no data sessions. When ALG is disabled, application identification relies on its signatures to identify the application of the control and data sessions. If a signature match is not found, the application is considered unknown. Application firewall handles applications based on the application identification result.

Redirecting Users

Although drop and reject actions are logged, application firewall does not notify clients when either action is taken. Clients are not aware that the webpage is not available and might keep trying to access the page. To provide an explanation for the action or to redirect the client to an informative webpage, use the **block-message** option with the **reject** or **deny** action in an application firewall rule.

```
...
then reject block-message
```

When traffic is rejected by the application firewall rule, a splash screen with the following default message is displayed to the user:

user-name, Application Firewall has blocked your request to application *application-name* at *dst-ip:dst-port* accessed from *src-ip:src-port*.

To help the user fully understand which request has been rejected or denied, the default message includes traffic-specific details, such as the username, application, and address information.

You can customize the redirect action by including additional text on the splash screen or by specifying a URL to which the user is redirected. To customize the block message, define the type and content in a block message profile defined in the rule set:

```
[edit security application-firewall profile deny-profile-1]
set block-message type custom-redirect-url content http://abc.company.com/information
```

The block message profile is identified for the rule set, and applied to one or more of the rules using the **block-message** option.

```
[edit security application-firewall rule-sets application-firewall-3]
```


Copyright © 2016, Juniper Networks, Inc.

```
set profile deny-profile-1
set rule redirect-on-deny
set match dynamic-application [junos:KAZZA junos:EDONKEY junos:YSMG]
set then deny block-message
```

In this example, any traffic matching one of the specified dynamic applications is denied, and the block message defined for rule set, deny-profile-1, is applied. Based on the profile for deny-profile-1, the user is redirected to the URL http://abc.company.com/information for further details.

Session Logging for Application Firewalls

With security policies, the permit action of the matched policy rule creates a session and logs a session create message. A reject or deny action logs a reject or deny message, but does not create a session.

When an application firewall is implemented, the permit action of the security policy creates a session before the application firewall rules are applied. If the dynamic application have been retrieved from the cache, this information is added to the session create message. If the application is in the process of being identified, the dynamic application fields specify UNKNOWN.

If traffic is rejected or denied by the application firewall, application firewall also closes the session. The reject or deny message actions are logged with the reason field containing one of the following phrases:

- **appfw deny** or **appfw deny redirect**
- **appfw reject** or **appfw reject redirect**
- **policy deny**
- **policy reject**

Application Firewall Support in Chassis Cluster

When the application ID is not identified during failover sessions, the ID is considered an unknown application ID. During this session, the traffic is processed based on the action defined in a rule specified for unknown. If there is no rule defined for unknown, then the default rule is applied.

NOTE: When an SRX Series device is operating in chassis cluster mode and application identification is enabled, pre-match state application IDs are not synced to other node. If there are any failover sessions, which were still under classification, will not have any application IDs assigned. This could result in application statistics and counters mismatch.

When the application ID is identified before sessions fail over, the same action taken before the failover is effective after the failover. The application firewall action taken before and after the failover depends on the application ID state, as shown in Table 39.

Copyright © 2016, Juniper Networks, Inc.

Table 39: Application Firewall Actions

Before Failover		After Failover	
Application ID State	Application Firewall Action	Application ID State	Application Firewall Action
Success	Deny	Success	Deny
Success	Permit	Success	Permit
Pending	—	UNKNOWN	Action based on the rule defined for unknown application

NOTE: In-service software upgrade (unified ISSU) is not supported due to lack of chassis cluster infrastructure support. Thus, the failover event is controlled through the application firewall policy by allowing or denying the unknown dynamic applications.

Related Documentation

- Example: Configuring an Application Group for Application Firewall on page 564
- Understanding Application Identification Techniques on page 485
- *Building Blocks Feature Guide for Security Devices*

Example: Configuring Application Firewall Rule Sets Within a Security Policy

This example shows how to configure application firewall rule sets within the security policy.

- Requirements on page 560
- Overview on page 560
- Configuration on page 561
- Verification on page 563

Requirements

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book with addresses for the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.

Overview

In Junos OS, the security policies provide firewall security functionality by enforcing rules for the traffic so that traffic passing through the device is permitted or denied based on the action defined in the rules. The application firewall support in the policies provides additional security control for dynamic applications.

Copyright © 2016, Juniper Networks, Inc.

The application firewall is defined by a collection of rule sets. These rule sets can be defined independently and shared across network security policies. A rule set defines the rules that match the application ID detected, based on the application signature.

This configuration example shows how to:

- Permit or deny selected traffic from the untrust zone to the trust zone, based on the application firewall rule sets defined with the rules matching the dynamic applications.

NOTE: On all SRX Series devices, J-Web pages for AppSecure Services are preliminary. We recommend using CLI for configuration of AppSecure features.

Configuration

CLI Quick Configuration

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

```
set security policies from-zone untrust to-zone trust policy policy1 match source-address
  1.1.1.0
set security policies from-zone untrust to-zone trust policy policy1 match
  destination-address 2.2.2.0
set security policies from-zone untrust to-zone trust policy policy1 match application
  junos-http
set security policies from-zone untrust to-zone trust policy policy1 then permit
  application-services application-firewall rule-set rs1
set security policies from-zone untrust to-zone trust policy policy2 match source-address
  1.1.1.0
set security policies from-zone untrust to-zone trust policy policy2 match
  destination-address 2.2.2.0
set security policies from-zone untrust to-zone trust policy policy2 match application any
set security policies from-zone untrust to-zone trust policy policy2 then permit
  application-services application-firewall rule-set rs2
set security application-firewall rule-sets rs1 rule r1 match dynamic-application
  [junos:KAZZA junos:EDONKEY junos:YSMG]
set security application-firewall rule-sets rs1 rule r1 then deny
set security application-firewall rule-sets rs1 default-rule permit
set security application-firewall rule-sets rs2 rule r1 match dynamic-application
  [junos:FACEBOOK-ACCESS junos:GOOGLE-TALK junos:MEEBO junos:UNKNOWN]
set security application-firewall rule-sets rs2 rule r1 then permit
set security application-firewall rule-sets rs2 default-rule deny
```

Step-by-Step Procedure

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

To configure two security policies with application firewall rule sets that permit or deny traffic from different dynamic applications:

1. Configure a policy to process the traffic that goes to the HTTP static ports with the application firewall rule set rs1.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies from-zone untrust to-zone trust policy policy1]
user@host# set match source-address 1.1.1.0
user@host# set match destination-address 2.2.2.0
user@host# set match application junos-http
user@host# set then permit application-services application-firewall rule-set rs1
```

2. Configure another policy to process any traffic that does not go to the HTTP static ports with the application firewall rule set rs2.

```
[edit security policies from-zone untrust to-zone trust policy policy2]
user@host# set match source-address 1.1.1.0
user@host# set match destination-address 2.2.2.0
user@host# set match application any
user@host# set then permit application-services application-firewall rule-set rs2
```

3. Define the application firewall rule set rs1 to deny traffic from selected dynamic applications.

```
[edit security application-firewall rule-sets rs1]
user@host# set rule r1 match dynamic-application [junos:KAZZA junos:EDONKEY
  junos:YSMG]
user@host# set rule r1 then deny
user@host# set default-rule permit
```

4. Define the application firewall rule set rs2 to permit traffic from selected dynamic applications.

```
[edit security application-firewall rule-sets rs2]
user@host# set rule r1 match dynamic-application [junos:FACEBOOK-ACCESS
  junos:GOOGLE-TALK junos:MEEBO junos:UNKNOWN]
user@host# set rule r1 then permit
user@host# set default-rule deny
```

Results From configuration mode, confirm your configuration by entering the **show security policies** and **show security application-firewall** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
  from-zone untrust to-zone trust {
    policy 1 {
      match {
        source-address 1.1.1.0;
        destination-address 2.2.2.0;
        application junos-http;
      }
      then {
        permit {
          application-services {
            application-firewall {
              rule-set rs1;
            }
          }
        }
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    policy 2 {
      match {
        source-address 1.1.1.0;
        destination-address 2.2.2.0;
        application any;
      }
      then {
        permit {
          application-services {
            application-firewall {
              rule-set rs2;
            }
          }
        }
      }
    }
  }
user@host# show security application-firewall
  rule-sets rs1 {
    rule r1 {
      match {
        dynamic-application [junos:KAZZA junos:EDONKEY junos:YSMG];
      }
      then {
        deny;
      }
    }
    default-rule {
      permit;
    }
  }
  rule-sets rs2 {
    rule r1 {
      match {
        dynamic-application [junos:FACEBOOK-ACCESS junos:GOOGLE-TALK
          junos:MEEBO junos:UNKNOWN];
      }
      then {
        permit;
      }
    }
    default-rule {
      deny;
    }
  }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Application Firewall Configuration on page 564

Copyright © 2016, Juniper Networks, Inc.

Verifying Application Firewall Configuration

Purpose Verify information about application firewall support enabled under the security policy.

Action To verify the security policy configuration enabled with application firewall, enter the **show security policies** and **show security policies detail** commands. To verify all the application firewall rule sets configured on the device, enter the **show security application-firewall rule-set all** command.

Meaning The output displays information about application firewall enabled policies configured on the system. Verify the following information.

- Rule set
- Rules
- Match criteria

Related Documentation

- Application Firewall Overview on page 555
- Understanding Application Identification Techniques on page 485
- *Building Blocks Feature Guide for Security Devices*

Example: Configuring an Application Group for Application Firewall

With application identification, multiple applications can be configured in a dynamic application groups for consistent reuse. AppFW rules permit and deny traffic by specifying application names, dynamic application group names, or both. By using predefined application groups, AppFW rules require no updating when new applications are added to common groups.

NOTE: **The application group is managed by the application identification module.**

This example shows how to configure application groups within the application firewall rule set.

- Requirements on page 565
- Overview on page 565
- Configuration on page 565
- Verification on page 567

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.

Overview

The following example configures network policies to control outbound traffic from the trust zone to the untrust zone. All traffic permitted by the policy is processed further with the specified application firewall. The application firewall denies outbound traffic from unknown applications. Outbound Google-Talk traffic is allowed, but all other known social networking traffic is denied. All other traffic is permitted.

The junos:GOOGLE-TALK application is included in the predefined group junos:social-networking. To allow junos:GOOGLE-TALK traffic and deny the rest of the group, the rule permitting junos:GOOGLE-TALK traffic must come before the rule denying traffic from the rest of the applications in the group.

This configuration example shows how to:

- Configure dynamic application groups in an application firewall.

Configuration

CLI Quick Configuration

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

```
set security application-firewall rule-sets social-network
set rule google-rule match dynamic-application junos:GOOGLE-TALK
set rule google-rule then permit
set rule denied-sites match dynamic-application-groups junos:social-networking
set rule denied-sites match dynamic-application junos:UNKNOWN
set rule denied-sites then deny
set default-rule permit
edit security policies from-zone trust to-zone untrust policy outbound-traffic
set match source-address any
set match destination-address any
set match application junos:HTTP
set then permit application-services application-firewall rule-set social-network
```

Step-by-Step Procedure

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

To configure application firewall rule-sets and security policies for outbound traffic:

1. Create the rule-set social-network.

```
[edit]
user@host# set security application-firewall rule-sets social-network
```

Copyright © 2016, Juniper Networks, Inc.

2. Define a rule to permit Google-Talk traffic.

```
[edit security application-firewall rule-sets social-network]
user@host# set rule google-rule match dynamic-application junos:GOOGLE-TALK
user@host# set rule google-rule then permit
```

3. Define a second rule that denies all other social-networking traffic and traffic from an unknown application.

```
[edit security application-firewall rule-sets social-network]
user@host# set rule denied-sites match dynamic-application-groups
  junos:social-networking
user@host# set rule denied-sites match dynamic-application junos:UNKNOWN
user@host# set rule denied-sites then deny
```

   Note that rule sequence is important. If the rules google-rule and denied-sites are reversed, Google-Talk traffic would never be permitted. The denied-sites rule would shadow google-rule.

4. Define the default-rule that permits all other traffic.

```
[edit security application-firewall rule-sets social-network]
user@host# user@host# set default-rule permit
```

5. Configure the outbound-traffic policy to apply the social-network rule-set to all outbound traffic.

```
[edit security policies from-zone trust to-zone untrust policy outbound-traffic]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application junos:HTTP
user@host# set then permit application-services application-firewall rule-set
  social-network
```

Results From configuration mode, confirm your configuration by entering the **show security application-firewall** and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security application-firewall
...
rule-sets social-network {
  rule google-rule {
    match {
      dynamic-application junos:GOOGLE-TALK;
    }
  }
    then {
      permit ;
    }
  rule denied-sites {
    match {
      dynamic-application-groups junos:social-networking
      dynamic-application junos:UNKNOWN;
    }
    then {
```


Copyright © 2016, Juniper Networks, Inc.

```
        deny ;
      }
    }
    default-rule {
      permit;
    }
  }
...

[edit]
user@host# show security policies
from-zone untrust to-zone trust {
  ...
  policy outbound-traffic {
    match {
      source-address any;
      destination-address any;
      application junos-http;
    }
    then {
      permit {
        application-services {
          application-firewall {
            rule-set social-network
          }
        }
      }
    }
  }
  ...
}
```

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

Verification

Verifying Application Firewall Configuration

Purpose Verify information about application grouping support under the application firewall policy.

Action

- To verify the application firewall policy configuration enabled with application grouping, from the operational mode, enter the **show security policies** and **show security policies detail** commands.
- To verify all the application firewall rule sets configured on the device, from the operational mode, enter the **show security application-firewall rule-set all** command.
- To verify the list of applications defined within the application group, from the operational mode, enter the **show services application-identification application-group** ***application-group-name*** command.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560
- Understanding Application Identification Techniques on page 485
- Security Policies Overview on page 1085

Example: Configuring Application Firewall When SSL Proxy Is Enabled

NOTE: If none of the services (AppFW, IDP, or AppTrack) are configured, then SSL proxy services are bypassed even if an SSL proxy profile is attached to a firewall policy.

This example describes how AppFW supports this AppID functionality when SSL proxy is enabled.

- Requirements on page 568
- Overview on page 568
- Configuration on page 568

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book with addresses for the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.
- Create a SSL proxy profile that enables SSL proxy by means of a policy. See "Configuring SSL Proxy" on page 540.

Overview

This example shows how to verify the functionality of AppFW when SSL proxy is enabled and a different action, deny or permit, is performed on plain text and encrypted traffic.

Configuration

CLI Quick Configuration

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

```
set security policies from-zone Z_1 to-zone Z_2 policy policy1 match source-address any
set security policies from-zone Z_1 to-zone Z_2 policy policy1 match destination-address
    any
set security policies from-zone Z_1 to-zone Z_2 policy policy1 match application junos-https
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone Z_1 to-zone Z_2 policy policy1 then permit
  application-services application-firewall rule-set appfw-rs-1
set security policies from-zone Z_1 to-zone Z_2 policy policy1 then permit
  application-services ssl-proxy profile-name ssl-profile-1
set security policies from-zone Z_1 to-zone Z_2 policy policy2 match source-address any
set security policies from-zone Z_1 to-zone Z_2 policy policy2 match destination-address
  any
set security policies from-zone Z_1 to-zone Z_2 policy policy2 match application junos-http
set security policies from-zone Z_1 to-zone Z_2 policy policy2 then permit
  application-services application-firewall rule-set appfw-rs-2
set security application-firewall rule-sets appfw-rs-1 rule rule1 match dynamic-application
  [junos:ORACLE]
user@host# set security application-firewall rule-sets appfw-rs-1 rule rule1 then permit
user@host# set security application-firewall rule-sets appfw-rs-1 default-rule deny
user@host# set security application-firewall rule-sets appfw-rs-2 rule rule1 match
  dynamic-application [junos:HULU]
user@host# set security application-firewall rule-sets appfw-rs-2 rule rule1 then deny
user@host# set security application-firewall rule-sets appfw-rs-2 default-rule permit
```

Step-by-Step Procedure

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

In this example, you configure two security policies with AppFW rule sets that permit or deny traffic from plain text or encrypted traffic:

- Allow the encrypted version of Oracle and deny any other encrypted traffic.
- Allow all HTTP traffic, except Hulu.

1. Configure a policy to process the traffic with AppFW rule set appfw-rs-1 and SSL proxy profile ssl-profile-1.

   ```
   [edit security policies from-zone Z_1 to-zone Z_2 policy policy1
   user@host# set match source-address any
   user@host# set match destination-address any
   user@host# set match application junos-https
   user@host# set then permit application-services application-firewall rule-set
     appfw-rs-1
   user@host# set then permit application-services ssl-proxy profile-name ssl-profile-1
   ```

2. Configure another policy with rule set appfw-rs-2.

   ```
   [edit security policies from-zone Z_1 to-zone Z_2 policy policy2
   user@host# set match source-address any
   user@host# set match destination-address any
   user@host# set match application junos-http
   user@host# set then permit application-services application-firewall rule-set
     appfw-rs-2
   ```

3. Define the AppFW rule set appfw-rs-1 to permit an encrypted version of Oracle and to deny any other encrypted traffic.

   ```
   [edit security application-firewall rule-sets appfw-rs1]
   user@host# set rule rule1 match dynamic-application [junos:ORACLE]
   user@host# set rule rule1 then permit
   user@host# set default-rule deny
   ```

Copyright © 2016, Juniper Networks, Inc.

4. Define the AppFW rule set appfw-rs-2 to allow all plain text traffic except Hulu.

```
[edit security application-firewall rule-sets appfw-rs2]
user@host# set rule rule1 match dynamic-application [junos:HULU]
user@host# set rule rule1 then deny
user@host# set default-rule permit
```

Results From configuration mode, confirm your configuration by entering the **show security policies** and **show security application-firewall** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

NOTE: For **application junos-https**, SSL proxy detects an SSL session based on the dynamic application identified for that session. If you know any webservers that are running nonstandard ports, you can use a custom Junos OS application to identify the application. However, if the webservers are not known, for example on the Internet, use **application any**. Non-SSL sessions that come across the policy rule are ignored by SSL proxy. A syslog SSL_PROXY_SESSION_IGNORE is sent out for these sessions. Juniper Networks recommends that you use application “any” with caution because this can result in a lot of traffic, incurring initial SSL proxy processing and thereby impacting performance.

Verifying Application Firewall In an SSL Proxy Enabled Policy

Purpose Verify that the application is configured correctly when SSL proxy is enabled in a policy.

Action From operational mode, enter the **show security policies** command.

The following output shows the options for the **show security flow session** command.

```
user@host> show security flow session ?

   Possible completions:
    <[Enter]>            Execute this command
    application          Application protocol name
    application-firewall Show application-firewall sessions
    application-firewall-rule-set  Show application firewall sessions matching
rule-set name
    brief                Show brief output (default)
    destination-port     Destination port (1..65535)
    destination-prefix   Destination IP prefix or address
    dynamic-application  Dynamic application name
    extensive            Show detailed output
+   encrypted            Show encrypted traffic
    family               Show session by family
    idp                  Show idp sessions
    interface            Name of incoming or outgoing interface
    nat                  Show sessions with network address translation
```


Copyright © 2016, Juniper Networks, Inc.

```
protocol              IP protocol number
resource-manager      Show sessions with resource manager
session-identifier    Show session with specified session identifier
source-port           Source port (1..65535)
source-prefix         Source IP prefix or address
summary               Show output summary
tunnel                Show tunnel sessions
 |                     Pipe through a command
```

To display SSL encrypted UNKNOWN sessions, use the **show security flow session application-firewall dynamic-application junos:SSL extensive** command.

To display all HTTPS sessions, use the **show security flow session application-firewall dynamic-application junos:HTTP encrypted extensive** command.

Related Documentation

- SSL Proxy Overview on page 531
- Application Firewall, IDP, and Application Tracking with SSL Proxy Overview on page 540
- Understanding Security Policy Elements on page 1091
- Security Policies Configuration Overview on page 1093
- Application Firewall Overview on page 555
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 30

Configuring Application Tracking

- Understanding AppTrack on page 573
- Example: Configuring AppTrack on page 574
- Example: Configuring AppTrack When SSL Proxy Is Enabled on page 580
- Disabling AppTrack on page 581

Understanding AppTrack

AppTrack, an application tracking tool, provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack messages are similar to session logs and use syslog or structured syslog formats. The message also includes an application field for the session. If AppTrack identifies a custom-defined application and returns an appropriate name, the custom application name is included in the log message. (If the application identification process fails or has not yet completed when an update message is triggered, the message specifies **none** in the application field.)

AppTrack supports both IPv4 and IPv6 addressing. Related messages display addresses in the appropriate IPv4 or IPv6 format.

User identity details such as user name and user role have been added to the AppTrack session create, session close, and volume update logs. These fields will contain the user name and role associated with the policy match. The logging of user name and roles is enabled only for security policies that provide UAC enforcement. For security policies without UAC enforcement, the user name and user role fields are displayed as N/A. The user name is displayed as unauthenticated user and user role is displayed as N/A, if the device cannot retrieve information for that session because there is no authentication table entry for that session or because logging of this information is disabled. The user role field in the log contains the list of all the roles performed by the user if match criteria is specific, authenticated user, or any, and the user name field in the log contains the correct user name. The user role field in the log will contain N/A if the match criteria and the user name field in the log contain unauthenticated user or unknown user.

Copyright © 2016, Juniper Networks, Inc.

If you enable AppTrack for a zone and specify a **session-update-interval** time, whenever a packet is received, AppTrack checks whether the time since the start of the session or since the last update is greater than the update interval. If so, AppTrack updates the counts and sends an update message to the host. If a short-lived session starts and ends within the update interval, AppTrack generates a message only at session close.

When you want the initial update message to be sent earlier than the specified update interval, use the **first-update-interval**. The **first-update-interval** lets you enter a shorter interval for the first update only. Alternatively, you can generate the initial update message at session start by using the **first-update** option.

The close message updates the statistics for the last time and provides an explanation for the session closure. The following codes are used:

TCP RST—RST received from either end.

TCP FIN—FIN received from either end.

Response received—Response received for a packet request (such as **icmp req-reply**).

ICMP error—ICMP error received (such as **dest unreachable**).

Aged out—Session aged out.

ALG—ALG closed the session.

IDP—IDP closed the session.

Parent closed—Parent session closed.

CLI—Session cleared by a CLI statement.

Policy delete—Policy marked for deletion.

Related Documentation

- Example: Configuring AppTrack on page 574
- Disabling AppTrack on page 581
- Understanding Application Identification Techniques on page 485

Example: Configuring AppTrack

This example shows how to configure the AppTrack tracking tool so you can analyze the bandwidth usage of your network.

- Requirements on page 575
- Overview on page 575
- Configuration on page 575
- Verification on page 577

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you configure AppTrack, it is important that you understand conceptual information about AppTrack and Junos OS application identification. See "Understanding AppTrack" on page 573 and "Understanding the Junos OS Application Identification Database" on page 488.

Overview

Application identification is enabled by default and is automatically turned on when you configure the AppTrack, AppFW, or IDP service. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility. STRM includes the support for AppTrack Reporting and includes several predefined search templates and reports.

Configuration

This example shows how to enable application tracking for the security zone named trust. The first log message is to be generated when the session starts, and update messages should be sent every 4 minutes after that. A final message should be sent at session end.

The example also shows how to add the remote syslog device configuration to receive AppTrack log messages in sd-syslog format. The source IP address that is used when exporting security logs is 5.0.0.254, and the security logs are sent to the host located at address 5.0.0.1.

NOTE: On all SRX Series devices, J-Web pages for AppSecure Services are preliminary. We recommend using CLI for configuration of AppSecure features.

CLI Quick Configuration

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

NOTE: Changing the session-update-interval and the first-update-interval is not necessary in most situations. The commands are included in this example to demonstrate their use.

```
user@host# set security log mode stream
user@host# set security log format sd-syslog
user@host# set security log source-address 5.0.0.254
user@host# set security log stream app-track-logs host 5.0.0.1
user@host#set security zones security-zone trust application-tracking
user@host#set security application-tracking session-update-interval 4
user@host#set security application-tracking first-update
```

Copyright © 2016, Juniper Networks, Inc.

NOTE: On SRX5600, and SRX5800 devices, if the syslog configuration does not specify a destination port, the default destination port will be the syslog port. If you specify a destination port in the syslog configuration, then that port will be used instead.

Step-by-Step Procedure

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

To configure AppTrack:

1. Add the remote syslog device configuration to receive Apptrack messages in sd-syslog format.

   ```
   [edit]
   user@host# set security log mode stream
   user@host# set security log format sd-syslog
   user@host# set security log source-address 5.0.0.254
   user@host# set security log stream app-track-logs host 5.0.0.1
   ```

2. Enable AppTrack for the security zone trust.

   ```
   [edit]
   user@host# set security zones security-zone trust application-tracking
   ```

3. (Optional) For this example, generate update messages every 4 minutes.

   ```
   [edit]
   user@host# set security application-tracking session-update-interval 4
   ```

   The default interval between messages is 5 minutes. If a session starts and ends within this update interval, AppTrack generates one message at session close. However, if the session is long-lived, an update message is sent every 5 minutes. The **session-update-interval *minutes*** is configurable as shown in this step.

4. (Optional) For this example, generate the first message when the session starts.

   ```
   [edit]
   user@host# set security application-tracking first-update
   ```

   By default, the first message is generated after the first session update interval elapses. To generate the first message at a different time than this, use the **first-update** option (generate the first message at session start) or the **first-update-interval *minutes*** option (generate the first message after the specified minutes). For example, enter the following command to generate the first message one minute after session start.

   ```
   [edit]
   user@host# set security application-tracking first-update-interval 1
   ```


Copyright © 2016, Juniper Networks, Inc.

NOTE: The **first-update** option and the **first-update-interval** *minutes* option are mutually exclusive. If you specify both, the **first-update-interval** value is ignored.

Once the first message has been generated, an update message is generated each time the session update interval is reached.

Results From configuration mode, confirm your configuration by entering the **show security** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
user@host# show security

...
application-tracking {
   first-update;
   session-update-interval 4;
}
log {
   mode stream;
   format sd-syslog;
   source-address 5.0.0.254;
   stream app-track-logs {
      host {
         5.0.0.1;
      }
   }
}
...

[edit]
user@host# show security zones
...
security-zone trust {
   ...
   application-tracking;
}
```

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

Verification

Use the STRM product on the remote logging device to view the AppTrack log messages.

Copyright © 2016, Juniper Networks, Inc.

To confirm that the configuration is working properly, you can also perform these tasks on the SRX Series device:

- Reviewing AppTrack Statistics on page 578
- Verifying AppTrack Counter Values on page 578
- Verifying Security Flow Session Statistics on page 578
- Verifying Application System Cache Statistics on page 579
- Verifying the Status of Application Identification Counter Values on page 579

Reviewing AppTrack Statistics

Purpose Review AppTrack statistics to view characteristics of the traffic being tracked.

Action From operational mode, enter the **show services application-identification statistics applications** command.

```
user@host> show services application-identification statistics applications
Last Reset: 2012-02-14 21:23:45 UTC

Application    Sessions              Bytes    Encrypted
HTTP                  1               2291          Yes
HTTP                  1                942           No
SSL                   1               2291          Yes
unknown               1                100           No
unknown               1                100          Yes
```

NOTE: For more information on the **show services application-identification statistics applications** command, see show services application-identification statistics applications.

Verifying AppTrack Counter Values

Purpose View the AppTrack counters periodically to monitor logging activity.

Action From operational mode, enter the **show security application-tracking counters** command.

```
user@host> show security application-tracking counters
AVT counters:                  Value
 Session create messages           1
 Session close messages            1
 Session volume updates            0
 Failed messages                   0
```

Verifying Security Flow Session Statistics

Purpose Compare byte and packet counts in logged messages with the session statistics from the **show security flow session** command output.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security flow session** command.

```
user@host> show security flow session
Flow Sessions on FPC6 PIC0:

Session ID: 120000044, Policy name: policy-in-out/4, Timeout: 1796, Valid
In: 4.0.0.1/39075 --> 5.0.0.1/21;tcp, If: ge-0/0/0.0, Pkts: 22, Bytes: 1032
Out: 5.0.0.1/21 --> 4.0.0.1/39075;tcp, If: ge-0/0/1.0, Pkts: 24, Bytes: 1442

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

Byte and packet totals in the session statistics should approximate the counts logged by AppTrack but might not be exactly the same. AppTrack counts only incoming bytes and packets. System-generated packets are not included in the total, and dropped packets are not deducted.

Verifying Application System Cache Statistics

Purpose Compare cache statistics such as IP address, port, protocol, and service for an application from the **show services application-identification application-system-cache** command output.

Action From operational mode, enter the **show services application-identification application-system-cache** command.

Verifying the Status of Application Identification Counter Values

Purpose Compare session statistics for application identification counter values from the **show services application-identification counter** command output.

Action From operational mode, enter the **show services application-identification counter** command.

Related Documentation

- Understanding AppTrack on page 573
- Disabling AppTrack on page 581
- Understanding Application Identification Techniques on page 485

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring AppTrack When SSL Proxy Is Enabled

This example describes how AppTrack supports AppID functionality when SSL proxy is enabled.

- Requirements on page 580
- Overview on page 580
- Configuration on page 580

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Create an SSL proxy profile that enables SSL proxy by means of a policy. See "Configuring SSL Proxy" on page 540.

Overview

You can configure AppTrack either in the to or from zones. This example shows how to configure AppTrack in a to zone in a policy rule when SSL proxy is enabled.

Configuration

CLI Quick Configuration

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

```
set security zones security-zone Z_1 application-tracking
set security policies from-zone Z_1 to-zone Z_2 policy policy1 match source-address any
set security policies from-zone Z_1 to-zone Z_2 policy policy1 match destination-address
  any
set security policies from-zone Z_1 to-zone Z_2 policy policy1 then permit
  application-services ssl-proxy profile-name ssl-profile-1
set security policies from-zone Z_1 to-zone Z_2 policy policy1 then permit
```

Step-by-Step Procedure

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

In this example, you configure application tracking and permit application services in an SSL proxy profile configuration.

1. Configure application tracking in a to-zone (you can also configure using a from-zone).

   ```
   [edit security policies]
   user@host# set security zones security-zone Z_1 application-tracking
   ```

2. Configure SSL proxy profile.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies from-zone Z_1 to-zone Z_2 policy policy1]
set match source-address any
set match destination-address any
set match application junos-https
set then permit application-services ssl-proxy profile-name ssl-profile-1
set then permit
```

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

```
from-zone Z_1 to-zone Z_2 {
    policy policy1 {
        match {
            source-address any;
            destination-address any;
        }
        then {
            permit {
                application-services {
                    ssl-proxy {
                        profile-name ssl-profile-1;
                    }
                }
            }
        }
    }
}
```

NOTE: Verify that the configuration is working properly. Verification in AppTrack works similarly to verification in AppFW. See the verification section of "Example: Configuring Application Firewall When SSL Proxy Is Enabled" on page 568.

Related Documentation

- SSL Proxy Overview on page 531
- Application Firewall, IDP, and Application Tracking with SSL Proxy Overview on page 540
- Understanding Security Policy Elements on page 1091
- Security Policies Configuration Overview on page 1093
- Example: Configuring AppTrack on page 574

Disabling AppTrack

Application tracking is enabled by default. You can disable application tracking without deleting the zone configuration.

To disable application tracking:

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security application-tracking disable
```

If application tracking has been previously disabled and you want to reenable it, delete the configuration statement that specifies disabling of application tracking:

```
user@host# delete security application-tracking disable
```

If you are finished configuring the device, commit the configuration.

To verify the configuration, enter the **show security application-tracking** command.

Related Documentation

- Understanding AppTrack on page 573
- Example: Configuring AppTrack on page 574
- Understanding Application Identification Techniques on page 485

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 31

Configuring Application QoS

- Understanding Application QoS (AppQoS) on page 583
- Example: Configuring AppQoS on page 589

Understanding Application QoS (AppQoS)

The application quality of service (AppQoS) feature expands the capability of Junos OS class of service (CoS) to include marking DSCP values based on Layer-7 application types, honoring application-based traffic through loss priority settings, and controlling transfer rates on egress PICs based on Layer-7 application types.

There are four ways to mark DSCP values on SRX Series devices:

- IDP attack action-based DSCP rewriters
- Layer 7 application-based DSCP rewriters
- ALG-based DSCP rewriters
- Firewall filter-based DSCP rewriters

IDP remarking is conducted at the ingress port based on IDP rules. Application remarking is conducted at the egress port based on application rules. Interface-based remarking also occurs at the egress port based on firewall filter rules. (See the *Class of Service Feature Guide for Security Devices* for a detailed description of Junos OS CoS features.)

The remarking decisions of these three rewriters can be different. If a packet triggers all three, the method that takes precedence is based on how deep into the packet content the match is conducted. IDP remarking has precedence over application remarking which has precedence over interface-based remarking.

If a packet triggers both AppQoS and ALG-based DSCP rewriters, then AppQoS takes precedence over ALG-based DSCP rewriters.

The AppQoS DSCP rewriter conveys a packet's quality of service through both the forwarding class and a loss priority. The AppQoS rate-limiting parameters control the transmission speed and volume for its associated queues.

- Unique Forwarding Classes and Queue Assignments on page 584
- Application-Aware DSCP Code-Point and Loss Priority Settings on page 584

Copyright © 2016, Juniper Networks, Inc.

- Rate Limiters and Profiles on page 586
- Rate-Limiter Assignment on page 587
- Rate-Limiter Action on page 589
- AppQoS Security Policy Configuration on page 589

Unique Forwarding Classes and Queue Assignments

The forwarding class provides three functions:

- Groups packets with like characteristics
- Assigns output queues
- Resolves conflicts with existing Junos OS firewall filter-based rewriters

Unique forwarding class names protect AppQoS remarking from being overwritten by interface-based rewrite rules. A firewall filter-based rewriter remarks a packet's DSCP value if the packet's forwarding class matches a class defined specifically for this rewriter. If the packet's forwarding class does not match any of the firewall filter-based rewriter's classes, the DSCP value is not remarked. To protect AppQoS values from being overwritten, therefore, use forwarding class names that are unknown to the firewall filter-based rewriter.

Each forwarding class is assigned to an egress queue that provides the appropriate degree of enhanced or standard processing. Many forwarding classes can be assigned to a single queue. Therefore, any queues defined for the device can be used by IDP, AppQoS, and firewall filter-based rewriters. It is the forwarding class name, not the queue, that distinguishes the transmission priority. (See the *Class of Service Feature Guide for Security Devices* for information about configuring queues and schedulers.)

For SRX5400, SRX5600, and SRX5800 devices, the AppQoS forwarding class names and queue assignments are defined with the **class-of-service** CLI configuration command:

```
[edit class-of-service]
user@host# forwarding-classes class forwarding-class-name queue-num queue-number
```

For SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 devices, the AppQoS forwarding class names and queue assignments are defined with the **class-of-service** CLI configuration command:

```
[edit class-of-service]
user@host# forwarding-classes queue queue-number forwarding-class-name
```

Application-Aware DSCP Code-Point and Loss Priority Settings

For AppQoS, traffic is grouped based on rules that associate a defined forwarding class with selected applications. The match criteria for the rule includes one or more applications. When traffic from a matching application encounters the rule, the rule action sets the forwarding class, and remarks the DSCP value and loss priority to values appropriate for the application.

Copyright © 2016, Juniper Networks, Inc.

A Differentiated Services (DiffServ) code point (DSCP) value is specified in the rule either by a 6-bit bitmap value or by a user-defined or default alias. Table 40 provides a list of Junos OS default DSCP alias names and bitmap values.

Table 40: Standard CoS Aliases and Bit Values

CoS Value Type	Alias	Bit Value
Expedited forwarding	ef	101110
Assured forwarding	af11	001010
Assured forwarding	af12	001100
Assured forwarding	af13	001110
Assured forwarding	af21	010010
Assured forwarding	af22	010100
Assured forwarding	af23	010110
Assured forwarding	af31	011010
Assured forwarding	af32	011100
Assured forwarding	af33	011110
Assured forwarding	af41	100010
Assured forwarding	af42	100100
Assured forwarding	af43	100110
Best effort	be	000000
	cs1	001000
	cs2	010000
	cs3	011000
	cs4	100000
	cs5	101000
Network control	nc1/cs6	110000
Network control	nc2/cs7	111000

Copyright © 2016, Juniper Networks, Inc.

The queue's scheduler uses the loss priority to control packet discard during periods of congestion by associating drop profiles with particular loss priority values. (See the *Class of Service Feature Guide for Security Devices* for information about configuring queues and schedulers.)

The rule applies a loss priority to the traffic groups. A high loss priority means a high probability that the packet could be dropped during a period of congestion. Four levels of loss priority are available:

- **high**
- **medium-high**
- **medium-low**
- **low**

The rule set is defined in the **class-of-service application-traffic-control** configuration command:

```
[edit class-of-service]
user@host# application-traffic-control rule-sets ruleset-name rule rule-name1 match
   application application-name application-name ...
user@host# application-traffic-control rule-sets ruleset-name rule rule-name1 match
   application-group application-group-name application-group-name ...
user@host# application-traffic-control rule-sets ruleset-name rule rule-name1 then
   forwarding-class fc-name
user@host# application-traffic-control rule-sets ruleset-name rule rule-name1 then
   dscp-code-point bitmap
user@host# application-traffic-control rule-sets ruleset-name rule rule-name1 then
   loss-priority loss-pri-value
```

Rate Limiters and Profiles

When congestion occurs, AppQoS implements rate limiting on all egress PICs on the device. If packets exceed the assigned limitations, they are dropped. *Rate limiters* maintain a consistent level of throughput and packet loss sensitivity for different classes of traffic. All egress PICs employ the same rate-limiting scheme.

The total bandwidth of a PIC is about 10 Gbps. Rate-limiter hardware for the PIC can provision up to 2 Gbps. Therefore, the upper bandwidth limit for rate limiting is 2^{31} bps.

A rate-limiter profile defines the limitations. It is a unique combination of **bandwidth-limit** and **burst-size-limit** specifications. The **bandwidth-limit** defines the maximum number of kilobits per second that can traverse the port. The **burst-size-limit** defines the maximum number of bytes that can traverse the port in a single burst. The **burst-size-limit** reduces starvation of lower priority traffic by ensuring a finite size for each burst.

Copyright © 2016, Juniper Networks, Inc.

AppQoS allows up to 16 profiles and up to 1000 rate limiters per device. Multiple rate limiters can use the same profile. In the following example, five rate limiters are defined using two profiles:

Rate Limiter Name	Profile	
	bandwidth-limit	burst-size-limit
limiter-1	200	26000
limiter-2	200	26000
limiter-3	200	26000
limiter-4	400	52000
limiter-5	400	52000

Rate limiters are defined with the **class-of-service application-traffic-control** configuration command.

```
[edit class-of-service]
user@host# application-traffic-control rate-limiters rate-limiter-name bandwidth-limit
   value-in-Kbps burst-rate-limit value-in-bytes
```

Rate-Limiter Assignment

Rate limiters are applied in rules based on the application of the traffic. Two rate limiters are applied for each session: **client-to-server** and **server-to-client**. This usage allows traffic in each direction to be provisioned separately.

Different AppQoS rules within the same rule set can share a rate limiter. In this case, the applications of those rules share the same bandwidth. There are no limitations on the number of rules in one rule set that can assign the same rate limiter.

The following examples show how the rate limiters defined in the preceding section could be assigned. For instance, a rule set could reuse a rate limiter in several rules and in one or both flow directions:

- rule-set-1
  - rule-1A
    - client-to-server limiter-1
    - server-to-client limiter-1
  - rule-1B
    - client-to-server limiter-1

Copyright © 2016, Juniper Networks, Inc.

    - server-to-client limiter-1

If the same profiles are needed in several rule sets, a sufficient number of rate limiters needs to be defined specifying the same **bandwidth-limit** and **burst-size-limit**. The two rule sets in the following example implement the same profiles by assigning different, but comparable, rate limiters.

- rule-set-2
  - rule-2A
    - client-to-server limiter-2
    - server-to-client limiter-2
  - rule-2B
    - client-to-server limiter-2
    - server-to-client limiter-4
- rule-set-3
  - rule-3A
    - client-to-server limiter-3
    - server-to-client limiter-3
  - rule-3B
    - client-to-server limiter-3
    - server-to-client limiter-5

A rate limiter is applied using the **class-of-service application-traffic-control rule-sets** command in the same way that a forwarding class, DSCP value, and loss priority are set.

```
[edit class-of-service]
user@host# application-traffic-control rule-sets rule-set-name rule rule-name1 then
  rate-limit client-to-server rate-limiter1 server-to-client rate-limiter2
```

If AppQoS and firewall filter-based rate limiting are both implemented on the egress PIC, both are taken into consideration. AppQoS rate limiting is considered first. Firewall filter-based rate limiting occurs after that.

NOTE: If packets are dropped from a PIC, the SRX Series device does not send notifications to the client or the server. The upper-level applications on the client and the server devices are responsible for retransmission and error handling.

Copyright © 2016, Juniper Networks, Inc.

Rate-Limiter Action

Based on the type of SRX Series device, AppQoS rules can be configured with different rate-limiter actions:

- Discard
  - When this option is selected, the out-of-profile packets are just dropped.
  - This is the default action type and need not be configured.
  - This option is supported on all SRX Series devices.
- Loss-priority-high
  - When this option is selected , it elevates the loss priority to maximum. In other words, it is a delayed drop; that is, the discard decision is taken at the egress output queue level. If there is no congestion, it allows the traffic even with maximum loss priority. But if congestion occurs, it drop these maximum loss priority packets first.
  - This option must be configured within the AppQoS rule (to override the default action) using the following command:

    ```
    [edit]
    user@host# set class-of-service application-traffic-control rule-sets rset-01 rule r1
      then rate-limit loss-priority-high
    ```

  - This option is supported only on branch SRX Series devices.

AppQoS Security Policy Configuration

The AppQoS rule set can be implemented in an existing policy or a specific application policy.

```
[edit]
user@host# security policies from-zone zone-name to-zone zone-name
[edit security policies from-zone zone-name to-zone zone-name]
user@host# policy policy-name match source-address IP-address
user@host# policy policy-name match destination-address IP-address
user@host# policy policy-name match application application-name application-name
user@host# policy policy-name then permit application-services application-traffic-control
  rule-set app-rule-set-name
```

Related Documentation

- Example: Configuring AppQoS on page 589
- Understanding Application Identification Techniques on page 485

Example: Configuring AppQoS

This example shows how to enable AppQoS prioritization and rate limiting within a policy.

- Requirements on page 590
- Overview on page 590

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 590
- Verification on page 593

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, AppQoS is implemented so that FTP applications are restricted to a level below the specified throughput while other applications are transmitted at a more conventional speed and loss priority level.

NOTE: On all SRX Series devices, J-Web pages for AppSecure Services are preliminary. We recommend using CLI for configuration of AppSecure features.

Configuration

Step-by-Step Procedure

To configure an AppQoS implementation:

1. Define one or more forwarding classes dedicated to AppQoS marking. In this example, a single forwarding class, my-app-fc, is defined and assigned to queue 0.

   For SRX1500, SRX5400, SRX5600, and SRX5800 devices, use the following command:

   ```
   [edit]
   user@host# set class-of-service forwarding-classes class my-app-fc queue-num
     0
   ```

   For SRX300, SRX320, SRX340, SRX345, and SRX550 devices, use the following command:

   ```
   [edit]
   user@host# set class-of-service forwarding-classes queue-num 0 my-app-fc
   ```

2. Define rate limiters. In this example, two rate limiters are defined.

   NOTE: For high-end SRX Series devices, you can define up to 1000 rate limiters for a device, but only 16 profiles (unique bandwidth-limit and burst-size-limit combinations).

   - test-r1 with a bandwidth of 100 Kbps and a burst limit of 13,000 bytes
   - test-r2 with a bandwidth of 200 Kbps and a burst limit of 26,000 bytes

   ```
   [edit]
   user@host# set class-of-service application-traffic-control rate-limiters test-rl
     bandwidth-limit 100
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set class-of-service application-traffic-control rate-limiters test-rl
  burst-size-limit 13000
user@host# set class-of-service application-traffic-control rate-limiters test-r2
  bandwidth-limit 200
user@host# set class-of-service application-traffic-control rate-limiters test-r2
  burst-size-limit 26000
```

3. Define AppQos rules and application match criteria. For this example, rule 0 in rule set ftp-test1 is applied to junos:FTP packets.

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 match application junos:FTP
```

4. Define the action for rule 0 when it encounters a junos:FTP packet. In this example, when a match is made, the packet is marked with the forwarding class my-app-fc, the DSCP value of af22, and a loss priority of low.

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then forwarding-class my-app-fc
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then dscp-code-point af22
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then loss-priority low
```

5. Assign rate limiters for rule 0 to traffic in each direction. In this case, the rate limiter test-r1 is set in both directions.

NOTE: Rate limiter test-r1 can be assigned to one or both traffic directions in rule 0. It could also be assigned in other rules within rule set ftp-test1. However, once test-r1 is assigned to rule set ftp-test1, it cannot be assigned in any other rule set.

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then rate-limit client-to-server test-r1
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then rate-limit server-to-client test-r1
```

6. Log the AppQoS event whenever this action is triggered:

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  0 then log
```

7. Define other rules to handle application packets that did not match the previous rule. In this example, a second and final rule applies to all remaining applications.

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  1 match application-any
```

8. Assign rate limiters for the second rule. In this example, any traffic that is not from FTP is assigned rate limiter test-r2 in both directions.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  1 then rate-limit client-to-server test-r2
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  1 then rate-limit server-to-client test-r2
user@host# set class-of-service application-traffic-control rule-sets ftp-test1 rule
  1 then log
```

9. Add the AppQoS implementation to a policy. In this example, policy p1 applies the rule set ftp-test1 to all traffic from the trust zone to the untrust zone.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy p1
[edit security policies from-zone trust to-zone untrust policy p1]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set then permit application-services application-traffic-control rule-set
  ftp-test1
```

Results From configuration mode, confirm your policy configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

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

```
...
policy p1 {
  match {
    source-address any;
    destination-address any;
    application any;
  }
  then {
    permit {
      application-services {
        application-traffic-control {
          rule-set ftp-test1
        }
      }
    }
  }
}
...
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Flow Session Configuration on page 593
- Verifying Session Statistics on page 593
- Verifying Rate-Limiter Statistics on page 594
- Verifying Rule Statistics on page 595

Verifying Flow Session Configuration

Purpose Verify that AppQoS is enabled.

Action From operational mode, enter the **show security flow session application-traffic-control extensive** command.

```
user@host> show security flow session application-traffic-control extensive
  Session ID: 3729, Status: Normal, State: Active
  Flag: 0x40
  Policy name: p1
  Source NAT pool: Null
  Dynamic application: junos:FTP
  Application traffic control rule-set: ftp-test1, Rule: rule0
  Maximum timeout: 300, Current timeout: 276
  Session State: Valid
  Start time: 18292, Duration: 603536
     In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
      Interface: reth1.0,
      Session token: 0x1c0, Flag: 0x0x21
      Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
      Port sequence: 0, FIN sequence: 0,
      FIN state: 0,
      Pkts: 21043, Bytes: 1136322
     Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
      Interface: .local..0,
      Session token: 0x80, Flag: 0x0x30
      Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
      Port sequence: 0, FIN sequence: 0,
      FIN state: 0,
      Pkts: 0, Bytes: 0
```

Meaning The entry for application traffic control identifies the rule set and rule of the current session.

Verifying Session Statistics

Purpose Verify that AppQoS session statistics are being accumulated at each egress node.

Action From operational mode, enter the **show class-of-service application-traffic-control counter** command.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The PIC number is always displayed as 0 for branch SRX Series devices.

```
user@host> show class-of-service application-traffic-control counter
pic: 2/1
  Counter type                                          Value
  Sessions processed                                    300
  Sessions marked                                       200
  Sessions honored                                      0
  Sessions rate limited                                 100
  Client-to-server flows rate limited                   100
  Server-to-client flows rate limited                   100

pic: 2/0
  Counter type                                          Value
  Sessions processed                                    400
  Sessions marked                                       300
  Sessions honored                                      0
  Sessions rate limited                                 200
  Client-to-server flows rate limited                   200
  Server-to-client flows rate limited                   200
```

Meaning The AppQoS statistics are maintained only if application-traffic-control service is enabled. The number of sessions processed, marked, and honored show that sessions are being directed based on configured AppQoS features. The rate-limiting statistics count the number of directional session flows that have been rate limited.

Verifying Rate-Limiter Statistics

Purpose Verify that bandwidth is being limited as expected when the FTP application is encountered.

Action From operational mode, enter the **show class-of-service application-traffic-control statistics rate-limiter** command.

NOTE: The PIC number is always displayed as 0 for branch SRX Series devices.

```
user@host> show class-of-service application-traffic-control statistics
rate-limiter
pic: 2/1
  Ruleset    Application  Client-to-server Rate(kbps)  Server-to-client Rate(kbps)

  ftp-test1     HTTP      test-r2          200         test-r2          200
  ftp-test1     HTTP      test-r2          200         test-r2          200
  ftp-test1     FTP       test-r1          100         test-r1          100
```


Copyright © 2016, Juniper Networks, Inc.

Meaning Real-time application bandwidth-limit information for each PIC is displayed by rule set. This command provides an indication of the applications being rate limited and the profile being applied.

Verifying Rule Statistics

Purpose Verify that the rule matches the rule statistics.

Action From operational mode, enter the **show class-of-service application-traffic-control statistics rule** command.

NOTE: The PIC number is always displayed as 0 for branch SRX Series devices.

```
user@host>show class-of-service application-traffic-control statistics rule
pic: 2/1
  Ruleset          Rule           Hits
  ftp-test1        0              100
  ftp-test1        1              200
  ...

pic: 2/0
  Ruleset          Rule           Hits
  ftp-test1        0              100
  ftp-test1        1              200
```

Meaning This command provides information on the number of (session) hits for a rule under each rule set.

Related Documentation

- Understanding Application Identification Techniques on page 485

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 32

Configuration Statements

- Class-of-Service Configuration Statement Hierarchy on page 599
- Security Configuration Statement Hierarchy on page 603
- Services Configuration Statement Hierarchy on page 605
- System Configuration Statement Hierarchy on page 612
- [edit security address-book] Hierarchy Level on page 643
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- [edit security log] Hierarchy Level on page 645
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 647
- [edit security zones] Hierarchy Level on page 651
- actions (Services SSL Proxy) on page 654
- actions (Services SSL Initiation) on page 655
- address-mapping (Application Identification) on page 656
- appfw-profile (System) on page 657
- appfw-rule on page 658
- appfw-rule-set on page 659
- application-firewall on page 660
- application (Application Identification) on page 662
- application-firewall (Application Services) on page 664
- application-identification on page 665
- application-group (Services) on page 667
- application-services (Security Policies) on page 668
- application-system-cache on page 669
- application-system-cache-timeout (Services) on page 669
- application-tracking on page 670
- application-tracking (Security Zones) on page 670
- application-traffic-control on page 671

Copyright © 2016, Juniper Networks, Inc.

- application-traffic-control (Application Services) on page 672
- block-message (Application Firewall) on page 673
- context (Application Identification) on page 675
- custom-ciphers on page 676
- default-rule on page 677
- direction (Application Identification) on page 678
- disable (Application Tracking) on page 678
- download (Services) on page 679
- dynamic-application on page 680
- dynamic-application-group on page 680
- enable-flow-tracing (Services) on page 681
- enable-performance-mode on page 681
- enable-session-cache on page 682
- file (Services) on page 682
- files (Services) on page 683
- file (System Logging) on page 684
- first-update on page 686
- first-update-interval on page 687
- flag (Services) on page 688
- format (Security Log) on page 688
- forwarding-classes (CoS) on page 689
- global-config (Services) on page 690
- icmp-mapping (Application Identification) on page 691
- ip (Application Identification) on page 691
- ip-protocol-mapping (Application Identification) on page 692
- initiation (Services) on page 693
- level (Services) on page 694
- log (Security) on page 695
- log (Services) on page 697
- match (Services) on page 698
- no-application-identification (Services) on page 698
- no-application-system-cache (Services) on page 699
- no-remote-trace (Services) on page 699
- over (Application Identification) on page 700
- policies on page 702
- policy (Security Policies) on page 707
- port-range (Application Identification) on page 709

Copyright © 2016, Juniper Networks, Inc.

- preferred-ciphers on page 710
- profile (Application Firewall) on page 710
- profile (Rule Sets) on page 711
- profile (Services) on page 712
- profile (SSL Initiation) on page 713
- profile (SSL Termination) on page 714
- protocol-version on page 715
- proxy (Services) on page 716
- rate-limiters on page 717
- renegotiation (Services) on page 718
- root-ca (Services) on page 718
- rule-sets (CoS AppQoS) on page 719
- rule-sets (Security Application Firewall) on page 721
- security-zone on page 722
- server-certificate (Services) on page 723
- session-update-interval on page 724
- size (Services) on page 724
- ssl (Services) on page 725
- ssl-encryption on page 727
- ssl-proxy (Application Services) on page 727
- statistics (Services) on page 728
- stream (Security Log) on page 729
- termination (Services) on page 730
- then (Security Application Firewall) on page 731
- trusted-ca (Services) on page 732
- traceoptions (Security Application Firewall) on page 733
- traceoptions (Services SSL) on page 735
- traceoptions (Services Application Identification) on page 736
- transport (Security Log) on page 738
- whitelist (Services) on page 739
- zones on page 740

Class-of-Service Configuration Statement Hierarchy

Use the statements in the **class-of-service** configuration hierarchy to configure class-of-services (CoS) features.

```
class-of-service {
  adaptive-shapers adaptive-shaper-name {
    trigger becn {
```

Copyright © 2016, Juniper Networks, Inc.

```
      shaping-rate (absolute-rate | percent percent);
    }
  }
  application-traffic-control {
    rate-limiters rate-limiter-name {
      bandwidth-limit kbps;
      burst-size-limit bytes;
    }
    rule-sets rule-set-name {
      rule rule-name {
          match {
            application [application-name];
            application-any;
            application-group [application-group-name];
            application-known;
            application-unknown;
          }
          then {
            dscp-code-point dscp-value;
            forwarding-class class-name;
            log;
            loss-priority (high | low |medium-high | medium-low);
            rate-limit {
              loss-priority-high;
              client-to-server rate-limiter;
              server-to-client rate-limiter;
            }
          }
        }
      }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier name
      {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-points [alias-or-bit-string ];
        }
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 |exp |ieee-802.1 |ieee-802.1ad |inet-precedence) alias-name{
```


Copyright © 2016, Juniper Networks, Inc.

```
      dscp-bits;
      }
}
drop-profiles profile-name {
  fill-level percent {
    drop-probability number;
  }
  interpolate {
    drop-probability [number];
    fill-level [percent];
  }
}
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map next-hop-map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [lsp-regular-expression];
      next-hop [next-hop-identifier];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps fragmentation-map-name {
  forwarding-class forwarding-class-name {
    drop-timeout milliseconds;
      (fragment-threshold bytes |no-fragmentation) ;
      multilink-class number;
    }
  }
host-outbound-traffic {
  dscp-code-point static-dscp-code-point;
  forwarding-class class-name;
  tcp {
    raise-internet-control-priority;
  }
}
interfaces interface-name {
  input-traffic-control-profile profile-name;
    output-traffic-control-profile profile-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
      output-traffic-control-profile-remaining profile-name;
      scheduler-map scheduler-map;
      shaping-rate bps;
      unit  logical-unit-number  {
        adaptive-shaper adaptive-shaper-name;
        classifiers {
          (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
        }
        forwarding-class class-name;
        input-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        loss-priority-maps {
          frame-relay-de {
            (lpmap-name | default);
          }
        }
        output-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        rewrite-rules {
          (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad
            |inet-precedence)
        }
        scheduler-map scheduler-map-name;
          shaping-rate {
          rate;
        }
        vc-shared-scheduler;
        virtual-channel-group group-name;
      }
    }
  }
  loss-priority-maps {
    frame-relay-de loss-priority-map-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-points [bit-string];
      }
    }
  }
  rewrite-rules {
    (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      rewrite-rule-name {
      forwarding-class forwarding-class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-point alias-or-bit-string;
        }
        import (default | rewrite-rule-name);
      }
    }
  }
  scheduler-maps scheduler-map-name {
    forwarding-class class-name {
      scheduler scheduler-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    schedulers scheduler-name {
      buffer-size {
        exact;
        (percent percent | remainder percent  | temporal microseconds) ;
      }
      drop-profile-map {
        loss-priority (any | high | low | medium-high | medium-low);
        protocol any ;
        drop-profile profile;
      }
      priority (high | low | medium-high | medium-low | strict-high);
      shaping-rate (absolute-rate | percent percent);
      transmit-rate <exact> (percent percent | rate bits | remainder percent);
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
    traffic-control-profiles profile-name {
      delay-buffer-rate ( absolute-rate | cps cells-per-second | percent percent);
        guaranteed-rate (absolute-rate | percent percent);
        overhead-accounting (bytes bytes | cell-mode | frame-mode);
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    tri-color;
    virtual-channel-groups virtual-channel-group-name {
      virtual-channel-name {
        default;
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    }
    virtual-channels virtual-channel-name;
  }
```

Related Documentation

- SSL Proxy Overview on page 531
- Understanding Interfaces on page 2441

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet

Copyright © 2016, Juniper Networks, Inc.

Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325

Copyright © 2016, Juniper Networks, Inc.

- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

Services Configuration Statement Hierarchy

Use the statements in the **services** configuration hierarchy to configure application identification, probes, and Unified Access Control.

```
services {
  application-identification {
    application application-name {
      address-mapping address-name {
        filter {
          ip ip-address-and-prefix-length;
          port-range {
            tcp [port];
            udp [port];
          }
        }
      }
      cacheable;
      description;
      icmp-mapping {
        code number;
        type number;
      }
      ip-protocol-mapping {
        protocol number;
      }
      order number;
      over protocol-type ;
      priority [high | low];
    }
    application-group group-name {
      application-groups application-group-name;
      applications application-name;
    }
    application-system-cache-timeout value;
    download {
      automatic {
        interval hours;
        start-time  MM-DD.hh:mm;
      }
      url url;
    }
    enable-performance-mode max-packet-threshold number;
    no-application-identification;
```

```
    no-application-system-cache;
    statistics {
      interval minutes;
    }
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level [all | error | info | notice | verbose | warning]
      no-remote-trace;
    }
  }
  captive-portal {
    authentication-profile-name authentication-profile-name;
    custom-options {
      banner-message string;
      footer-bgcolor hex-color-value;
      footer-message string;
      footer-text-color hex-color-value;
      form-header-bgcolor hex-color-value;
      form-header-message string;
      form-header-text-color hex-color-value;
      form-reset-label label name;
      form-submit-label label name;
      header-bgcolor hex-color-value;
      header-logo filename;
      header-message string;
      header-text-color hex-color-value;
      post-authentication-url url-string;
    }
    interface (all | interface-name) {
      quiet-period seconds;
      retries number-of-retries;
      server-timeout seconds;
      session-expiry seconds;
      supplicant (multiple | single | single-secure);
    }
    secure-authentication (http | https);
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
    }
  }
  flow-monitoring {
    version9 {
```


Copyright © 2016, Juniper Networks, Inc.

```
      template template-name {
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        ipv4-template;
        ipv6-template;
        option-refresh-rate {
          packets packets;
          seconds seconds;
        }
        template-refresh-rate {
          packets packets;
          seconds seconds;
        }
      }
    }
  }
  ip-monitoring {
    policy policy-name {
      match {
        rpm-probe [probe-name];
      }
      no-preempt ;
      then {
        interface interface-name (disable | enable);
        preferred-route {
          route destination-address {
            next hop next-hop;
            preferred-metric metric;
            }
          routing-instances name;
        }
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  rpm {
    bgp {
      data-fill data;
      data-size size;
      destination-port port;
      history-size size;
      logical-system logical-system-name <routing-instances routing-instance-name>;
      moving-average-size number-of-samples;
      probe-count count;
      probe-interval seconds;
      probe-type type;
```

```
    routing-instances {
      routing-instance-name;
    }
    test-interval seconds;
  }
  probe owner {
    test test-name {
      data-fill data;
      data-size size;
      destination-interface interface-name;
      destination-port port;
      dscp-code-point dscp-bits;
      hardware-timestamp;
      history-size size;
      inet6-options {
        source-address address;
      }
      moving-average-size number;
      next-hop next-hop;
      one-way-hardware-timestamp;
      probe-count count;
      probe-interval seconds;
      probe-type type;
      routing-instance instance-name;
      source-address address;
      target {
        address ipv4-address;
        url url;
        inet6-address ipv6-address;
        inet6-url url;
      }
      test-interval interval;
      thresholds {
        egress-time microseconds;
        ingress-time microseconds;
        jitter-egress microseconds;
        jitter-ingress microseconds;
        jitter-rtt microseconds;
        rtt microseconds;
        std-dev-egress microseconds;
        std-dev-ingress microseconds;
        std-dev-rtt microseconds;
        successive-loss count;
        total-loss count;
      }
      traps [ trap-names];
    }
  }
probe-limit number;
  probe-server {
    icmp {
      destination-interface interface-name;
    }
    tcp {
      destination-interface interface-name;
      port port-number;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    udp {
      destination-interface interface-name;
      port port-number;
    }
  }
service-device-pools {
  pool pool-name {
    interface service-device-name;
  }
}
service-interface-pools {
  pool pool-name {
    interface service-interface-name;
  }
}
ssl {
  initiation {
    profile profile-name {
      actions {
        ignore-server-auth-failure;
      }
      client-certificate;
      custom-ciphers [cipher];
      enable-flow-tracing;
      enable-session-cache;
      preferred-ciphers (custom | medium | strong | weak);
      protocol-version (all | tls1 | tls11 | tls12);
      trusted-ca (all | [ca-profile] );
    }
  }
  proxy {
    global-config {
      session-cache-timeout seconds;
    }
    profile profile-name {
      actions {
        crl {
          disable;
          if-not-present (allow | drop);
          ignore-hold-instruction-code;
        }
        disable-session-resumption;
        ignore-server-auth-failure;
        logs {
          all;
          errors;
          info;
          sessions-allowed;
          sessions-dropped;
          sessions-ignored;
          sessions-whitelisted;
          warning;
        }
        renegotiation {
          (allow | allow-secure | drop);
```

```
            }
          }
          custom-ciphers [cipher];
          enable-flow-tracing;
          preferred-ciphers (custom | medium | strong | weak);
          root-ca root-certificate;
          trusted-ca (all | [ca-profile] );
          whitelist [global-address-book-addresses];
        }
      }
      termination {
        profile profile-name {
          custom-ciphers [cipher];
          enable-flow-tracing;
          enable-session-cache;
          preferred-ciphers (custom | medium | strong | weak);
          protocol-version (all | tls1 | tls11 | tls12);
          server-certificate certificate-identifier;
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        level [brief | detail | extensive | verbose];
        no-remote-trace;
      }
    }
    unified-access-control {
      captive-portal redirect-policy-name {
        redirect-traffic (all | unauthenticated);
        redirect-url redirect-url;
      }
      certificate-verification [ optional | required | warning ];
      infranet-contoller host-name {
        address ip-address;
        ca-profile [ca-profile];
        interface interface-name;
        password password;
        port port-number;
        server-certificate-subject subject;
      }
      interval seconds;
      test-only-mode;
      timeout seconds;
      timeout-action (close | no-change | open);
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
```


Copyright © 2016, Juniper Networks, Inc.

```
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
user-identification {
    active-directory-access {
        domain domain-name {
            user username;
                password password;
            domain-controller domain-controller-name {
                address domain-controller-address;
            }
            ip-user-mapping {
                discovery-method {
                    wmi {
                        event-log-scanning-interval seconds;
                        initial-event-log-timespan hours;
                    }
                }
            }
            user-group-mapping {
                ldap {
                    address ip-address {
                        port port;
                    }
                    authentication-algorithm {
                        simple;
                    }
                    base base;
                    ssl;
                    user username {
                        password password;
                    }
                }
            }
        }
        authentication-entry-timeout minutes;
        filter {
            include address;
            exclude address;
        }
        no-on-demand-probe;
        wmi-timeout seconds;
        traceoptions {
            file file;
            flag {
                active-directory-authentication;
                all;
                configuration;
                db;
                ip-user-mapping;
                ip-user-probe;
                ipc;
```

```
            user-group-mapping;
            wmic;
          }
          level {
            all;
            error;
            info;
            notice;
            verbose;
            warning;
          }
          no-remote-trace;
        }
      }
      wireless-wan {
        adapter adapter-name {
          adapter-type cx-bridge;
          ip-address ip-address;
          modem {
            usb1 description description;
            usb2 description description;
            usb3 description description;
          }
        }
      }
    }
```

Related Documentation

- Understanding AppSecure Services on page 483
- Understanding Unified Access Control on page 5269

System Configuration Statement Hierarchy

Use the statements in the **system** configuration hierarchy to configure system management functions including addresses of the Domain Name System (DNS) servers; device's hostname, address, and domain name; health monitoring; interface filtering; properties of the device's auxiliary and console ports; security profiles for logical systems; time zones and Network Time Protocol (NTP) properties; trace options; and user login accounts, including user authentication and the root-level user account. Statement descriptions that are exclusive to the SRX Series devices running Junos OS are described in this section.

```
system {
  accounting {
    destination {
      radius {
        server server-address {
          accounting-port port-number;
          max-outstanding-requests number;
          port number;
          retry number;
          secret password;
          source-address address;
          timeout seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      tacplus {
        server server-address {
          port port-number;
          secret password;
          single-connection;
          source-address source-address;
          timeout seconds;
        }
      }
    }
    events [change-log interactive-commands login];
    traceoptions {
      file {
        filename;
        files number;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  allow-v4mapped-packets;
  archival {
    configuration {
      archive-sites url {
        password password;
      }
      transfer-interval interval;
      transfer-on-commit;
    }
  }
  arp {
    aging-timer minutes;
    gratuitous-arp-delay seconds;
    gratuitous-arp-on-ifup;
    interfaces {
      interface name {
        aging-timer minutes;
      }
    }
    passive-learning;
    purging;
  }
  authentication-order [password radius tacplus];
  auto-configuration {
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
```

```
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  auto-snapshot;
  autoinstallation {
    configuration-servers {
      url {
        password password;
      }
    }
    interfaces {
      interface-name {
        bootp;
        rarp;
      }
    }
    usb {
      disable;
    }
  }
  auto-snapshot;
  backup-router {
    address;
    destination [network];
  }
  commit {
    server {
      commit-interval seconds;
      days-to-keep-error-logs days;
      maximum-aggregate-pool number;
      maximum entries number;
      traceoptions {
        file {
          filename;
          files number;
          microsecond-stamp;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
    synchronize;
  }
  compress-configuration-files;
  default-address-selection;
  diag-port-authentication {
    encrypted-password passsword;
    plain-text-password;
  }
  domain-name domain-name;
  domain-search [domain-list];
  donot-disable-ip6op-ondad;
```


Copyright © 2016, Juniper Networks, Inc.

```
dump-device (boot-device | compact-flash | usb);
dynamic-profile-options {
  versioning;
}
encrypt-configuration-files;
extensions {
  providers {
    provider-id {
      license-type license deployment-scope [deployments];
    }
  }
  resource-limits {
    package package-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size mbytes;
          locked-in mbytes;
          resident-set-size mbytes;
          socket-buffers mbytes;
          stack-size mbytes;
        }
      }
    }
    process process-ui-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size mbytes;
          locked-in mbytes;
          resident-set-size mbytes;
          socket-buffers mbytes;
          stack-size mbytes;
        }
      }
    }
  }
}
fips {
  level (0 | 1 | 2 | 3 | 4);
```

Copyright © 2016, Juniper Networks, Inc.

```
}
host-name hostname;
inet6-backup-router {
  address;
  destination destination;
}
internet-options {
  icmpv4-rate-limit {
    bucket-size seconds;
    packet-rate packets-per-second;
  }
  icmpv6-rate-limit {
    bucket-size seconds;
    packet-rate packets-per-second;
  }
  (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
  ipv6-duplicate-addr-detection-transmits number;
  (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
  ipv6-path-mtu-discovery-timeout minutes;
  no-tcp-reset (drop-all-tcp | drop-tcp-with-syn-only);
  no-tcp-rfc1323;
  no-tcp-rfc1323-paws;
  (path-mtu-discovery | no-path-mtu-discovery);
  source-port upper-limit upper-limit;
  (source-quench | no-source-quench);
  tcp-drop-synfin-set;
  tcp-mss bytes;
}
kernel-replication;
license {
  autoupdate {
  url url;
  password password;
  }
  renew {
    before-expiration number;
    interval interval-hours;
  }
  traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
```


Copyright © 2016, Juniper Networks, Inc.

```
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
login {
    announcement text;
    class class-name {
        access-end hh:mm;
        access-start hh:mm;
        allow-commands regular-expression;
        allow-configuration regular-expression;
        allow-configuration-regexps [regular-expression];
        allowed-days [day];
        deny-commands regular-expression;
        deny-configuration regular-expression;
        deny-configuration-regexps [regular-expression];
        idle-timeout minutes;
        logical-system logical-system;
        login-alarms;
        login-script script;
        login-tip;
        permissions [permissions ];
        security-role (audit-administrator | crypto-administrator | ids-administrator |
            security-administrator);
    }
    deny-sources {
        address [address-or-hostname];
    }
    message text;
    }
    password {
        change-type (character-set | set-transitions);
        format (des | md5 | sha1);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        lockout-period time;
        maximum-time seconds;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication {
            encrypted-password password;
            load-key-file url;
            plain-text-password;
            ssh-dsa public-key;
            ssh-rsa public-key;
```

```
        }
        class class-name;
        full-name complete-name;
        uid uid-value;
    }
}
log-vital {
    interval minutes;
    files days;
    storage-limit percentage;
    file-size Mbytes;
    add oid{
        comment comment;
    }
    group {
        operating;
        idp;
        storage;
        cluster-counter;
        screen zone-name;
        spu spu-name;
        }
    }
max-configuration-rollbacks number;
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server ip-address;
nd-maxmcast-solicit value;
nd-retrasmit-timer value;
no-compress-configuration-files;
no-debugger-on-alt-break;
no-multicast-echo;
no-neighbor-learn;
no-ping-record-route;
no-ping-time-stamp;
no-redirects;
no-saved-core-context;
ntp {
    authentication-key key-number {
        type md5;
        value password;
    }
    boot-server address;
    broadcast broadcast-address  {
        key key;
        ttl value;
        version version;
    }
    broadcast-client;
    multicast-client {
        address;
    }
    peer peer-address {
        key key;
        prefer;
        version version;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    server server-address {
      key key;
      prefer;
      version version;
    }
    source-address source-address;
    trusted-key [key-number];
  }
  pic-console-authentication {
    encrypted-password password;
    plain-text-password;
  }
  ports {
    auxiliary {
      disable;
      insecure;
      type (ansi | small-xterm | vt100 | xterm);
    }
    console {
      disable;
      insecure;
      log-out-on-disconnect;
      type (ansi | small-xterm | vt100 | xterm);
    }
  }
  processes {
    802.1x-protocol-daemon {
      command binary-file-path;
      disable;
    }
    adaptive-services {
      command binary-file-path;
      disable;
      failover (alternate-media | other-routing-engine);
    }
    alarm-control {
      command binary-file-path;
      disable;
      failover (alternate-media | other-routing-engine);
    }
    application-identification {
      command binary-file-path;
      disable;
      failover (alternate-media | other-routing-engine);
    }
    application-security {
      command binary-file-path;
      disable;
      failover (alternate-media | other-routing-engine);
    }
    audit-process {
      command binary-file-path;
      disable;
    }
    auto-configuration {
```

```
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  bootp {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  chassis-control {
    disable;
    failover alternate-media;
  }
  class-of-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  craft-control {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  database-replication {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  datapath-trace-service {
    disable;
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  dhcp {
    command binary-file-path;
    disable;
  }
  dhcp-service {
    disable;
    failover (alternate-media | other-routing-engine);
    interface-traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
```


Copyright © 2016, Juniper Networks, Inc.

```
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    traceoptions {
        file {
            filename ;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
dialer-services {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
diameter-service {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
disk-monitoring {
    command binary-file-path;
    disable;
}
dynamic-flow-capture {
    command binary-file-path;
    disable;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
ecc-error-logging {
  command binary-file-path;
  disable;
}
ethernet-connectivity-fault-management {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
ethernet-link-fault-management {
  command binary-file-path;
  disable;
}
ethernet-switching {
  command binary-file-path;
  disable;
}
event-processing {
  command binary-file-path;
  disable;
}
fipsd {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
firewall {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
firewall-authentication-service {
  disable;
}
forwarding {
  command binary-file-path;
  disable;
}
general-authentication-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
gprs-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
```


Copyright © 2016, Juniper Networks, Inc.

```
}
group-key-member {
   disable;
}
group-key-server {
   disable;
}
idp-policy {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ilmi {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
inet-process {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
init {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
interface-control {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ipmi {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ipsec-key-management {
   (disable | enable);
}
jsrp-service {
   disable;
}
jtasktest {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
kernel-replication {
   command binary-file-path;
   disable;
}
l2-learning {
   command binary-file-path;
   disable;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
l2cpd-service {
  command binary-file-path;
  disable;
}
lacp {
  command binary-file-path;
  disable;
}
lldpd-service {
  command binary-file-path;
  disable;
}
logical-system-mux {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
logical-system-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
mib-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mobile-ip {
  command binary-file-path;
  disable;
}
mountd-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mspd {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
multicast-snooping {
  command binary-file-path;
  disable;
}
named-service {
  disable;
```


Copyright © 2016, Juniper Networks, Inc.

```
    failover (alternate-media | other-routing-engine);
}
neighbor-liveness {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
network-security {
    disable;
}
network-security-trace {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
nfsd-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
ntp {
    disable;
    failover (alternate-media | other-routing-engine);
}
ntpd-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
peer-selection-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
periodic-packet-services {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pgcp-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pgm {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pic-services-logging {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
ppp {
    command binary-file-path;
```

```
    disable;
}
pppoe {
    command binary-file-path;
    disable;
}
process-monitor {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error |info |notice | verbose | warning);
        no-remote-trace;
    }
}
profilerd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
r2cp {
    command binary-file-path;
    disable;
}
redundancy-interface-process {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
remote-operations {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
resource-cleanup {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error |info |notice | verbose | warning);
        no-remote-trace;
    }
}
routing {
```


Copyright © 2016, Juniper Networks, Inc.

```
    disable;
    failover (alternate-media | other-routing-engine);
  }
  sampling {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  sbc-configuration-process {
    disable;
    failover (alternate-media | other-routing-engine);
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  sdk-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error |info |notice | verbose | warning);
      no-remote-trace;
    }
  }
  secure-neighbor-discovery {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  security-log {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  send {
    disable;
  }
  service-deployment {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
shm-rtsdbd {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
simple-mail-client-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
smtpd-service {
   disable;
}
snmp {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
static-subscribers {
   disable;
}
statistics-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
subscriber-management {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
subscriber-management-helper {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
system-health-management {
   disable;
}
system-log-vital {
   disable;
}
tunnel-oamd {
   command binary-file-path;
   disable;
}
uac-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
usb-control {
   command binary-file-path;
   disable;
}
virtualization-service {
```


Copyright © 2016, Juniper Networks, Inc.

```
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  vrrp {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  wan-acceleration {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  watchdog {
    enable;
    disable;
    timeout value;
  }
  web-management {
    disable;
    failover (alternate media | other-routing-engine);
  }
  wireless-lan-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  wireless-wan-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
```

```
        flag flag;
        no-remote-trace;
      }
    }
  proxy {
    password password;
    port port-number;
    server url;
    username user-name;
  }
  radius-options {
    attributes {
      nas-ip-address nas-ip-address;
      }
    password-protocol mschap-v2;
  }
  radius-server server-address {
    accounting-port number;
    max-outstanding-requests number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
  }
  root-authentication {
    encrypted-password password;
    load-key-file url;
    plain-text-password;
    ssh-dsa public-key {
      <from pattern-list>;
    }
    ssh-rsa public-key {
      <from pattern-list>;
    }
  }
  saved-core-context;
  saved-core-files number;
  scripts {
    commit {
      allow-transients;
      direct-access;
      file filename {
        checksum (md5 | sha-256 | sha1);
        optional;
        refresh;
        refresh-from url;
        source url;
      }
      refresh;
      refresh-from url;
      traceoptions {
        file {
          filename;
          files number;
          size maximum-file-size;
```


Copyright © 2016, Juniper Networks, Inc.

```
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
    load-scripts-from-flash;
    op {
      file filename {
        arguments name {
          description text;
        }
        checksum (md5 | sha-256 | sha1);
        command filename-alias;
        description cli-help-text;
        refresh;
        refresh-from url;
        source url;
      }
      no-allow-url;
      refresh;
      refresh-from url;
      traceoptions {
        file {
          filename;
          files number;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
  }
  security-profile security-profile-name {
      address-book {
        maximum amount;
        reserved amount;
      }
      appfw-profile {
        maximum amount;
        reserved amount;
      }
      appfw-rule {
        maximum amount;
        reserved amount;
      }
      appfw-rule-set {
        maximum amount;
        reserved amount;
      }
      auth-entry {
        maximum amount;
        reserved amount;
      }
      cpu {
        reserved percent;
```

```
}
dslite-softwire-initiator {
    maximum amount;
    reserved amount;
}
flow-gate {
    maximum amount;
    reserved amount;
}
flow-session {
    maximum amount;
    reserved amount;
}
idp-policy idp-policy-name;
logical-system logical-system-name;
nat-cone-binding {
    maximum amount;
    reserved amount;
}
nat-destination-pool {
    maximum amount;
    reserved amount;
}
nat-destination-rule {
    maximum amount;
    reserved amount;
}
nat-interface-port-ol {
    maximum amount;
    reserved amount;
}
nat-nopat-address {
    maximum amount;
    reserved amount;
}
nat-pat-address {
    maximum amount;
    reserved amount;
}
nat-pat-portnum {
    maximum amount
    reserved amount
}
nat-port-ol-ipnumber {
    maximum amount;
    reserved amount;
}
nat-rule-referenced-prefix {
    maximum amount;
    reserved amount;
}
nat-source-pool {
    maximum amount;
    reserved amount;
}
nat-source-rule {
```


Copyright © 2016, Juniper Networks, Inc.

```
            maximum amount;
            reserved amount;
        }
        nat-static-rule {
            maximum amount;
            reserved amount;
        }
        policy {
            maximum amount;
            reserved amount;
        }
        policy-with-count {
            maximum amount;
            reserved amount;
        }
        root-logical-system;
        scheduler {
            maximum amount;
            reserved amount;
        }
        zone {
            maximum amount;
            reserved amount;
        }
    }
    security-profile-resources {
        cpu-control;
        cpu-control-target percent;
    }
    services {
        database-replication {
            traceoptions {
                file {
                    filename ;
                    files number;
                    match regular-expression;
                    size maximum-file-size;
                    (world-readable | no-world-readable);
                }
                flag flag;
                no-remote-trace;
            }
        }
        dhcp {
            boot-file filename;
            boot-server (address | hostname);
            default-lease-time (infinite | seconds);
            domain-name domain-name;
            domain-search dns-search-suffix;
            maximum-lease-time (infinite | seconds);
            name-server ip-address;
            next-server ip-address;
            option option-identifier-code array type-name [ type-values ] | byte 8-bit-value | flag
                (false | off | on | true) | integer signed-32-bit-value | ip-address address | short
                signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
                unsigned-short 16-bit-value);
```

Copyright © 2016, Juniper Networks, Inc.

```
    pool subnet-ip-address/mask {
      address-range {
        high address;
        low address;
      }
      boot-file filename;
      boot-server (address | hostname);
      default-lease-time (infinite | seconds);
      domain-name domain-name;
      domain-search dns-search-suffix;
      exclude-address ip-address;
      maximum-lease-time (infinite | seconds);
      name-server ip-address;
      next-server ip-address;
      option option-identifier-code array type-name [ type-values ] | byte 8-bit-value |
        flag (false | off | on | true) | integer signed-32-bit-value | ip-address address |
        short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
        unsigned-short 16-bit-value);
      propagate-ppp-settings interface-name;
      propagate-settings interface-name;
      router ip-address;
      server-identifier dhcp-server;
      sip-server {
        address ip-address;
        name sip-server-name;
      }
      wins-server ip-address;
    }
    propagate-ppp-settings interface-name;
    propagate-settings interface-name;
    router ip-address;
    server-identifier dhcp-server;
    sip-server {
      address ip-address;
      name sip-server-name;
    }
    static-binding mac-address;
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
    wins-server ip-address;
  }
  dhcp-local-server {
    dhcpv6 {
      authentication {
        password password;
        username-include {
```


Copyright © 2016, Juniper Networks, Inc.

```
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix;
    }
}
dynamic-profile {
    profile-name;
    aggregate-clients {
        merge;
        replace;
    }
    junos-default-profile;
    use-primary dynamic-profile-name;
}
group group-name {
    authentication {
        password password;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
        }
    }
    dynamic-profile {
        profile-name;
        aggregate-clients {
            merge;
            replace;
        }
        junos-default-profile;
        use-primary dynamic-profile;
    }
    interface interface-name {
        dynamic-profile {
            profile-name;
            aggregate-clients {
                merge;
                replace;
            }
            junos-default-profile;
```

```
      use-primary dynamic-profile-name;
    }
    exclude;
    overrides {
      delegated-pool pool-name;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      rapid-commit ;
    }
    service-profile service-profile-name
    trace ;
    upto interface-name;
  }
  liveness-detection {
    failure-action {
      clear-binding;
      clear-binding-if-interface-up;
      log-only;
    }
    method {
      bfd {
        detection-time {
          threshold milliseconds;
        }
        holddown-interval interval;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        session-mode (automatic | multihop | single-hop);
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
      version (0 | 1 | automatic);
    }
  }
  overrides {
    delegated-pool pool-name;
    interface-client-limit number;
    process-inform {
      pool pool-name;
    }
    rapid-commit ;
  }
  reconfigure {
    attempts number;
    clear-on-abort;
    strict;
    timeout number;
    token token-name;
    trigger {
      radius-disconnect;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            service-profile service-profile-name;
        }
        liveness-detection {
            failure-action {
                clear-binding;
                clear-binding-if-interface-up;
                log-only;
            }
            method {
                bfd {
                    detection-time {
                        threshold milliseconds;
                    }
                    holddown-interval interval;
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    session-mode (automatic | multihop | single-hop);
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                version (0 | 1 | automatic);
            }
        }
        overrides {
            delegated-pool pool-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            rapid-commit ;
        }
        reconfigure {
            attempts number;
            clear-on-abort;
            strict;
            timeout number;
            token token-name;
            trigger {
                radius-disconnect;
            }
        }
        service-profile service-profile-name;
    }
    group group-name {
        interface interface-name {
            exclude;
            upto upto-interface-name;
        }
    }
}
dns {
    dns-proxy {
```

```
            cache hostname inet ip-address;
            default-domain domain-name {
              forwarders ip-address;
            }
            interface interface-name;
            propogate-setting (enable | disable);
            view view-name {
              domain domain-name {
                forward-only;
                forwarders ip-address;
              }
              match-clients subnet-address;
            }
          }
        }
        dnssec {
          disable;
          dlv {
            domain-name domain-name trusted-anchor  trusted-anchor;
            }
        secure-domains domain-name;
          trusted-keys (key dns-key | load-key-file url);
        forwarders {
          ip-address;
        }
        max-cache-ttl seconds;
        max-ncache-ttl seconds;
        traceoptions {
          category {
            category-type;
          }
          debug-level level;
          file {
            filename;
            files number;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          level (all | error | info | notice | verbose | warning);
          no-remote-trace;
        }
      }
      dynamic-dns {
        client hostname {
          agent agent-name;
          interface interface-name;
          password server-password;
          server server-name;
          username user-name;
        }
      }
      finger {
        connection-limit number;
        rate-limit number;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
ftp {
  connection-limit number;
  rate-limit number;
}
netconf {
  ssh {
    connection-limit number;
    port port-number;
    rate-limit number;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    on-demand;
  }
}
outbound-ssh {
  client client-id {
    address {
      port port-number;
      retry number;
      timeout value;
    }
    device-id device-id;
    keep-alive {
      retry number;
      time-out value;
    }
    reconnect-strategy (in-order |sticky);
    secret secret;
    services {
      netconf;
    }
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
service-deployment {
  local-certificate certificate-name;
  servers server-address {
```

```
        port port-number;
        security-options {
            ssl3;
            tls;
        }
        user user-name;
    }
    source-address source-address;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
ssh {
    ciphers [cipher];
    client-alive-count-max number;
    client-alive-interval seconds;
    connection-limit number;
    hostkey-algorithm {
        (ssh-dss | no-ssh-dss);
        (ssh-ecdsa |no-ssh-ecdsa);
        (ssh-rsa | no-ssh-rsa);
    }
    key-exchange [algorithm];
    macs [algorithm];
    max-sessions-per-connection number;
    protocol-version {
        v1;
        v2;
    }
    rate-limit number;
    root-login (allow | deny | deny-password);
    (tcp-forwarding | no-tcp-forwarding);
}
subscriber-management {
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    maintain-subscriber interface-delete;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
subscriber-management-helper {
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
telnet {
  connection-limit number;
  rate-limit number;
}
web-management {
  control {
    max-threads number;
  }
  http {
    interface [interface-name];
    port port-number;
  }
  https {
    interface [interface-name];
    local-certificate name;
    pki-local-certificate name;
    port port-number;
    system-generated-certificate;
  }
  management-url url;
  session {
    idle-timeout minutes;
    session-limit number;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
xnm-clear-text {
  connection-limit number;
  rate-limit number;
}
xnm-ssl {
```

```
      connection-limit number;
      local-certificate name;
      rate-limit number;
    }
  }
  static-host-mapping hostname {
    alias [host-name-alias];
    inet [ip- address];
    inet6 [ipv6- address];
    sysid system-identifier;
  }
  syslog {
    allow-duplicates;
    archive {
      binary-data;
      files number;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    console {
      (any | facility) severity;
    }
    file filename {
      allow-duplicates;
      archive {
        archive-sites url {
          password password;
        }
        (binary-data| no-binary-data);
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
      }
      structure-data {
        brief;
      }
      (any | facility) severity;
    }
    host (hostname | other-routing-engine) {
      (any | facility) severity;
    }
    log-rotate-frequency minutes;
    source-address source-address;
    time-format {
      millisecond;
      year;
    }
    user (username | *) {
      (any | facility) severity;
    }
  }
  tacplus-options {
    (exclude-cmd-attribute | no-cmd-attribute-value);
    service-name service-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    tacplus-server server-address {
      port port-number;
      secret password;
      single-connection;
      source-address source-address;
      timeout seconds;
    }
    time-zone (GMThour-offset | time-zone);
    tracing {
      destination-override {
        syslog {
          host address;
          }
      }
    }
    use-imported-time-zones;
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security address-book] Hierarchy Level

```
security {
  address-book (book-name | global) {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
    attach {
      zone zone-name;
    }
    description text;
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Address Books on page 1069

[edit security application-firewall] Hierarchy Level

```
security {
  application-firewall {
    profile profile-name {
      block-message type {
        custom-text content custom-html-text;
        custom-redirect-url content custom-redirect-url;
      }
    }
    rule-sets rule-set-name {
      default-rule {
        (deny [block-message] | permit | reject [block-message]);
      }
      profile profile-name;
      rule rule-name {
        match {
          dynamic-application [system-application];
          dynamic-application-group [system-application-group];
          ssl-encryption (any | yes | no);
        }
        then {
          (deny [block-message] | permit | reject [block-message]);
        }
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Application Firewall Overview on page 555

[edit security application-tracking] Hierarchy Level

```
security {
  application-tracking {
    disable;
    (first-update | first-update-interval first-update-interval);
    session-update-interval session-update-interval;
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Example: Configuring AppTrack on page 574

[edit security log] Hierarchy Level

```
security {
  log {
    cache {
      exclude exclude-name {
        destination-address destination-address;
        destination-port destination-port;
        event-id event-id;
        failure;
        interface-name interface-name;
        policy-name policy-name;
        process process-name;
        protocol protocol;
        source-address source-address;
        source-port source-port;
        success;
        user-name user-name;
      }
      limit value;
    }
    disable;
    event-rate rate;
    facility-override;
    file {
      files max-file-number;
      name file-name;
      path binary-log-file-path;
      size maximum-file-size;
    }
    format (binary | sd-syslog | syslog);
    mode (event | stream);
    source-address source-address | source-interface interface-name;
    stream stream-name {
      category(all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec
        | idp | rtlog | pst-ds-lite | appqos | secintel);
      filter {
        threat-attack;
      }
      format (binary | sd-syslog | syslog | welf);
      host {
        ip-address;
        port port-number;
      }
      severity (alert | critical | debug | emergency | error | info | notice | warning);
    }
    traceoptions {
      file {
        file-name;
        files max-file-number;
```

Copyright © 2016, Juniper Networks, Inc.

```
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
        transport {
          protocol (udp | tcp | tls);
          tls-profile tls-profile-name;
          tcp-connections tcp-connections;
        }
        utc-time-stamp;
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- SSL Proxy Overview on page 531

[edit security pki] Hierarchy Level

```
security {
  pki {
    auto-re-enrollment {
      cmpv2 {
        certificate-id certificate-id-name {
          ca-profile-name ca-profile-name ;
          re-enroll-trigger-time-percentage percentage ;
          re-generate-keypair;
        }
      }
      scep {
        certificate-id certificate-id-name {
          ca-profile-name ca-profile-name ;
          challenge-password password ;
          re-enroll-trigger-time-percentage percentage ;
          re-generate-keypair;
        }
      }
    }
    ca-profile ca-profile-name {
      administrator {
        e-mail-address e-mail-address;
      }
      ca-identity ca-identity ;
      enrollment {
          retry number;
          retry-interval seconds;
          url url-name;
      }
      revocation-check {
        crl {
          disable {
            on-download-failure;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            refresh-interval hours;
            url url-name;
          }
          disable;
          ocsp {
            connection-failure (disable | fallback-crl);
            disable-responder-revocation-check;
            nonce-payload (enable | disable);
            url ocsp-url;
          }
          use-ocsp;
        }
        routing-instance routing-instance-name ;
      }
      traceoptions {
        file filename {
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          destination-address-excluded;
          source-address {
            [address];
            any;
```

```
      any-ipv4;
      any-ipv6;
      }
    source-address-excluded;
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
          captive-portal captive-portal;
        }
        utm-policy policy-name;
      }
      destination-address {
        drop-translated;
        drop-untranslated;
      }
      firewall-authentication {
        pass-through {
          access-profile profile-name;
          client-match user-or-group-name;
          ssl-termination-profile profile-name;
          web-redirect;
          web-redirect-to-https;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
  global {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        from-zone {
          [zone-name];
          any;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
```

```
    to-zone {
      [zone-name];
      any;
    }
}
scheduler-name scheduler-name;
then {
  count {
    alarm {
      per-minute-threshold number;
      per-second-threshold number;
    }
  }
  deny;
  log {
    session-close;
    session-init;
  }
  permit {
    application-services {
      application-firewall {
        rule-set rule-set-name;
      }
      application-traffic-control {
        rule-set rule-set-name;
      }
      gprs-gtp-profile profile-name;
      gprs-sctp-profile profile-name;
      idp;
      redirect-wx | reverse-redirect-wx;
      ssl-proxy {
        profile-name profile-name;
      }
      uac-policy {
        captive-portal captive-portal;
      }
      utm-policy policy-name;
    }
    destination-address {
      drop-translated;
      drop-untranslated;
    }
    firewall-authentication {
      pass-through {
        access-profile profile-name;
        client-match user-or-group-name;
        ssl-termination-profile profile-name;
        web-redirect;
        web-redirect-to-https;
      }
      user-firewall {
        access-profile profile-name;
        domain domain-name
        ssl-termination-profile profile-name;
      }
      web-authentication {
```


Copyright © 2016, Juniper Networks, Inc.

```
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable);
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

[edit security zones] Hierarchy Level

```
security {
  zones {
    functional-zone {
      management {
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
        interfaces interface-name  {
```

```
            host-inbound-traffic {
              protocols protocol-name {
                except;
              }
              system-services service-name {
                except;
              }
            }
          }
          screen screen-name;
        }
      }
      security-zone zone-name {
        address-book {
          address address-name {
            ip-prefix {
              description text;
            }
            description text;
            dns-name domain-name {
              ipv4-only;
              ipv6-only;
            }
            range-address lower-limit to upper-limit;
            wildcard-address ipv4-address/wildcard-mask;
          }
          address-set address-set-name {
            address address-name;
            address-set address-set-name;
            description text;
          }
        }
        application-tracking;
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
        interfaces interface-name {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
        tcp-rst;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

}

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Security Zones and Interfaces Overview on page 1049

Copyright © 2016, Juniper Networks, Inc.

actions (Services SSL Proxy)

Syntax

```
actions {
  crl {
    disable;
    if-not-present (allow | drop);
    ignore-hold-instruction-code;
  }
  disable-session-resumption;
  ignore-server-auth-failure;
  logs {
    all;
    errors;
    info;
    sessions-allowed;
    sessions-dropped;
    sessions-ignored;
    sessions-whitelisted;
    warning;
  }
  renegotiation {
    (allow | allow-secure | drop);
  }
}
```

Hierarchy Level [edit services ssl proxy profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the logging and traffic related actions.

Options

- **crl**—Specify the certificate revocation actions.
  - **disable**—Disable CRL verification.
  - **if-not-present**—Specify actions for sessions.
    - **allow**—Allow sessions when CRL information is not available.
    - **drop**—Drop sessions when CRL information is not available.
  - **ignore-hold-instruction-code**—Ignore the unconfirmed (on hold) revocation status, and accept a certificate.
- **disable-session-resumption**—Disable session resumption.
- **ignore-server-auth-failure**—Ignore server authentication failure.
- **log**—Specify the logging actions.
  - **all**—Log all events.
  - **errors**—Log all error events.
  - **info**—Log all information events.
  - **sessions-allowed**—Log SSL session allowed events after an error.

Copyright © 2016, Juniper Networks, Inc.

- **sessions-dropped**—Log only SSL session dropped events.
- **sessions-ignored**—Log session ignored events.
- **sessions-whitelisted**—Log SSL session whitelisted events.
- **warning**—Log all warning events.

- **renegotiation**—Specify the renegotiation options.
  - **allow**—Allow secure and nonsecure renegotiation.
  - **allow-secure**—Allow secure negotiation only.
  - **drop**—Drop session on renegotiation request.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

actions (Services SSL Initiation)

Syntax
```
actions {
    ignore-server-auth-failure;
}
```

Hierarchy Level
[edit services ssl initiation profile *profile-name*]

Release Information
Statement introduced in Junos OS Release 12.1X44-D10.

Description
Specify the logging and traffic related actions.

Options
- **ignore-server-auth-failure**—Ignore server authentication failure.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

address-mapping (Application Identification)

Syntax

```
address-mapping address-name {
  filter {
    ip ip-address-and-prefix-length;
    port-range {
      tcp [port];
      udp [port];
    }
  }
}
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Defines an application by the IP address and the port range of the traffic.

Options filter—Specify the application matching criteria by the IP address of the application or the port range to match TCP or UDP destination port.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509
- *AppSecure Services Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

appfw-profile (System)

Syntax

```
appfw-profile {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the application firewall profile quota of a logical system.

Options

- **maximum *amount***—Specify the maximum allowed quota value.

  Range: 0 through 1024

- **reserved *amount***—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

appfw-rule

Syntax

```
appfw-rule {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the number of application firewall rule configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to the master logical system and the user logical systems
- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

 Copyright © 2016, Juniper Networks, Inc.

appfw-rule-set

Syntax

```
appfw-rule-set {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the number of application firewall rule set configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to the master logical system and the user logical systems
- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-firewall

Syntax

```
application-firewall {
  profile profile-name {
    block-message type {
      custom-text content custom-html-text;
      custom-redirect-url content custom-redirect-url;
    }
  }
  rule-sets rule-set-name {
    default-rule {
      (deny [block-message] | permit | reject [block-message]);
    }
    profile profile-name;
    rule rule-name {
      match {
        dynamic-application [system-application];
        dynamic-application-groups [system-application-group];
        ssl-encryption (any | yes | no);
      }
      then {
        (deny [block-message] | permit | reject [block-message]);
      }
    }
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (world-readable | no-world-readable);
      size maximum-file-size;
    }
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 11.1. Updated with the **ssl-encryption** and **reject** options in Junos OS Release 12.1X44-D10. Updated with the **block-message** option in Junos OS Release 12.1X45-D10.

Description Specify the profile options, rule set and rule specifications, and trace options to be used for application firewall implementations.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Application Firewall Overview on page 555

application (Application Identification)

Syntax

```
application application-name {
  address-mapping address-name {
    filter {
      ip ip-address-and-prefix-length;
      port-range {
        tcp [port];
        udp [port];
      }
    }
  }
  cacheable;
  description;
  icmp-mapping {
    code number;
    type number;
  }
  ip-protocol-mapping {
    protocol number;
  }
  over protocol-type  {
    signature name  {
      member name  {
        context {
          http-get-url-parsed-param-parsed;
          http-header-content-type;
          http-header-cookie;
          http-header-host;
          http-header-user-agent;
          http-post-url-parsed-param-parsed;
          http-post-variable-parsed ;
          http-url-parsed;
          http-url-parsed-param-parsed;
          ssl-server-name;
          stream;
        }
        direction {
          any;
          client-to-server;
          server-to-client;
        }
        pattern pattern;
      }
      port-range value;
  priority [high | low];
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Configure a custom application definition for the desired application name that will be used by the system to identify the application as it passes through the device. Custom

Copyright © 2016, Juniper Networks, Inc.

application definitions can be used for applications that are not part of the Juniper Networks predefined application database.

Options

application *application-name*—Name of the custom application signature. Must be a unique name with a maximum length of 63 characters.

NOTE: Application names are case insensitive.

cacheable—Enable caching of application identification results. By enabling this option, you can cache the application detection result in an ASC table. If there is an entry in the ASC table, based on the destination IP address, protocol, and the port, we can identify AppID without again sending packet to engine. This option is not supported for address-based, IP protocol-based, and ICMP-based custom application signatures.

description—Description of the application.

order *number*—Specify the order for the custom application. Lower order has higher priority. This option is used when multiple custom applications match the criteria.

priority [high | low]—Specify the priority over other signature applications.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- *AppSecure Services Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

application-firewall (Application Services)

Syntax

```
application-firewall {
  rule-set rule-set-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure application firewall rule sets with rules defining match criteria and the action to be performed.

Options **rule-set** *rule-set-name*—Name of the rule set that contains application firewall specification rules.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-identification

Syntax

```
application-identification {
  application application-name {
    address-mapping address-name {
      filter {
        ip ip-address-and-prefix-length;
        port-range {
          tcp [port];
          udp [port];
        }
      }
    }
    cacheable;
    description;
    icmp-mapping {
      code number;
      type number;
    }
    ip-protocol-mapping {
      protocol number;
    }
    over protocol-type  {
      signature name  {
        member name  {
          context {
            http-get-url-parsed-param-parsed;
            http-header-content-type;
            http-header-cookie;
            http-header-host;
            http-header-user-agent;
            http-post-url-parsed-param-parsed;
            http-post-variable-parsed ;
            http-url-parsed;
            http-url-parsed-param-parsed;
            ssl-server-name;
            stream;
          }
          direction {
            any;
            client-to-server;
            server-to-client;
          }
          pattern pattern;
        }
        port-range value {
    priority [high | low];
  }
  application-group group-name {
    application-groups application-group-name;
    applications application-name;
  }
  application-system-cache-timeout value;
  download {
```

```
    automatic {
      interval hours;
      start-time  MM-DD.hh:mm;
    }
    url url;
  }
  enable-performance-mode max-packet-threshold number;
  no-application-identification;
  no-application-system-cache;
  statistics {
    interval minutes;
  }
  traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level [all | error | info | notice | verbose | warning]
    no-remote-trace;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 10.2. Custom application definition option introduced in Junos OS Release 15.1X49-D40.

Description Configure application identification options to identify the TCP or UDP application session running on nonstandard ports to match the application properties of transiting network traffic.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Application Identification Techniques on page 485

 Copyright © 2016, Juniper Networks, Inc.

application-group (Services)

Syntax

```
application-group group-name {
    application-groups application-group-name;
    applications application-name;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify any number of associated predefined applications, user-defined applications, and other groups for ease of use in configuring application-based policies.

An application group is hierarchical: a tree structure of groups with applications as the leaf nodes.

Options

group-name—Name of the group. This name is used in policy configuration statements in place of multiple predefined applications, user-defined applications, or other groups.

application-groups *application-group-name*— Name of an application group to be assigned to this group. There is no maximum number of groups that can be assigned to a group. Use multiple commands to assign multiple groups.

applications *application-name*—Name of an application to be assigned to this group. An application can remain unassigned or be assigned to a group, but it cannot be assigned to more than one group. There is no maximum number of applications that can be assigned to a group. Use multiple commands to assign multiple groups.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Custom Application Group for Junos OS Application Identification for Simplified Management on page 518

Copyright © 2016, Juniper Networks, Inc.

application-services (Security Policies)

Syntax

```
application-services {
  application-firewall {
    rule-set rule-set-name;
  }
  application-traffic-control {
    rule-set rule-set-name;
  }
  gprs-gtp-profile profile-name;
  gprs-sctp-profile profile-name;
  idp;
  redirect-wx | reverse-redirect-wx;
  ssl-proxy {
    profile-name profile-name;
  }
  uac-policy {
    captive-portal captive-portal;
  }
  utm-policy policy-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement modified in Junos OS Release 11.1.

Description Enable application services within a security policy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-system-cache

Syntax application-system-cache;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.2.

Description When a session is created, specify an application ID to match the application properties of transiting network traffic. The application port mappings are saved in the application system cache.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding the Application System Cache on page 523

application-system-cache-timeout (Services)

Syntax application-system-cache-timeout *value*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 9.2. Support for application identification in the services hierarchy added in Junos OS Release 10.2.

Description Specify the timeout value in seconds for the application system cache entries. Note the cache is not cleared when IDP policy is loaded. Users need to manually clear or wait for the cache entries to expire.

NOTE: On SRX Series devices, when you change the timeout value for the application system cache entries using the command **set services application-identification application-system-cache-timeout**, the cache entries need to be cleared to avoid inconsistency in timeout values of existing entries.

Options *value*—Timeout value for the application system cache entries.

Range: 0 through 1,000,000 seconds

Default: 3600 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding the Application System Cache on page 523

Copyright © 2016, Juniper Networks, Inc.

application-tracking

Syntax

```
application-tracking {
  disable;
  (first-update | first-update-interval first-update-interval);
  session-update-interval session-update-interval;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 10.2. Support for **disable** added in Junos OS Release 11.4.

Description AppTrack, an application tracking tool, is a form of statistical profiling. Enabling this feature for a zone logs flow statistics (the byte count, packet count, and start and end times for a session) at session end. You can modify the logging time and log frequency with command options. Periodically, a network management tool, such as STRM, collects the logged statistics sent by each network device for bandwidth usage analysis of the network.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

application-tracking (Security Zones)

Syntax

```
application-tracking;
```

Hierarchy Level [edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable application tracking support for the zone.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574
- [edit security zones] Hierarchy Level on page 324

Copyright © 2016, Juniper Networks, Inc.

application-traffic-control

Syntax

```
application-traffic-control {
  rate-limiters rate-limiter-name{
    bandwidth-limit kbps;
      burst-size-limit bytes;
    }
  }
  rule-sets ruleset-name {
    rule rule-name {
      match {
        application application-name1;
        application-any;
        application-group application-group-name;
        application-known;
        application-unknown;
      }
      then {
        dscp-code-point dscp-value;
        forwarding-class forwarding-class-name;
        log;
        loss-priority [ high | medium-high | medium-low | low ];
        rate-limit {
          loss-priority-high;
          client-to-server rate-limiter-name;
          server-to-client rate-limiter-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 11.4.

Description Mark DSCP values for outgoing packets or apply rate limits based on the specified Layer 7 application types.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

application-traffic-control (Application Services)

Syntax

```
application-traffic-control {
  rule-set rule-set-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enables AppQoS, application-aware quality of service, as specified in the rules of the specified rule set.

Options

- **rule-set *rule-set-name***—Name of the rule set that contains application-aware traffic control specification rules.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589
- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

block-message (Application Firewall)

Syntax

```
block-message type {
  custom-text content custom-html-text;
  custom-redirect-url content custom-redirect-url;
}
```

Hierarchy Level [edit security application-firewall profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Defines the profile of the notification to be sent to clients when HTTP or HTTPS traffic is blocked by a reject or deny action from an application firewall.

NOTE: The block message option is not supported for non-HTTP traffic. In these instances, if the action is drop or reject, the traffic is silently dropped or rejected. The user is not informed of the action and no redirection occurs. The associated system log message identifies the action taken for this traffic.

When the **block-message** option is specified, a splash screen and message inform the client that the traffic has been blocked. The default message text is:

> "*username*, Application Firewall has blocked your request to application *application-name* at *dest-ip:dest-port* accessed from *src-ip:source-port* "

The variables in the message are replaced with specific traffic values. For clarity, the prefix **junos:** is truncated from the application name.

Options Use the following option pairs to customize the default message or to redirect the client to a custom webpage instead of the default splash screen.

NOTE: Both the type and content fields must be used to add custom text or redirect the client to a URL.

- **type**—(Optional) The message type to be displayed after a reject or deny action.
  - **custom-text**—Text message in HTML to be added to the default text. If **custom-text** is specified, the splash screen displays both the default block message and the custom-defined block message.

    When specified, the user is redirected when a reject or deny action is taken during one of the following HTTP methods: GET, POST, OPTIONS, HEAD, PUT, DELETE, TRACE, CONNECT, PROPFIND, PROPPATCH, LOCK, UNLOCK, COPY, MOVE, MKCOL, BCOPY, BDELETE, BCOPY, BMOVE, BPROPFIND, BPROPPATCH, POLL, SEARCH, SUBSCRIBE, and UNSUBSCRIBE. If the reject or deny action occurs during a different HTTP method, the traffic is silently dropped.

Copyright © 2016, Juniper Networks, Inc.

- **custom-redirect-url**—URL redirection.

- **content**—(Optional) Message content for the selected message type.

  NOTE: The content value must match the type option selected: custom-text requires text, and custom-redirect-url requires a URL value.

  - *custom-text*—Custom text to be added to the splash screen. Custom text is inserted below the default message. Add the characters \n to insert a line break in the displayed text.
  - *custom-redirect-url*—The URL of the webpage to which the client is directed. When traffic is rejected or denied, the client is redirected to the specified webpage for further action. The URL can be hosted on either the SRX Series device or an external server.

    Enter the redirect URL in quotation marks for an HTTP or HTTPS site, as shown in the following examples:

    "http://*custom-redirect-url*"
    "https://*custom-redirect-url*"

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589

Copyright © 2016, Juniper Networks, Inc.

context (Application Identification)

Syntax

```
context {
  http-get-url-parsed-param-parsed;
  http-header-content-type;
  http-header-cookie;
  http-header-host;
  http-header-user-agent;
  http-post-url-parsed-param-parsed;
  http-post-variable-parsed ;
  http-url-parsed;
  http-url-parsed-param-parsed;
  ssl-server-name;
  stream;
}
```

Hierarchy Level [edit services application-identification application *application-name* over *protocol-type* signature *name* member *name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify context for matching application running over TCP, UDP, or Layer 7.

Options

http-get-url-parsed-param-parsed—The decoded, normalized GET URL in an HTTP request along with the decoded CGI parameters (if any).

http-header-content-type —The content-type header in an HTTP transaction.

http-header-cookie—The cookie header in an HTTP transaction.

http-header-host —The host header in an HTTP transaction.

http-header-user-agent—The user-agent header in an HTTP transaction.

http-post-url-parsed-param-parsed —The decoded, normalized POST URL in an HTTP request along with the decoded CGI parameters (if any).

http-post-variable-parsed—The decoded POST URL or form data variables.

http-url-parsed—The decoded, normalized URL in an HTTP request.

http-url-parsed-param-parsed—The decoded, normalized URL in an HTTP request along with the decoded CGI parameters (if any).

ssl-server-name —Server name in the TLS server name extension or the SSL server certificate. This is also known as Server Name Indication (SNI).

stream —TCP or UDP stream data.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509

custom-ciphers

Syntax

```
custom-ciphers [ rsa-with-rc4-128-md5 RSA | 128bit rc4 | md5 hash rsa-with-rc4-128-sha
    RSA | 128bit rc4 |sha hash rsa-with-des-cbc-sha RSA | des cbc | sha hash
    rsa-with-3des-ede-cbc-sha RSA | 3des ede/cbc | sha hash rsa-with-aes-128-cbc-sha RSA
    | 128 bit aes/cbc | sha hash rsa-with-aes-256-cbc-sha RSA | 256 bit aes/cbc |sha hash
    rsa-export-with-rc4-40-md5 RSA-export | 40 bit rc4 | md5 hash
    rsa-export-with-des40-cbc-sha RSA-export | 40 bit des/cbc | sha hash rsa-with-null-md5
    RSA | no symmetric cipher | md5 hash rsa-with-null-sha RSA | no symmetric cipher | sha
    hash];
```

Hierarchy Level

[edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Description

Display the custom cipher list.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

default-rule

Syntax

```
default-rule {
  (deny [block-message] | permit | reject [block-message]);
}
```

Hierarchy Level [edit security application-firewall rule-sets *rule-set-name*]

Release Information Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 with the **reject** option. The **block-message** option added in Junos OS Release 12.1X45-D10.

Description Configure the default rule that defines the actions to be performed on a packet that does not match any defined rule.

Note that an application firewall is applied after a session has already been created by the security firewall. When traffic is rejected or denied by an application firewall, therefore, logs contain a session open message, a session reject or deny message, and a session close message.

Options

- **deny**—Block the traffic at the firewall. The device drops the packet. No message is returned to the sender.
  - **block-message**—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the **profile** option for this rule set, including the **block-message** option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.
- **permit**—Permit traffic at the firewall.
- **reject**—Block the traffic at the firewall. For TCP traffic, by default the device drops the packet and returns a TCP reset (RST) message to the source host and to the server in some cases. For UDP and other protocol traffic, by default the device drops the packet and returns an ICMP "destination unreachable, port unreachable" message to both the client and the server.
  - **block-message**—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the **profile** option for this rule set, including the **block-message** option displays a default message or customized message, or redirects the user for rejected HTTP or HTTPS traffic. All other traffic is dropped as specified in the default action for the **reject** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

Copyright © 2016, Juniper Networks, Inc.

direction (Application Identification)

Syntax

```
direction {
  any;
  client-to-server;
  server-to-client;
}
```

Hierarchy Level [edit services application-identification application *application-name* over *protocol-type* signature *name* member *name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description The connection direction of the packets to apply pattern matching.

Options

any—The directions of packets are either from client-side to server-side or from server-side to client-side.

client-to-server—The direction of packets are from client-side to server-side.

server-to-client—The direction of packets are from server-side to client-side.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509

disable (Application Tracking)

Syntax

```
disable;
```

Hierarchy Level [edit security application-tracking]

Release Information Statement introduced in Junos OS Release 11.4.

Description Disable application tracking on a device without deleting the zone configuration. Application tracking is enabled by default.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

download (Services)

Syntax

```
download {
  automatic {
    interval hours;
    start-time MM-DD.hh:mm;
  }
url url;
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure automatic download for the application identification services application package. The application package contains definitions for known applications, such as: DNS, Facebook, FTP, Skype, and SNMP. The application package is extracted from the IDP signature database located at https://signatures.juniper.net. If you do not have access to the default download site from your device, you can use the URL option to download from a different location.

NOTE: You need to download the application package before configuring application identification services.

Options

- automatic—Download the application package automatically at a certain time of day or at intervals.
- *interval*—Download the application package at intervals.

  Range: 6 through 720 hours

- *start-time*—Start time in which the application package will be download. Format is MM-DD.hh:mm. Example: 04-15.09:00 will start the download on April 15 at 9 AM.
- *url*—Use this option to change the default download location of the application package.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Scheduling the Application Signature Package Updates on page 499

Copyright © 2016, Juniper Networks, Inc.

dynamic-application

Syntax dynamic-application [*system-application*];

Hierarchy Level [edit security application-firewall rule-sets *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 11.1.

Description Specify the dynamic application names for match criteria.

Options ***system-application***—Set of system applications for match criteria.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Application Firewall Overview on page 555

dynamic-application-group

Syntax dynamic-application-group [*system-application-group*];

Hierarchy Level [edit security application-firewall rule-sets *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the dynamic application group to match.

Options ***system-application-group***—Set of groups defining one or more system applications for match criteria.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

enable-flow-tracing (Services)

Syntax enable-flow-tracing;

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Enable flow tracing for the profile.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

enable-performance-mode

Syntax enable-performance-mode max-packet-threshold *number*;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Set the deep packet inspection (DPI) in performance mode with default packet inspection limit as two packets, including both client-to-server and server-to-client directions.

Options **max-packet-threshold** *number*—Set the maximum packet threshold for DPI performance mode.

Range: 1 through 100.

Default: 2.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Improving the Application Traffic Throughput on page 528
- show services application-identification status on page 825

Copyright © 2016, Juniper Networks, Inc.

enable-session-cache

Syntax enable-session-cache;

Hierarchy Level [edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Enable SSL session cache.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

file (Services)

Syntax
```
file file-name; {
  files;
  match;
  no-world-readable size;
  world-readable;
}
```

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the trace file information.

Options
- **files**—Specify the maximum number of trace files. Range: 2 through 1000.
- **match**—Specify the regular expression for lines to be logged.
- **no-world-readable size**—Do not allow any user to read the log file.
- **size**—Specify the maximum trace file size. Range: 10,240 to 1,073,741,824.
- **world-readable**—Allow any user to read the log file.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

files (Services)

Syntax files *files*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the maximum number of trace files.

Options **files**—Specify the maximum number of trace files.

Range: 2 to 1000

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

file (System Logging)

Syntax

```
file filename {
  allow-duplicates;
  any (alert | any | critical | emergency | error | info | none | notice | warning);
  archive {
    archive-sites {
      url password;
    }
    (binary-data | no-binary-data);
    files number;
    size size;
    start-time start-time;
    transfer-interval transfer-interval;
    (world-readable | no-world-readable);
  }
  authorization (alert | any | critical | emergency | error | info | none | notice | warning);
  change-log (alert | any | critical | emergency | error | info | none | notice | warning);
  conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
  daemon (alert | any | critical | emergency | error | info | none | notice | warning);
  dfc (alert | any | critical | emergency | error | info | none | notice | warning);
  explicit-priority;
  external (alert | any | critical | emergency | error | info | none | notice | warning);
  firewall (alert | any | critical | emergency | error | info | none | notice | warning);
  ftp (alert | any | critical | emergency | error | info | none | notice | warning);
  interactive-commands (alert | any | critical | emergency | error | info | none | notice | warning);
  kernel (alert | any | critical | emergency | error | info | none | notice | warning);
  match "regular-expression";
  ntp (alert | any | critical | emergency | error | info | none | notice | warning);
  pfe (alert | any | critical | emergency | error | info | none | notice | warning);
  security (alert | any | critical | emergency | error | info | none | notice | warning);
  structured-data {
    brief;
  }
  user (alert | any | critical | emergency | error | info | none | notice | warning);
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 12.1X47 for SRX Series.

Description Specify the file in which to log data.

Options

- *filename*—Specify the name of the file in which to log data.
- *allow-duplicates*—Do not suppress the repeated messages.
- *any*—Specify all facilities information.
  - *alert*—Specify the conditions that should be corrected immediately.
  - *critical*—Specify the critical conditions.
  - *emergency*—Specify the conditions that cause security functions to stop.
  - *error*—Specify the general error conditions.

Copyright © 2016, Juniper Networks, Inc.

    - *info*—Specify the information about normal security operations.
    - *none*—Do not specify any messages.
    - *notice*—Specify the conditions that should be handled specifically.
    - *warning*—Specify the general warning conditions.
- *archive*—Specify the archive file information.
    - *archive-sites*—Specify a list of destination URLs for the archived log files.
        - *url*—Specify the primary and failover URLs to receive archive files.
    - *binary-data*—Mark file such that it contains binary data.
    - *no-binary-data*—Do not mark the file such that it contains binary data.
    - *files*—Specify the number of files to be archived. Range: 1 through 1000 files.
    - *size*—Specify the size of files to be archived. Range: 65,536 through 1,073,741,824 bytes.
    - *world-readable*—Allow any user to read the log file.
    - *no-world-readable*—Do not allow any user to read the log file.
    - *start-time*—Specify the start time for file transmission. Enter the start time in the yyyy-mm-dd.hh:mm format.
    - *transfer-interval*—Specify the frequency at which to transfer the files to archive sites.
- *authorization*—Specify the authorization system.
- *change-log*—Specify the configuration change log.
- *conflict-log*—Specify the configuration conflict log.
- *daemon*—Specify the various system processes.
- *dfc*—Specify the dynamic flow capture.
- *explicit-priority*—Include the priority and facility in messages.
- *external*—Specify the local external applications.
- *firewall*—Specify the firewall filtering system.
- *ftp*—Specify the FTP process.
- *interactive-commands*—Specify the commands executed by the UI.
- *kernel*—Specify the kernel information.
- *match*—Specify the regular expression for lines to be logged.
- *ntp*—Specify the NTP process.
- *pfe*—Specify the Packet Forwarding Engine.
- *security*—Specify the security-related information.

Copyright © 2016, Juniper Networks, Inc.

- *structured-data*—Log the messages in structured log format.
  - *brief*—Omit English language text from the end of the logged message.
- *user*—Specify the user processes.
  - *info*—Specify the informational messages.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Junos OS System Log Overview*
- *syslog (System)*

first-update

Syntax
first-update;

Hierarchy Level
[edit security application-tracking]

Release Information
Statement introduced in Junos OS Release 10.2.

Description
Generate an AppTrack start message when a new session begins. (A final message is produced at session end with any option.) This option overrides the **first-update-interval** option if both are specified.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

first-update-interval

Syntax first-update-interval *first-update-interval*;

Hierarchy Level [edit security application-tracking]

Release Information Statement introduced in Junos OS Release 10.2.

Description For long-lived sessions being monitored by AppTrack, configure this value to issue the first update message after a specified number of minutes.

NOTE: The **first-update-interval** setting is disregarded if the **first-update** option is set to log the first message at session start.

Options ***minutes***—Maximum number of minutes after session start for the first update message to be sent. This value must be smaller than the **session-update-interval** setting.

Default: 1

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

flag (Services)

Syntax flag (*all* | *cli-configuration* | *initiation* | *proxy* | *selected-profile* | *termination*);

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the tracing flag parameters.

Options
- *all*—Trace all the parameters.
- *cli-configuration*—Trace CLI configuration events.
- *initiation*—Trace initiation service events.
- *proxy*—Trace proxy service events.
- *selected-profile*—Trace events for profiles with enable-flow-tracing set.
- *termination*—Trace termination service events.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

format (Security Log)

Syntax format (binary | sd-syslog | syslog)

Hierarchy Level [edit security log]

Release Information Statement introduced prior to Junos OS Release 10.0. Statement updated in Junos OS Release 12.1.

Description Set the default log format for event mode security logging on the device.

Options
- **binary**—Binary encoded text to conserve resources.
- **sd-syslog**—Structured system log file.
- **syslog**—Traditional system log file.

Default: syslog.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

Copyright © 2016, Juniper Networks, Inc.

forwarding-classes (CoS)

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- *class-name*—Display the forwarding class name assigned to the internal queue number.

  NOTE: This option is supported only on high-end SRX Series devices, including the SRX1500, SRX5400, SRX5600, and SRX5800.

  NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **policing-priority**—Layer 2 policing. One forwarding class can be configured as **premium** and others are configured as **normal**.
- **priority**—Fabric priority value:
  - **high**—Forwarding class's fabric queuing has high priority.
  - **low**—Forwarding class's fabric queuing has low priority.
- *queue-number*—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, either **high** or **low**.

  NOTE: The spu-priority option is only supported on SRX1500, and SRX5000 line devices.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppQoS on page 589

global-config (Services)

Syntax
```
global-config {
  session-cache-timeout seconds;
}
```

Hierarchy Level
[edit services ssl proxy]

Release Information
Statement introduced in Junos OS Release 12.1X44-D10.

Description
Specify the global proxy configuration.

Options
session-cache-timeout—Specify the session cache timeout.

Range: 300 to 3600 seconds

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

icmp-mapping (Application Identification)

Syntax

```
icmp-mapping {
  code number;
  type number;
}
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify the Internet Control Message Protocol (ICMP) value for an application to match. The ICMP mapping technique maps standard ICMP message types and optional codes to a unique application name. This mapping technique lets you differentiate between various types of ICMP messages.

Options

- **code *number***—Numeric value of an ICMP code. The code field provides further information about the associated type field.
- **type *number***—Numeric value of an ICMP type. The type field identifies the ICMP message.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509

ip (Application Identification)

Syntax ip *ip-address-and-prefix-length*;

Hierarchy Level [edit services application-identification application *application-name* address-mapping *address-name* filter]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify the IP address and the prefix length of the application for address mapping.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509

Copyright © 2016, Juniper Networks, Inc.

ip-protocol-mapping (Application Identification)

Syntax

```
ip-protocol-mapping {
    protocol number;
}
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify the IP protocol value for an application to match. Standard IP protocol numbers can map an application to IP traffic. As with address mapping, to ensure adequate security, use IP protocol mapping only in your private network for trusted servers.

Options **protocol *number***—Numeric value of an IP protocol

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Junos OS Application Identification Custom Application Signatures on page 509

 Copyright © 2016, Juniper Networks, Inc.

initiation (Services)

Syntax

```
initiation {
  profile profile-name {
    actions {
      ignore-server-auth-failure;
    }
    client-certificate;
    custom-ciphers [cipher];
    enable-flow-tracing;
    enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | tls1 | tls11 | tls12);
    trusted-ca (all | [ca-profile] );
  }
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) initiation support service.

Options

- **client-certificate**—Local certificate.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

level (Services)

Syntax level [*brief* | *detail* | *extensive* | *verbose]*];

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the level of debugging the output.

Options
- *brief*—Specify brief debugging output.
- *detail*—Specify detailed debugging output.
- *extensive*—Specify extensive debugging output.
- *verbose*—Specify verbose debugging output.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

log (Security)

Syntax

```
log {
  cache {
    exclude exclude-name {
      destination-address destination-address;
      destination-port destination-port;
      event-id event-id;
      failure;
      interface-name interface-name;
      policy-name policy-name;
      process process-name;
      protocol protocol;
      source-address source-address;
      source-port source-port;
      success;
      user-name user-name;
    }
    limit value;
  }
  disable;
  event-rate rate;
  facility-override;
  file {
    files max-file-number;
    name file-name;
    path binary-log-file-path;
    size maximum-file-size;
  }
  format (binary | sd-syslog | syslog);
  mode (event | stream);
  rate-cap rate-cap-value;
  (source-address source-address | source-interface interface-name);
  stream stream-name {
    category (all | content-security);
    format (binary | sd-syslog | syslog | welf);
    host {
      ip-address;
      port port-number;
    }
    severity (alert | critical | debug | emergency | error | info | notice | warning);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
  transport {
```

Copyright © 2016, Juniper Networks, Inc.

```
    protocol (udp | tcp | tls);
    tls-profile tls-profile-name;
    tcp-connections tcp-connections;
  }
  utc-time-stamp;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.2.

Description You can set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

- **disable**—Disable the security logging for the device.
- **event-rate** *rate*—Limits the rate (0 through 1500) at which logs will be streamed per second.
- **rate-cap** *rate-cap-value*—Works with event mode only. Limits the rate (0 through 5000) at which data plane logs will be generated per second.
- **source-address** *source-address*—Specify a source IP address or IP address used when exporting security logs.
- **source-interface** *interface-name*—Specify a source interface name, which is mandatory to configure **stream**.

NOTE: The **source-address** and **source-interface** are alternate values. Using one of the options is mandatory.

- **utc-time-stamp**—Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

log (Services)

Syntax

```
log {
  all;
  errors;
  info;
  sessions-allowed;
  sessions-dropped;
  sessions-ignored;
  sessions-whitelisted;
  warning;
}
```

Hierarchy Level [edit services ssl proxy profile *profile-name* actions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the logging actions.

Options

- **all**—Log all events.
- **errors**—Log all error events.
- **info**—Log all information events.
- **sessions-allowed**—Log SSL session allowed events after an error.
- **sessions-dropped**—Log only SSL session dropped events.
- **sessions-ignored**—Log session ignored events.
- **sessions-whitelisted**—Log SSL session whitelisted events.
- **warning**—Log all warning events.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

match (Services)

Syntax match *match*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the regular expression for lines to be logged.

Options **match**—Specify the regular expression for lines to be logged.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

no-application-identification (Services)

Syntax no-application-identification;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Disable the TCP/UDP application identification of applications running on nonstandard ports.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Disabling and Reenabling Junos OS Application Identification on page 507

Copyright © 2016, Juniper Networks, Inc.

no-application-system-cache (Services)

Syntax no-application-system-cache;

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Application identification information is saved in the application system cache to improve performance. This cache is updated when a different application is identified. This caching is turned on by default. Use the **no-application-system-cache** statement to turn it off.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Deactivating Application System Cache Information for Application Identification (CLI Procedure) on page 524

no-remote-trace (Services)

Syntax no-remote-trace;

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Disable remote tracing.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

over (Application Identification)

Syntax

```
over protocol-type  {
  signature name  {
    member name  {
      context {
        http-get-url-parsed-param-parsed;
        http-header-content-type;
        http-header-cookie;
        http-header-host;
        http-header-user-agent;
        http-post-url-parsed-param-parsed;
        http-post-variable-parsed ;
        http-url-parsed;
        http-url-parsed-param-parsed;
        ssl-server-name;
        stream;
      }
      direction {
        any;
        client-to-server;
        server-to-client;
      }
      pattern pattern;
    }
    port-range value;
```

Hierarchy Level [edit services application-identification application *application-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify application running over TCP, UDP, or Layer 7.

Options

signature *name* —Name of the custom application signature. Must be a unique name with a maximum length of 63 characters.

member *name* —Member name for a custom application signature. Custom signatures can contain multiple members that define attributes for an application. (The supported member name range is m01 through m15.)

context—Service-specific context, such as http-header-content-type.

direction—Connection direction of the packets to match pattern

patterns—(Optional) Deterministic finite automaton (DFA) pattern matched on the context. The DFA pattern specifies the pattern to be matched for the signature. Maximum length is 128.

port-range—Port range. This option is applicable for TCP or UDP-based applications only.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Understanding Junos OS Application Identification Custom Application Signatures on page 509

policies

Syntax

```
policies {
  default-policy (deny-all | permit-all);
  from-zone zone-name to-zone zone-name {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
        deny;
        log {
          session-close;
          session-init;
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
}
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
```

```
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      from-zone {
        [zone-name];
        any;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
      to-zone {
        [zone-name];
        any;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
          redirect-wx | reverse-redirect-wx;
          ssl-proxy {
            profile-name profile-name;
          }
          uac-policy {
```


Copyright © 2016, Juniper Networks, Inc.

```
                captive-portal captive-portal;
              }
              utm-policy policy-name;
            }
            destination-address {
              drop-translated;
              drop-untranslated;
            }
            firewall-authentication {
              pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
              }
              web-authentication {
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable) ;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Release Information

Statement introduced in Junos OS Release 8.5.
Support for the **services-offload** option added in Junos OS Release 11.4.
Support for the **source-identitiy** option added in Junos OS Release 12.1.
Support for the **description** option added in Junos OS Release 12.1.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.
Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.
Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description

Configure network security policies.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy (Security Policies)

Syntax

```
policy policy-name {
  description description;
  match {
    application {
      [application];
      any;
    }
    destination-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
```

Copyright © 2016, Juniper Networks, Inc.

```
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              web-redirect;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. The **services-offload** option added in Junos OS Release 11.4. Statement updated with the **source-identity** option and the **description** option added in Junos OS Release 12.1. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Define a security policy.

 Copyright © 2016, Juniper Networks, Inc.

Options ***policy-name***—Name of the security policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320
- *Building Blocks Feature Guide for Security Devices*

port-range (Application Identification)

Syntax
```
port-range {
  tcp [port];
  udp [port];
}
```

Hierarchy Level [edit services application-identification application *application-name* address-mapping *address-name* filter]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify a port to match a TCP or UDP destination port.

Options
- **tcp [*port*]**—Define the TCP port range for the application.
- **udp [*port*]**—Define the UDP port range for the application.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Understanding Junos OS Application Identification Custom Application Signatures on page 509

Copyright © 2016, Juniper Networks, Inc.

preferred-ciphers

Syntax preferred-ciphers (custom | medium | strong | weak);

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Select the preferred ciphers.

Options
- **custom**—Configure custom cipher suite and order of preference.
- **medium**—Use ciphers with key strength of 128 bits or greater.
- **strong**—Use ciphers with key strength of 168 bits or greater.
- **weak**—Use ciphers with key strength of 40 bits or greater.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- SSL Proxy Overview on page 531

profile (Application Firewall)

Syntax
```
profile profile-name {
  block-message type {
    custom-text content custom-html-text;
    custom-redirect-url content custom-redirect-url;
  }
}
```

Hierarchy Level [edit security application-firewall]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Define the profile of the response to be issued when an application firewall rule set blocks HTTP or HTTPS traffic with a **deny** or **reject** action. You can display a default or custom message, or redirect traffic to a URL where an explanation or further action is provided.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

profile (Rule Sets)

Syntax profile *profile-name*;

Hierarchy Level [edit security application-firewall rule-sets *rule-set-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specifies the profile of the block message to be used for any deny or reject action in the rule set that specifies the **block-message** option.

Options *profile-name*—Name of the block-message profile to be used for this rule set.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

profile (Services)

Syntax

```
profile profile-name {
  actions {
    crl {
      disable;
      if-not-present (allow | drop);
      ignore-hold-instruction-code;
    }
    disable-session-resumption;
    ignore-server-auth-failure;
    logs {
      all;
      errors;
      info;
      sessions-allowed;
      sessions-dropped;
      sessions-ignored;
      sessions-whitelisted;
      warning;
    }
    renegotiation {
      (allow | allow-secure | drop);
    }
  }
  custom-ciphers [cipher];
  enable-flow-tracing;
  preferred-ciphers (custom | medium | strong | weak);
  root-ca root-certificate;
  trusted-ca (all | [ca-profile] );
  whitelist [global-address-book-addresses];
}
```

Hierarchy Level [edit services ssl proxy]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the SSL server profile.

Options *profile-name*—Specify the profile identifier.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

profile (SSL Initiation)

Syntax

```
profile profile-name {
  actions {
    ignore-server-auth-failure;
  }
  client-certificate;
  custom-ciphers [cipher];
  enable-flow-tracing;
  enable-session-cache;
  preferred-ciphers (custom | medium | strong | weak);
  protocol-version (all | tls1 | tls11 | tls12);
  trusted-ca (all | [ca-profile] );
}
```

Hierarchy Level [edit services ssl initiation]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the name of the profile for SSL initiation support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

profile (SSL Termination)

Syntax

```
profile profile-name {
  custom-ciphers [cipher];
  enable-flow-tracing;
  enable-session-cache;
  preferred-ciphers (custom | medium | strong | weak);
  protocol-version (all | tls1 | tls11 | tls12);
  server-certificate certificate-identifier;
}
```

Hierarchy Level [edit services ssl initiation]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the name of the profile for SSL termination support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

protocol-version

Syntax protocol-version (all | tls1 | tls11 | tls12);

Hierarchy Level [edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10. The **tls11** and **tls12** options are introduced in 15.1X49-D30.

Description Specify the accepted SSL protocol version.

Options

- **all**—Accept all versions of TLS.
- **TLS version 1.0**—Accept TLS version 1.0. It provides secure communication over networks by providing privacy and data integrity between communicating applications
- **TLS version 1.1**—Accept TLS version 1.1. This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.
- **TLS version 1.2** —Accept TLS version 1.2. This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- SSL Proxy Overview on page 531

Copyright © 2016, Juniper Networks, Inc.

proxy (Services)

Syntax

```
proxy {
  global-config {
    session-cache-timeout seconds;
  }
  profile profile-name {
    actions {
      crl {
        disable;
        if-not-present (allow | drop);
        ignore-hold-instruction-code;
      }
      disable-session-resumption;
      ignore-server-auth-failure;
      logs {
        all;
        errors;
        info;
        sessions-allowed;
        sessions-dropped;
        sessions-ignored;
        sessions-whitelisted;
        warning;
      }
      renegotiation {
        (allow | allow-secure | drop);
      }
    }
    custom-ciphers [cipher];
    enable-flow-tracing;
    preferred-ciphers (custom | medium | strong | weak);
    root-ca root-certificate;
    trusted-ca (all | [ca-profile] );
    whitelist [global-address-book-addresses];
  }
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) proxy support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

rate-limiters

Syntax

```
rate-limiters {
  rate-limiter-name {
    bandwidth-limit value-in-kbps;
    burst-size-limit value-in-bytes;
  }
}
```

Hierarchy Level [edit class-of-service application-traffic-control]

Release Information Statement introduced in Junos OS Release 11.4.

Description Share the available bandwidth and burst size of a device's PICs by defining rate limiter profiles and applying them in AppQoS rules.

Options

- ***rate-limiter-name***—Name of the rate limiter. It is applied in AppQoS rules to share device resources based on quality of service requirements.

  The combination of rate limiting parameters, namely bandwidth- limit and burst-size-limit rate limit, make up the rate limiter profile. A maximum of 16 profiles are allowed per device. The same profile can be used by multiple rate limiters. For example, a profile with a bandwidth-limit of 200 Kbps and a burst-limit of 130,000 bytes, could be used in several rate limiters.

  A maximum of 1000 rate limiters can be created. Rate limiters are defined for the device, and are assigned in rules in a rule set. A single rate limiter can be used multiple times within the same rule set. However, the rate limiter cannot be used in another rule set.

- **bandwidth-limit *value-in-Kbps***—Maximum number of kilobits to be transmitted per second for this rate limiter. Up to 2 GB of bandwidth can be provisioned among multiple rate limiters to share the resource proportionally.

- **burst-size-limit *value-in-bytes***—Maximum number of bytes to be transferred in a single burst or time-slice. This limit ensures that a high-priority transmission does not keep a lower priority transmission from transmitting.

NOTE: The number of bandwidth-limit and burst-size-limit combinations cannot exceed 16.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589

Copyright © 2016, Juniper Networks, Inc.

renegotiation (Services)

Syntax renegotiation (allow | allow-secure | drop);

Hierarchy Level [edit services ssl proxy profile *profile-name* actions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the renegotiation options.

Options
- **allow**—Allow secure and nonsecure renegotiation.
- **allow-secure**—Allow secure negotiation only.
- **drop**—Drop session on renegotiation request.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

root-ca (Services)

Syntax root-ca *root-certificate*;

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Root certificate for interdicting server certificates in proxy mode.

Options *root-ca-name*—Specify root certificate for interdicting server certificates in proxy mode.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

rule-sets (CoS AppQoS)

Syntax

```
rule-sets {
  rule-set-name {
    rule rule-name {
      match {
        application application-name;
        application-any;
        application-group application-group-name;
        application-known;
        application-unknown;
      }
      then {
        dscp-code-point dscp-value ;
        forwarding-class forwarding-class-name;
        log;
        loss-priority [ high | medium-high | medium-low | low ];
        rate-limit {
          loss-priority-high;
          client-to-server rate-limiter-name;
          server-to-client rate-limiter-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service application-traffic-control]

Release Information Statement introduced in Junos OS Release 11.4.

Description Defines AppQoS rule sets and the rules that establish priorities based on quality of service requirements for the associated applications. AppQoS rules can be included in policy statements to implement application-aware quality of service control.

Options

- ***rule-set-name***—Name used to refer to a collection of AppQoS rules.
- **rule *rule-name***—Name applied to the match criteria and resulting actions that control the quality of service provided to any matching applications.
- **application *application-name***—Name of the application to be used as match criteria for the rule.
- **application-any** —Any application encountering this rule. Note that when you use this specification, all application matching ends. Any application rule following this one will never be encountered.
- **application-group *application-group-name***—Group of applications to be used as match criteria for the rule. Both applications and application groups can be match criteria for a single rule.
- **application-known**—Match criteria specifying any session that is identified, but its corresponding application is not specified.

Copyright © 2016, Juniper Networks, Inc.

- **application-unknown**—Match criteria specifying any session that is not identified.
- **forwarding-class *forwarding-class-name***—The AppQoS class with which matching applications will be marked. This field identifies the rewriter that has marked the DSCP value . Therefore, the AppQoS forwarding class must be different from those used by IDP or firewall filters. With this class specified, firewall filter class will not overwrite the existing DSCP value.
- **dscp-code-point**—DSCP alias or bit map with which matching applications will be marked to establish the output queue. This value can be marked by rewriters from IDP, AppQoS, or a firewall filter. The forwarding-class value identifies which rewriter has re-marked the packet with the current DSCP value. If a packet triggers all three rewriters, IDP takes precedence over AppQoS, which takes precedence over a firewall filter.
- **loss-priority**—Loss priority with which matching applications will be marked. This value is used to determine the likelihood that a packet would be dropped when encountering congestion. A high loss priority means that there is an 80% chance of packet loss in congestion. Possible values are high, medium-high, medium-low and low.
- **rate-limit**—Rate limiters to be associated with client-to-server and with server-to-client traffic for this application. The rate limiter profile defines maximum speed and volume limits for matching applications.
- **log**—AppQoS event logging.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589

Copyright © 2016, Juniper Networks, Inc.

rule-sets (Security Application Firewall)

Syntax

```
rule-sets rule-set-name {
  default-rule {
    (deny [block-message] | permit | reject [block-message]);
  }
  profile profile-name;
  rule rule-name {
    match {
      dynamic-application [system-application];
      dynamic-application-groups [system-application-group];
      ssl-encryption (any | yes | no);
    }
    then {
      (deny [block-message] | permit | reject [block-message]);
    }
  }
}
```

Hierarchy Level [edit security application-firewall]

Release Information Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 to include the **ssl-encryption** and **reject** options. The **block-message** options added in Junos OS Release 12.1X45-D10.

Description Configure the set of rules for the application firewall.

Options *rule-set-name*—Name of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring an Application Group for Application Firewall on page 564

Copyright © 2016, Juniper Networks, Inc.

security-zone

Syntax

```
security-zone zone-name {
  address-book {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
  }
  application-tracking;
  description text;
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
  interfaces interface-name {
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
  }
  screen screen-name;
  tcp-rst;
}
```

Hierarchy Level [edit security zones]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description Define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Copyright © 2016, Juniper Networks, Inc.

Options ***zone-name*** —Name of the security zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security zones] Hierarchy Level on page 324
- Security Zones and Interfaces Overview on page 1049
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

server-certificate (Services)

Syntax server-certificate *server-certificate*;

Hierarchy Level [edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the local certificate identifier.

Options **server-certificate**—Specify the name of the local certificate identifier.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

session-update-interval

Syntax session-update-interval *session-update-interval*;

Hierarchy Level [edit security application-tracking]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the interval between session update messages for long-lived sessions being monitored by AppTrack. Byte count, packet count, and start and end times are updated and logged when the amount of time between session start or the previous update and the current time exceeds the interval.

Options ***session-update-interval***—Minutes between updates.

Default: 5

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppTrack on page 574

size (Services)

Syntax size *size*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the maximum trace file size.

Options **size**—Specify the maximum trace file size.

Range: 10,240 to 1,073,741,824.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

ssl (Services)

Syntax

```
ssl {
  initiation {
    profile profile-name {
      actions {
        ignore-server-auth-failure;
      }
      client-certificate;
      custom-ciphers [cipher];
      enable-flow-tracing;
      enable-session-cache;
      preferred-ciphers (custom | medium | strong | weak);
      protocol-version (all | tls1 | tls11 | tls12);
      trusted-ca (all | [ca-profile] );
    }
  }
  proxy {
    global-config {
      session-cache-timeout seconds;
    }
    profile profile-name {
      actions {
        crl {
          disable;
          if-not-present (allow | drop);
          ignore-hold-instruction-code;
        }
        disable-session-resumption;
        ignore-server-auth-failure;
        log {
          all;
          errors;
          info;
          sessions-allowed;
          sessions-dropped;
          sessions-ignored;
          sessions-whitelisted;
          warning;
        }
        renegotiation {
          (allow | allow-secure | drop);
        }
      }
      custom-ciphers [cipher];
      enable-flow-tracing;
      preferred-ciphers (custom | medium | strong | weak);
      root-ca root-certificate;
      trusted-ca (all | [ca-profile] );
      whitelist [global-address-book-addresses];
    }
  }
  termination {
    profile profile-name {
```

```
        custom-ciphers [cipher];
        enable-flow-tracing;
        enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | tls1 | tls11 | tls12);
        server-certificate certificate-identifier;
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      level [brief | detail | extensive | verbose];
      no-remote-trace;
    }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

ssl-encryption

Syntax ssl-encryption (any | no | yes);

Hierarchy Level [edit security application-firewall rule-sets *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Distinguishes between encrypted and unencrypted SSL traffic as match criteria for the rule. In application firewall usage, this option lets you specify different actions for encrypted and unencrypted SSL traffic.

Options
- **any**—Matches both encrypted and unencrypted SSL traffic.
- **no**—Matches unencrypted SSL traffic only.
- **yes**—Matches encrypted SSL traffic only.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

ssl-proxy (Application Services)

Syntax
```
ssl-proxy {
  profile-name profile-name
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 12.1.

Description Enable SSL proxy and identify the name of the SSL proxy profile to be used.

Options *profile-name*—SSL proxy profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

statistics (Services)

Syntax

```
statistics {
  interval interval-number;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the interval, in minutes, for statistics collection.

Options

interval *interval-number*—Length of time, in minutes, that application statistics are collected.

Range: 1 through 1440 minutes

Default: 1 minute

NOTE: For SRX Series devices, the maximum number of interval periods for which statistics are stored is 8.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Onbox Application Identification Statistics on page 527

Copyright © 2016, Juniper Networks, Inc.

stream (Security Log)

Syntax

```
stream stream-name {
  category (all | content-security)
  format (binary | sd-syslog | syslog | welf)
  host {
    <ipaddr> ip-address;
    port port-number;
  }
  severity (alert | critical | debug | emergency | error | info | notice | warning);
}
```

Hierarchy Level [edit security log]

Release Information Statement modified in Junos OS Release 9.2.

Description Set stream settings for a security log. You can set a maximum of three streams.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

termination (Services)

Syntax

```
termination {
  profile profile-name {
    custom-ciphers [cipher];
    enable-flow-tracing;
    enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | tls1 | tls11 | tls12);
    server-certificate certificate-identifier;
  }
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) termination support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

then (Security Application Firewall)

Syntax

```
then {
   (deny [block-message] | permit | reject [block-message]);
}
```

Hierarchy Level [edit security application-firewall rule-set *rule-set-name* rule *rule-name*]

Release Information Statement introduced in Junos OS Release 11.1. Statement updated in Junos OS Release 12.1X44-D10 with the **reject** option. The **block-message** option added in Junos OS Release 12.1X45-D10.

Description Specify the action to be performed when traffic matches the associated match criteria.

Note that an application firewall is applied after a session has already been created by the security firewall. When traffic is rejected or denied by an application firewall, therefore, logs contain a session open message, a session reject or deny message, and a session close message.

Options

- **deny**—Block the traffic at the firewall. The device drops the packet. By default, no message is returned to the sender.
  - **block-message**—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the **profile** option for this rule set, including the **block-message** option displays a default message or customized message, or redirects the user for denied HTTP or HTTPS traffic. All other traffic is dropped silently.
- **permit**—Permit traffic at the firewall.
- **reject**—Block the traffic at the firewall. For TCP traffic, by default the device drops the packet and returns a TCP reset (RST) message to the source host. For UDP and other protocol traffic, by default the device drops the packet and returns an ICMP "destination unreachable, port unreachable" message to both the client and the server.
  - **block-message**—(Optional) In application firewall rules, provide information to the user regarding blocked traffic. Depending on the content of the **profile** option for this rule set, including the **block-message** option displays a default message or customized message, or redirects the user for rejected HTTP or HTTPS traffic. All other traffic is dropped as specified in the default action for the **reject** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring an Application Group for Application Firewall on page 564

Copyright © 2016, Juniper Networks, Inc.

trusted-ca (Services)

Syntax trusted-ca (all | [*ca-profile*]);

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the list of trusted certificate authority profiles.

Options
- *trusted-ca-name*—Specify the certificate authority profile name.
- **all**—Select all certificate authority profiles.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Application Firewall)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security application-firewall]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure trace options for the application firewall.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
  - **compilation**—Trace rule set compilation events
  - **configuration**—Trace configuration events
  - **ipc**—Trace process inter communication events
  - **lookup**—Trace rule set lookup events
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Services SSL)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  level [brief | detail | extensive | verbose];
  no-remote-trace;
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the trace file information.

Options

- *file-name*—Specify the name of file in which to write trace information.
- **files**—Specify the maximum number of trace files. Range: 2 to 1000.
- **match**—Specify the regular expression for lines to be logged.
- **no-world-readable size**—Do not allow any user to read the log file.
- **size**—Specify the maximum trace file size. Range: 10,240 to 1,073,741,824.
- **world-readable**—Allow any user to read the log file.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Services Application Identification)

Syntax

```
traceoptions {
  file {
    filename ;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag all;
  level (all | error | info | notice | verbose | warning)
  no-remote-trace;
}
```

Hierarchy Level [edit services application-identification]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure tracing operations for application identification services.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
- **level**—Set the level of debugging the output option.
  - **all**—Match all levels.
  - **error**—Match error conditions.
  - **info**—Match informational messages.
  - **notice**—Match conditions that should be handled specially
  - **verbose**—Match verbose messages.
  - **warning**—Match warning messages.
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Understanding Application Identification Techniques on page 485

Copyright © 2016, Juniper Networks, Inc.

transport (Security Log)

Syntax

```
transport {
  protocol (udp | tcp | tls);
  tls-profile tls-profile-name;
  tcp-connections tcp-connections;
}
```

Hierarchy Level [edit security log]

Release Information Statement introduced in Junos OS Release 12.1X46-D25.

Description Configure security log transport options.

Options

protocol—Specify the type of transport protocol to be used to log the data.

- **UDP**—Set the transport protocol to UDP.
- **TCP**—Set the transport protocol to TCP.
- **TLS**—Set the transport protocol to TLS.

Default: UDP.

tls-profile *tls-profile-name*—Specify the TLS profile name.

tcp-connections *tcp-connections*—Specify the number of TCP connections per SPU.

Range: 1 through 5.

Default: 1.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding AppTrack on page 573

 Copyright © 2016, Juniper Networks, Inc.

whitelist (Services)

Syntax whitelist [global-address-book-addresses];

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the addresses exempted from the SSL proxy.

Options
- *whitelist-address*—Specify address from the global address book.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

zones

Syntax

```
zones {
  functional-zone {
    management {
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
      interfaces interface-name  {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
    }
  }
  security-zone zone-name {
    address-book {
      address address-name {
        ip-prefix {
          description text;
        }
        description text;
        dns-name domain-name {
          ipv4-only;
          ipv6-only;
        }
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
      }
      address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
      }
    }
    application-tracking;
    description text;
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
```


Copyright © 2016, Juniper Networks, Inc.

```
            except;
          }
        }
        interfaces interface-name {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
        tcp-rst;
      }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description A zone is a collection of interfaces for security purposes. All interfaces in a zone are equivalent from a security point of view. Configure the following zones:

- Functional zone—Special-purpose zone, such as a management zone that can host dedicated management interfaces.
- Security zone—Most common type of zone that is used as a building block in policies.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 33

Operational Commands

- clear security application-firewall rule-set statistics
- clear security application-firewall rule-set statistics logical-system
- clear services application-identification application-statistics
- clear services application-identification application-statistics cumulative
- clear services application-identification application-statistics interval
- clear services application-identification application-system-cache (Junos OS)
- clear services application-identification counter (Values)
- clear services ssl proxy statistics
- request security pki ca-certificate ca-profile-group load
- request security pki local-certificate export
- request security pki local-certificate generate-self-signed
- request security pki local-certificate load
- request services application-identification application
- request services application-identification download
- request services application-identification download status
- request services application-identification group
- request services application-identification install
- request services application-identification install status
- request services application-identification proto-bundle-status
- request services application-identification uninstall
- request services application-identification uninstall status
- show class-of-service application-traffic-control counter
- show class-of-service application-traffic-control statistics rate-limiter
- show class-of-service application-traffic-control statistics rule
- show security application-firewall rule-set
- show security application-firewall rule-set logical-system
- show security application-tracking counters
- show security flow session

Copyright © 2016, Juniper Networks, Inc.

- show security flow session application-firewall
- show security pki ca-certificate
- show security pki local-certificate (View)
- show security policies
- show services application-identification application
- show services application-identification application-system-cache (View)
- show services application-identification commit-status
- show services application-identification counter (AppSecure)
- show services application-identification group
- show services application-identification statistics applications
- show services application-identification statistics application-groups
- show services application-identification status
- show services application-identification version
- show services ssl proxy statistics

Copyright © 2016, Juniper Networks, Inc.

clear security application-firewall rule-set statistics

Syntax **clear security application-firewall rule-set statistics**

Release Information Command introduced in Junos OS Release 11.1.

Description Clear all the security application firewall rule set statistics information.

Required Privilege Level clear

Related Documentation
- show security application-firewall rule-set on page 774

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security application-firewall rule-set statistics logical-system

Syntax The master, or root, administrator can issue the following statements:

clear security application-firewall rule-set statistics [logical-system *logical-system-name* | all | root-logical-system]

The user logical system administrator can issue the following statement:

clear security application-firewall rule-set statistics all

Release Information Command introduced in Junos OS Release 11.4.

Description Clear all security application firewall rule set statistics.

NOTE: User logical system administrators can clear statistics only for the logical systems they can access. For information about master and user administrator roles in logical systems, see "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Options ***logical-system-name***—Name of a specific logical system.

all—(default) Clear all rule set statistics for a specific logical system or all logical systems.

root-logical-system—Clear application firewall rule set statistics on the root logical system (master administrator only).

Required Privilege Level clear

Related Documentation

- show security application-firewall rule-set logical-system on page 777

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear services application-identification application-statistics

Syntax clear services application-identification application-statistics

Release Information Statement introduced in Junos OS Release 11.4.

Description Clears all Junos OS application statistics such as cumulative, interval, applications, and application groups.

Required Privilege Level clear

Related Documentation

- show services application-identification statistics applications on page 821
- show services application-identification statistics application-groups on page 823
- clear services application-identification application-statistics interval on page 749
- clear services application-identification application-statistics cumulative on page 748

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear services application-identification application-statistics cumulative

Syntax clear services application-identification application-statistics cumulative

Release Information Statement introduced in Junos OS Release 11.4.

Description Clear all Junos OS application cumulative statistics.

Required Privilege Level clear

Related Documentation

- show services application-identification statistics applications on page 821
- show services application-identification statistics application-groups on page 823
- clear services application-identification application-statistics on page 747
- clear services application-identification application-statistics interval on page 749

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear services application-identification application-statistics interval

Syntax clear services application-identification application-statistics interval

Release Information Statement introduced in Junos OS Release 11.4.

Description Clear all Junos OS application interval statistics.

Required Privilege Level clear

Related Documentation

- show services application-identification statistics applications on page 821
- show services application-identification statistics application-groups on page 823
- clear services application-identification application-statistics on page 747
- clear services application-identification application-statistics cumulative on page 748

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear services application-identification application-system-cache (Junos OS)

Syntax clear services application-identification application-system-cache
<node (*node-id* | all | local | primary) >

Release Information Command introduced in Junos OS Release 10.2. Command syntax updated in Junos OS Release 12.1.

Description Clear Junos OS application identification application system cache.

Options

- none—Clear the application system cache on the device.
- **node**—(Optional) For chassis cluster configurations, clear application system cache on the specified nodes.
  - ***node-id***—Specific node number
  - **all**—All nodes
  - **local**—Local node
  - **primary**—Primary node

Required Privilege Level clear

Related Documentation

- show services application-identification application-system-cache (View) on page 813

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear services application-identification counter (Values)

Syntax clear services application-identification counter <ssl-encrypted-sessions>

Release Information Command introduced in Junos OS Release 10.2. Command updated in Junos OS Release 12.1-X47-D15.

Description Reset all the Junos OS application identification counter values.

Options **ssl-encrypted-sessions**—Reset application identification counter values for SSL encrypted sessions.

Required Privilege Level clear

Related Documentation
- show services application-identification counter (AppSecure) on page 816

List of Sample Output clear services application-identification counter on page 751

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services application-identification counter

```
user@host> clear services application-identification counter
clear_counter_class: counters cleared, status = 0
```

clear services ssl proxy statistics

Syntax **clear services ssl proxy statistics**

Release Information Command introduced in Junos OS Release 12.1.

Description Clear services SSL proxy statistics.

Options **none**—Clear the ssl proxy statistics.

Required Privilege Level clear

Related Documentation
- show services ssl proxy statistics on page 829

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate ca-profile-group load

Syntax request security pki ca-certificate ca-profile-group load ca-group-name *ca-group-name* filename [*path/filename* | default]

Release Information Command introduced in Junos OS Release 12.1; **default** option added in Junos OS Release 12.1X47-D10.

Description For SSL forward proxy, you need to load trusted CA certificates on your system. By default, Junos OS provides a list of trusted CA certificates that include default certificates used by common browsers. Alternatively, you can define your own list of trusted CA certificates and import them on to your system.

Use this command to load the default certificates or to specify a path and filename of trusted CA certificates that you define.

Options **ca-group-name *ca-group-name***—Load the specified CA group profile.

filename *path/filename*—Directory location and filename of the trusted CA certificates defined by you.

filename default—Load the trusted CA certificates available by default.

Required Privilege Level maintenance

Related Documentation
- show security pki ca-certificate on page 793
- Understanding Certificates and PKI on page 6347

List of Sample Output
request security pki ca-certificate ca-profile-group load (default) on page 753
request security pki ca-certificate ca-profile-group load (path/filename) on page 754

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate ca-profile-group load (default)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-default
filename default

Do you want to load this CA certificate ? [yes,no] (no) yes
Loading 157 certificates for group 'ca-default'.
ca-default_1: Loading done.
ca-default_2: Loading done.
ca-default_3: Loading done.
.....
```

Sample Output

request security pki ca-certificate ca-profile-group load (path/filename)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-manual
filename /var/tmp/firefox-all.pem

Do you want to load this CA certificate ? [yes,no] (no) yes

Loading 196 certificates for group 'ca-manual'.
ca-manual_1_sysgen: Loading done.
ca-manual_2_sysgen: Loading done.
ca-manual_3_sysgen: Loading done.
ca-manual_4_sysgen: Loading done.
ca-manual_5_sysgen: Loading done.
ca-manual_6_sysgen: Loading done.

...
ca-manual_195_sysgen: Loading done.
ca-manual_196_sysgen: Loading done.
ca-profile-group 'ca-manual' successfully loaded. Success[193] Skipped[3]
```


Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate export

Syntax request security pki local-certificate export

Release Information Command introduced in Junos OS Release 12.1.

Description Export a generated self-signed certificate from the default location (var/db/certs/common/local) to a specific location within the device.

Options

certificate id *certificate-id-name*—Name of the local digital certificate.

filename *path/filename*—Target directory location and filename of the CA digital certificate.

type (der | pem)—Certificate format: DER (distinguished encoding rules) or PEM (privacy-enhanced mail).

Required Privilege Level maintenance

Related Documentation

- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki local-certificate export on page 755

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate export

```
user@host> request security pki local-certificate export filename /var/tmp/my-cert.pem
certificate-id nss-cert type pem
certificate exported successfully
```

Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate generate-self-signed

Syntax request security pki local-certificate generate-self-signed certificate-id *certificate-id-name* domain-name *domain-name* ip-address *ip-address* email *email-address* subject *subject-distinguished-name*

Release Information Command introduced in Junos OS Release 9.1.

Description Manually generate a self-signed certificate for the given distinguished name.

Options

certificate-id ***certificate-id-name***—Name of the local digital certificate and the public/private key pair.

domain-name ***domain-name***—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

email ***email-address***—E-mail address of the certificate holder.

ip-address ***ip-address***—IP address of the router.

subject ***subject-distinguished-name***—Distinguished name format that contains the common name, department, company name, state, and country:

- **CN**—Common name
- **OU**—Organizational unit name
- **O**—Organization name
- **ST**—State
- **C**—Country

Required Privilege Level maintenance
security

Related Documentation

- show security pki local-certificate (View) on page 797

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name example.net email user@example.net
Self-signed certificate generated and loaded successfully
```


Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate load

Syntax request security pki local-certificate load certificate-id *certificate-id-name* filename *path*

Release Information Command introduced in Junos OS Release 7.5.

Description Manually load a local digital certificate from a specified location.

Options

certificate-id *certificate-id-name*—Name of the public/private key pair mapped to the local digital certificate.

filename *path/filename*—Directory location and filename of the local digital certificate provided by the CA.

Required Privilege Level maintenance

List of Sample Output request security pki local-certificate load on page 757

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate load

```
user@host> request security pki local-certificate load filename /tmp/router2-cert certificate-id
local-entrust2
Local certificate local-entrust2 loaded successfully
```

Copyright © 2016, Juniper Networks, Inc.

request services application-identification application

Syntax

request services application-identification application copy *predefined-application-name*

request services application-identification application [disable | enable]
predefined-application-name <no-commit>

Release Information

Command introduced in Junos OS Release 11.4.

Description

Copy, disable, or enable a predefined application signature.

Options

copy —(Optional) Copy a predefined application signature from the database to the configuration and change the name (for example, my:FTP). The ID and order will be generated automatically. Do not name your custom application signature with the "junos" prefix; this prefix is reserved for predefined application signatures. You can copy the same predefined application signature only once; duplicate custom signatures are not allowed. The **copy** command does not initiate signature recompilation.

NOTE: In configuration mode, if an uncommitted action is pending, the **request services application-identification application copy** command will fail. Uncompiled application signatures are shown as uncommitted in the **show services application-identification application [summary | detail]** command.

disable—(Optional) Disable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration. Include the <no-commit> keyword to defer signature recompilation.

The following conditions apply:

- You cannot disable a predefined application signature that is referenced by an active security policy or custom application signature. First modify or deactivate the policy or custom application signature.
- If you disable an application signature, for example, junos:HTTP, that has nested applications, the nested applications will not be recognized.

enable—(Optional) Enable a predefined application signature, initiate signature recompilation, and commit all pending uncompiled signatures to the configuration. Include the <no-commit> keyword to defer signature recompilation.

no-commit—(Optional) Enables you to enter multiple **enable** and **disable** commands before initiating signature recompilation (which takes some time) and committing the configuration. Uncompiled application signatures are shown as uncommitted in the **show services application-identification application [summary | detail]** command.

Required Privilege Level

maintenance

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- show services application-identification application on page 810

Output Fields

When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification application copy

```
user@host> request services application-identification application junos:FTP copy
application package 63 copied successfully.
```

request services application-identification application disable

```
user@host> request services application-identification application disable junos:163
 Please wait while we are re-compiling signatures ...
 Please wait while we are re-compiling signatures ...
 Please wait while we are re-compiling signatures ...
 Please wait while we are re-compiling signatures ...
 Disable application junos:163 succeed.
```

request services application-identification application disable no-commit

```
user@host> request services application-identification application disable
junos:FACEBOOK-SOCIALRSS no-commit
Disable application junos:FACEBOOK-SOCIALRSS succeed. It is not committed yet.
```

Copyright © 2016, Juniper Networks, Inc.

request services application-identification download

Syntax request services application-identification download <version>;

Release Information Statement introduced in Junos OS Release 10.2. Statement modified in Junos OS Release 11.4.

Description Manually download the application package for Junos OS application identification. The application package is extracted from the IDP signature database and contains signature definitions for known applications, such as: DNS, Facebook, FTP, Skype, and SNMP.

Options **version**—(Optional) Download a specific version of the application package from the Juniper security website. If you do not enter a version, the most recent version will be downloaded.

Required Privilege Level maintenance

Related Documentation
- request services application-identification download status on page 761
- request services application-identification install on page 764

List of Sample Output request services application-identification download on page 760

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services application-identification download

```
user@host> request services application-identifications download
Please use command "request services application-identification download status"
 to check status
```


Copyright © 2016, Juniper Networks, Inc.

request services application-identification download status

Syntax request services application-identification download status

Release Information Statement introduced in Junos OS Release 10.2. Statement modified in Junos OS Release 11.4.

Description Check the download status of the application signature package. The downloaded application package is saved under /var/db/appid/sec-download/.

Required Privilege Level maintenance

Related Documentation
- request services application-identification download on page 760

List of Sample Output request services application-identification download status on page 761

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification download status

```
user@host> request services application-identifications download status
Application package 1608 is downloaded successfully.
```

Copyright © 2016, Juniper Networks, Inc.

request services application-identification group

Syntax request services application-identification group [copy |disable | enable]
predefined-application-group-name

Release Information Command introduced in Junos OS Release 11.4.

Description Copy, disable, or enable a predefined application group.

Options **copy** —(Optional) Copy a predefined application signature group from the database to the configuration and change the name (for example, my:FTP). The ID and order will be generated automatically. Do not name your custom application signature group with the "junos" prefix; this prefix is reserved for predefined application signature groups. You can copy the same predefined application signature group only once; duplicate custom signature groups are not allowed.

NOTE: In configuration mode, if an uncommitted action is pending, the **request services application-identification group copy** command will fail.

disable—(Optional) Disable a predefined application signature group.

NOTE: You cannot disable a predefined application signature group that is referenced by an active security policy or custom application signature group. First modify or deactivate the policy or custom application signature group.

enable—(Optional) Enable a predefined application signature group.

Required Privilege Level maintenance

Related Documentation
- show services application-identification group on page 819

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

request services application-identification group

```
user@host> request services application-identification group disable
junos:infrastructure:networking
Disable application group junos:infrastructure:networking succeed.
```


Copyright © 2016, Juniper Networks, Inc.

request services application-identification group

```
user@host> request services application-identification group enable
junos:infrastructure:networking
Enable application group junos:infrastructure:networking succeed.
```

request services application-identification group

```
user@host> request services application-identification group copy junos:infrastructure:networking
Please wait while we are copying group ...
Copy application group junos:infrastructure:networking succeed.
```

Copyright © 2016, Juniper Networks, Inc.

request services application-identification install

Syntax request services application-identification install

Release Information Statement introduced in Junos OS Release 11.4.

Description Install the downloaded predefined application signature package.

NOTE: You can also use the **request security idp security-package [download | install]** command to download and install the predefined application signature package along with the IDP security package.

Required Privilege Level maintenance

Related Documentation

- request services application-identification install status on page 765
- request services application-identification download on page 760

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
user@host> request services application-identification install
Please use command "request services application-identification install status"
to check status and use command "request services application-identification
proto-bundle-status" to check protocol bundle status
```


Copyright © 2016, Juniper Networks, Inc.

request services application-identification install status

Syntax request services application-identification install status

Release Information Statement introduced in Junos OS Release 11.4.

Description Display the status of the install operation.

Required Privilege Level maintenance

Related Documentation
- request services application-identification install on page 764

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
user@host> request services application-identification install status
Install application package version (1776) succeed.
```

request services application-identification proto-bundle-status

Syntax request services application-identification proto-bundle-status

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Display the status of the install operation of the protocol bundle.

Required Privilege Level maintenance

Related Documentation
- request services application-identification install on page 764

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
user@host> request services application-identification proto-bundle-status
Protocol Bundle Version (1.30.4-22.005 (build date Jan 17 2014)) and application
 secpack version (2345) is loaded and activated.
```


Copyright © 2016, Juniper Networks, Inc.

request services application-identification uninstall

Syntax request services application-identification uninstall

Release Information Statement introduced in Junos OS Release 10.2. Statement modified in Junos OS Release 10.4. Statement modified in Junos OS Release 11.4.

Description Uninstall the predefined application package.

The uninstall operation will fail if any active security policies reference predefined application signatures or predefined application signature groups in the Junos OS configuration.

Required Privilege Level maintenance

Related Documentation

- request services application-identification install on page 764

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
user@host> request services application-identification uninstall
Please use command "request services application-identification uninstall status"
 to check status and use command "request services application-identification
proto-bundle-status" to check protocol bundle status
```

Copyright © 2016, Juniper Networks, Inc.

request services application-identification uninstall status

Syntax request services application-identification uninstall status

Release Information Statement introduced in Junos OS Release 11.4.

Description Display the status of the uninstall operation.

Required Privilege Level maintenance

Related Documentation
- request services application-identification uninstall on page 767

Output Fields When you enter this command, the system provides feedback on the status of your request.

Sample Output

```
user@host> request services application-identification uninstall status
Uninstall application package version (1776) succeed.
```


Copyright © 2016, Juniper Networks, Inc.

show class-of-service application-traffic-control counter

Syntax show class-of-service application-traffic-control counter

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS DSCP marking and honoring statistics based on Layer 7 application classifiers.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control counter on page 769

Output Fields Table 41 lists the output fields for the **show class-of-service application-traffic-control counter** command. Output fields are listed in the approximate order in which they appear.

Table 41: show class-of-service application-traffic-control counter Output Fields

Field Name	Field Description
pic	PIC number of the accumulated statistics. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Sessions processed	The number of sessions where the class of service was checked.
Sessions marked	The number of sessions marked based on application-aware DSCP marking.
Sessions honored	The number of sessions honored based on application-aware traffic honoring.
Sessions rate limited	The number of sessions that have been rate limited.
Client-to-server flows rate limited	The number of client-to-server flows that have been rate limited.
Server-to-client flows rate limited	The number of server-to-client flows that have been rate limited.

Sample Output

show class-of-service application-traffic-control counter

```
user@host> show class-of-service application-traffic-control counter
pic: 2/1
  Counter type                                  Value
  Sessions processed                            300
  Sessions marked                               200
  Sessions honored                              0
  Sessions rate limited                         100
  Client-to-server flows rate limited           100
  Server-to-client flows rate limited           70
```

Copyright © 2016, Juniper Networks, Inc.

```
pic: 2/0
  Counter type                                    Value
  Sessions processed                              400
  Sessions marked                                 300
  Sessions honored                                0
  Sessions rate limited                           200
  Client-to-server flows rate limited             200
  Server-to-client flows rate limited             100
```


Copyright © 2016, Juniper Networks, Inc.

show class-of-service application-traffic-control statistics rate-limiter

Syntax show class-of-service application-traffic-control statistics rate-limiter

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS real-time run information about application rate limiting of current or recent sessions.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control statistics rate-limiter on page 771

Output Fields Table 42 lists the output fields for the **show class-of-service application-traffic-control statistics rate-limiter** command. Output fields are listed in the approximate order in which they appear.

Table 42: show class-of-service application-traffic-control statistics rate-limiter Output Fields

Field Name	Field Description
pic	PIC number. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Ruleset	The rule set applied on the session.
Application	The application match for applying the rule set.
Client-to-server	The rate limiter applied from client to server.
Rate(kbps)	The rate in the client-to-server direction
Server-to-client	The rate limiter applied from server to client.
Rate(kbps)	The rate in the server-to-client direction.

Sample Output

show class-of-service application-traffic-control statistics rate-limiter

```
user@host> show class-of-service application-traffic-control statistics rate-limiter
pic: 2/1
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client
Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000     my-http-s2c-rl
20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000     my-http-s2c-rl-2
30000000
```

Copyright © 2016, Juniper Networks, Inc.

```
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl
50000
  ...

pic: 2/0
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client
Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000    my-http-s2c-rl
20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000    my-http-s2c-rl-2
30000000
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl
50000
```


Copyright © 2016, Juniper Networks, Inc.

show class-of-service application-traffic-control statistics rule

Syntax show class-of-service application-traffic-control statistics rule

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS counters identifying rule hits.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control statistics rule on page 773

Output Fields Table 43 lists the output fields for the **show class-of-service application-traffic-control statistics rule** command. Output fields are listed in the approximate order in which they appear.

Table 43: show class-of-service application-traffic-control statistics rule Output Fields

Field Name	Field Description
pic	PIC number where the rule is applied. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Ruleset	The rule set containing the rule.
Rule	The rule to which the statistic applies.
Hits	The number of times a match for the rule was encountered.

Sample Output

show class-of-service application-traffic-control statistics rule

```
user@host> show class-of-service application-traffic-control statistics rule
pic: 2/0
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       100
  my-ruleset-1     httpp-rule     100
  my-ruleset-2     telnet-rule    300
  my-ruleset-2     smtp-rule      300
  ...

pic: 2/1
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       200
  my-ruleset-1     httpp-rule     300
  my-ruleset-2     telnet-rule    400
  my-ruleset-2     smtp-rule      500
```

Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set

Syntax show security application-firewall rule-set (*<rule-set-name>* | all)

Release Information Command introduced in Junos OS Release 11.1. Updated in Junos OS Release 12.1X44-D10 with output format changes. Updated in Junos OS Release 12.1X45-D10 with redirection counters.

Description Display information about the specified rule set defined in the application firewall.

Options ***rule-set-name***—Name of the rule set.

all—Display information about all the application firewall rule sets.

Required Privilege Level view

Related Documentation
- clear security application-firewall rule-set statistics on page 745

List of Sample Output
show security application-firewall rule-set my_ruleset1 on page 775
show security application-firewall rule-set all on page 775

Output Fields Table 44 lists the output fields for the **show security application-firewall rule-set** command. Output fields are listed in the approximate order in which they appear.

Table 44: show security application-firewall rule-set Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system of the rule set.
Profile	The redirect profile to be used for rules requiring redirection for reject or deny actions.

Rule	Name of the rule
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **SSL-Encryption**—Setting for SSL traffic.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	  ◦ **reject**
	  ◦ **redirect**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.
	• **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Copyright © 2016, Juniper Networks, Inc.

Table 44: show security application-firewall rule-set Output Fields *(continued)*

Field Name	Field Description
Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule. • **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.
Number of sessions with appid pending	Number of sessions that are pending application identification processing

Sample Output

show security application-firewall rule-set my_ruleset1

```
user@host>show security application-firewall rule-set my_ruleset1
Rule-set: my_ruleset1
     Rule: rule1
         Dynamic Applications: junos:facebook, junos:messenger
         Dynamic Application Groups: junos:web, junos:chat
         SSL-Encryption: any
         Action: deny or redirect
         Number of sessions matched: 10
         Number of sessions redirected: 10
     Default rule: permit
         Number of sessions matched: 200
         Number of sessions redirected: 0
     Number of sessions with appid pending: 2
```

Sample Output

show security application-firewall rule-set all

```
user@host> show security application-firewall rule-set all
Rule-set: appfw
    Logical system: root-logical-system
    Profile: lsy2_pf555
    Rule: 2
        Dynamic Applications: junos:HTTP
        SSL-Encryption: any
        Action:deny or redirect
        Number of sessions matched: 2
        Number of sessions redirected: 2
    Rule: 1
        Dynamic Applications: junos:FTP
        SSL-Encryption: any
        Action:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Default rule:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Number of sessions with appid pending: 0
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set logical-system

Syntax The master, or root, administrator can issue the following statements:

show security application-firewall rule-set all
show security application-firewall rule-set *rule-set-name* | all | logical-system *logical-system-name* | all | root-logical-system [*logical-system-name* | all]

The user logical system administrator can issue the following statement:

show security application-firewall rule-set all

Release Information Command introduced in Junos OS Release 11.4.

Description Display information about application firewall rule set(s) associated with a specific logical system, all logical systems, or the root logical system configured on a device.

NOTE: The master administrator can configure and view application firewall rule sets for the root logical system and all user logical systems configured on the device. User logical system administrators can configure and view application firewall rule set information only for the user logical systems for which they have access. For information about master and user administrator roles in logical systems, see "Understanding Logical Systems for SRX Series Services Gateways" on page 3205.

Options ***rule-set-name***—Name of a specific rule set.

logical-system-name—Name of a specific logical system.

all—(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.

root-logical-system—Display application firewall rule set information for the root logical system (master administrator only).

Required Privilege Level view

Related Documentation

- clear security application-firewall rule-set statistics logical-system on page 746

List of Sample Output

show security application-firewall rule-set logical-system all on page 778
show security application-firewall rule-set all on page 779

Output Fields Table 45 lists the output fields for the **show security application-firewall rule-set logical-system** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 45: show security application-firewall rule-set logical-system Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system.

Rule	Name of the rule.
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.

Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule.
Number of sessions with appid pending	Number of sessions that are pending with the application ID processing.

Sample Output

show security application-firewall rule-set logical-system all

```
root@host> show security application-firewall rule-set logical-system all

Rule-set: root_rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FTP
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 10
```


Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set all

```
root@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Copyright © 2016, Juniper Networks, Inc.

show security application-tracking counters

Syntax show security application-tracking counters

Release Information Command introduced in Junos OS Release 10.2.

Description Display the status of AppTrack counters.

Required Privilege Level view

Related Documentation
- Example: Configuring AppTrack on page 574

Output Fields Table 46 lists the output fields for the **show security application-tracking counters** command. Output fields are listed in the approximate order in which they appear.

Table 46: show security application-tracking counters

Field Name	Field Description
Session create messages	The number of log messages generated when a session was created.
Session close messages	The number of log messages generated when a session was closed.
Session volume updates	The number of log messages generated when an update interval was exceeded.
Failed messages	The number of messages that were not generated due to memory or session constraints.

Sample Output

show security application-tracking counters

```
user@host> show security application-tracking counters
AVT counters:                       Value
 Session create messages              0
 Session close messages               0
 Session volume updates               0
 Failed messages                      0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session

Syntax **show security flow session**
[*filter*] **[brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 8.5. Support for filter and view options added in Junos OS Release 10.2. Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2. Policy ID filter added in Junos OS Release 12.3X48-D10. Support for connection tag added in Junos OS Release 15.1X49-D40.

Description Display information about all currently active security sessions on the device.

NOTE: For the normal flow sessions, the **show security flow session** command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. Because of this, the output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific **show** command for examples of the filtered output.

  application—Predefined application name

  application-firewall—Application firewall enabled

  application-firewall-rule-set—Application firewall enabled with the specified rule set

  application-traffic-control—Application traffic control session

  application-traffic-control-rule-set—Application traffic control rule set name and rule name

  conn-tag—Session connection tag

  destination-port—Destination port

  destination-prefix—Destination IP prefix or address

  dynamic-application—Dynamic application

  dynamic-application-group—Dynamic application

  encrypted—Encrypted traffic

  extensive—Display detailed output

Copyright © 2016, Juniper Networks, Inc.

family—Display session by family

idp—IDP enabled sessions

interface—Name of incoming or outgoing interface

logical-system (all | *logical-system-name*)—Name of a specific logical system or **all** to display all logical systems

nat—Display sessions with network address translation

policy-id—Display session information based on policy ID; the range is 1 through 4,294,967,295

protocol—IP protocol number

resource-manager—Resource manager

root-logical-system—Display root logical system as default

security-intelligence—Display security intelligence sessions

services-offload—Display services offload sessions

session-identifier—Display session with specified session identifier

source-port—Source port

source-prefix—Source IP prefix

summary—Display output summary

tunnel—Tunnel sessions

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session all on page 1902

List of Sample Output
show security flow session on page 784
show security flow session brief on page 785
show security flow session extensive on page 785
show security flow session summary on page 785

Output Fields Table 26 lists the output fields for the **show security flow session** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 47: show security flow session Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Application traffic control rule-set	AppQoS rule set for this session.
Rule	AppQoS rule for this session.
Forwarding class	The AppQoS forwarding class name for this session that distinguishes the transmission priority
DSCP code point	Differentiated Services (DiffServ) code point (DSCP) value remarked by the matching rule for this session.

Copyright © 2016, Juniper Networks, Inc.

Table 47: show security flow session Output Fields *(continued)*

Field Name	Field Description
Loss priority	One of four priority levels set by the matching rule to control discarding a packet during periods of congestion. A high loss priority means a high probability that the packet could be dropped during a period of congestion.
Rate limiter client to server	The rate-limiter profile assigned to the client-to-server traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Rate limiter server to client	The rate-limiter profile assigned to the server-to-client traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Maximum number of sessions permitted.

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
 Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/10000/2000 --> 203.0.113.11/10000/1000;udp, Conn Tag: 0x0, If:
 reth0.0, Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session application-firewall

Syntax

show security flow session application-firewall
< dynamic-application (*dyn-app-name* | junos:UNKNOWN) >
< dynamic-application-group (*dyn-app-group* | junos:UNASSIGNED) >
< application-firewall-rule-set *rule-set-name* >
< rule *rule-name* >
< brief | extensive | summary >

Release Information

Command introduced in Junos OS Release 11.2.

Description

Display all sessions where application firewall is enabled.

Include options to filter the output and display only those enabled sessions with the specified features.

Options

- **dynamic-application (*dyn-app-name* | junos:UNKNOWN)**—Display only those enabled sessions with the specified dynamic application. Enter **junos:UNKNOWN** to display all enabled sessions where no dynamic application can be determined.
- **dynamic-application-group (*dyn-app-group* | junos:UNASSIGNED)**— Display only those enabled session with the specified dynamic application group. Enter **junos:UNASSIGNED** to display all enabled sessions where no dynamic application group can be determined.
- **application-firewall-rule-set *rule-set-name***—Display only those enabled sessions that match the specified rule set.
- **rule *rule-name***—Display only those enabled sessions that match the specified rule.
- **brief | extensive | summary**—Specify the level of detail for the display.

  The output fields for the **brief** and **summary** options are the same as those of the **show security flow session** command. Only the **extensive** display is different and is shown in the following output table and examples.

Required Privilege Level

view

Related Documentation

- Example: Configuring an Application Group for Application Firewall on page 564
- show security flow session on page 453

List of Sample Output

show security flow session application-firewall extensive on page 789
show security flow session application-firewall dynamic-application junos:FTP extensive on page 789
show security flow session application-firewall dynamic-application junos:UNKNOWN extensive on page 790
show security flow session application-firewall dynamic-application-group junos:WEB extensive on page 791
show security flow session application-firewall application-firewall-rule-set rule-set1 extensive on page 791

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 48 lists the output fields for the **show security flow session application-firewall extensive** command. Output fields are listed in the approximate order in which they appear in the extensive display.

Table 48: show security flow session application-firewall extensive Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to display more information about a session.
Status	Session status.
State	Current state of the session: Active, Pending, Closed, Unknown.
Flag	Internal flag depicting the state of the session. It is used for debugging purposes.
Policy name	The name of the policy that permitted the traffic.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the dynamic application of the session. If the dynamic application has yet to be determined, the output indicates Pending. If the dynamic application cannot be determined, the output indicates junos:UNKNOWN.
Dynamic application group	Name of the dynamic application group of the session. If the dynamic application cannot be determined, the output indicates junos:UNASSIGNED.
Dynamic nested application	Name of the dynamic nested application of the session if one exists. If the dynamic nested application is yet to be determined, the output indicates Pending. If the dynamic nested application cannot be determined, the output indicates junos:UNKNOWN.
Application firewall rule-set	Name of the rule set that the session matched.
Rule	Name of the rule that the session matched. If the match has not yet been made, the output indicates Pending. If the rule has been deleted since the match was made, the output indicates the rule is invalid.
Maximum timeout	Maximum amount of idle time allowed for the session.
Current timeout	Number of seconds that the current session has been idle.
Session State	Session state.
Start time	Time when the session was created. Start time is indicated as an offset from the system start time.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).

Copyright © 2016, Juniper Networks, Inc.

Table 48: show security flow session application-firewall extensive Output Fields *(continued)*

Field Name	Field Description
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions per PIC that fit the display criteria.

Sample Output

show security flow session application-firewall extensive

The displayed information is similar to the **show security flow session** output but includes dynamic application and application firewall details for the session.

```
user@host> show security flow session application-firewall extensive
Flow Sessions on FPC9 PIC0:

     Session ID: 3729, Status: Normal, State: Active
      Policy name: self-traffic-policy/1
      Source NAT pool: Null
     Dynamic application: junos:HTTP, Dynamic nested application: junos:FACEBOOK

      Application firewall rule-set: rule-set1, Rule: rule2
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

     Total sessions: 1
```

show security flow session application-firewall dynamic-application junos:FTP extensive

Entering a specific dynamic application in the command line filters the output and displays only those sessions with the specified application.

```
user@host> show security flow session application-firewall dynamic-application junos:FTP
extensive
Flow Sessions on FPC3 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:FTP
      Application firewall rule-set: rule-set1, Rule: rule1
```

Copyright © 2016, Juniper Networks, Inc.

```
Maximum timeout: 300, Current timeout: 276
Session State: Valid
Start time: 18292, Duration: 603536
   In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
    Interface: reth1.0,
    Session token: 0x1c0, Flag: 0x0x21
    Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 21043, Bytes: 1136322
   Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
    Interface: .local..0,
    Session token: 0x80, Flag: 0x0x30
    Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0

Total sessions: 1
```

show security flow session application-firewall dynamic-application junos:UNKNOWN extensive

Using the keyword **junos:UNKNOWN** displays those enabled sessions where the dynamic application cannot be determined.

```
user@host> show security flow session application-firewall dynamic-application junos:UNKNOWN
extensive
Flow Sessions on FPC9 PIC0:

     Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
      Dynamic application: junos:UNKNOWN
      Application firewall rule-set: rule-set1, Rule:rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 21043, Bytes: 1136322
         Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
          Interface: .local..0,
          Session token: 0x80, Flag: 0x0x30
          Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
          Port sequence: 0, FIN sequence: 0,
          FIN state: 0,
          Pkts: 0, Bytes: 0

      Session ID: 180013339, Policy name: policy1/4, Timeout: 1776, Valid
     Dynamic application: junos:HTTP, Dynamic nested application: junos:UNKNOWN

      Application firewall rule-set: rule-set1, Rule:rule1
      Maximum timeout: 300, Current timeout: 276
      Session State: Valid
      Start time: 18292, Duration: 603536
         In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
          Interface: reth1.0,
          Session token: 0x1c0, Flag: 0x0x21
          Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 21043, Bytes: 1136322
   Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
    Interface: .local..0,
    Session token: 0x80, Flag: 0x0x30
    Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0

Total sessions: 2
```

show security flow session application-firewall dynamic-application-group junos:WEB extensive

Entering a specific dynamic application group in the command line filters the output and displays only those sessions with the specified application group.

```
user@host> show security flow session application-firewall dynamic-application-group junos:WEB
extensive
Flow Sessions on FPC9 PIC0:

Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
 Dynamic application: junos:HOTMAIL
 Application firewall rule-set: rule-set1, Rule: rule1
 Maximum timeout: 300, Current timeout: 276
 Session State: Valid
 Start time: 18292, Duration: 603536
    In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
     Interface: reth1.0,
     Session token: 0x1c0, Flag: 0x0x21
     Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 21043, Bytes: 1136322
    Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
     Interface: .local..0,
     Session token: 0x80, Flag: 0x0x30
     Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 0, Bytes: 0

 Total sessions: 1
```

show security flow session application-firewall application-firewall-rule-set rule-set1 extensive

Specifying a rule set name reduces the display to only those sessions matching the specified rule set.

```
user@host> show security flow session application-firewall application-firewall-rule-set rule-set1
extensive
Flow Sessions on FPC9 PIC0:

Session ID: 180013338, Policy name: policy1/4, Timeout: 1776, Valid
 Dynamic application: junos:FTP
 Application firewall rule-set: rule-set1, Rule: rule1
 Maximum timeout: 300, Current timeout: 276
 Session State: Valid
 Start time: 18292, Duration: 603536
```

```
    In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
     Interface: reth1.0,
     Session token: 0x1c0, Flag: 0x0x21
     Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 21043, Bytes: 1136322
    Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
     Interface: .local..0,
     Session token: 0x80, Flag: 0x0x30
     Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 0, Bytes: 0

Session ID: 180013339, Policy name: policy1/4, Timeout: 1776, Valid
Dynamic application: junos:HTTP, Dynamic nested application: junos:FACEBOOK

 Application firewall rule-set: rule-set1, Rule: rule2
 Maximum timeout: 300, Current timeout: 276
 Session State: Valid
 Start time: 18292, Duration: 603536
    In: 201.34.0.4/1 --> 224.0.0.13/1;pim,
     Interface: reth1.0,
     Session token: 0x1c0, Flag: 0x0x21
     Route: 0x0, Gateway: 201.34.0.4, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 21043, Bytes: 1136322
    Out: 224.0.0.13/1 --> 201.34.0.4/1;pim,
     Interface: .local..0,
     Session token: 0x80, Flag: 0x0x30
     Route: 0xfffd0000, Gateway: 224.0.0.13, Tunnel: 0
     Port sequence: 0, FIN sequence: 0,
     FIN state: 0,
     Pkts: 0, Bytes: 0


 Total sessions: 2
```


Copyright © 2016, Juniper Networks, Inc.

show security pki ca-certificate

Syntax show security pki ca-certificate
<brief | detail>
<ca-profile *ca-profile-name*>

Release Information Command introduced in Junos OS Release 7.5.

Description Display information about certificate authority (CA) digital certificates installed in the router.

Options **none**—(Same as brief) Display information about all CA digital certificates.

brief | detail—(Optional) Display the specified level of output.

ca-profile *ca-profile-name*—(Optional) Display information about only the specified CA profile.

Required Privilege Level view

List of Sample Output show security pki ca-certificate on page 794
show security pki ca-certificate detail on page 795

Output Fields Table 49 lists the output fields for the **show security pki ca-certificate** command. Output fields are listed in the approximate order in which they appear.

Table 49: show security pki ca-certificate Output Fields

Field Name	Field Description	Level of Output
Certificate identifier	Name of the digital certificate.	All levels
Certificate version	Revision number of the digital certificate.	**detail**
Serial number	Unique serial number of the digital certificate.	**detail**
Issued by	Authority that issued the digital certificate.	**none brief**
Issued to	Device that was issued the digital certificate.	**none brief**
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the authority. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 49: show security pki ca-certificate Output Fields *(continued)*

Field Name	Field Description	Level of Output
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Common name**—Name of the requestor. • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **State**—State of origin. • **Country**—Country of origin.	**detail**
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.	All levels
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption(1024 bits)**.	All levels
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.	**detail**
Fingerprint	Secure Hash Algorithm (SHA1) and Message Digest 5 (MD5) hashes used to identify the digital certificate.	**detail**
Distribution CRL	Distinguished name information and the URL for the certificate revocation list (CRL) server.	**detail**
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Key encipherment**.	**detail**

Sample Output

show security pki ca-certificate

```
user@host> show security pki ca-certificate
Certificate identifier: entrust
  Issued to: example, Issued by: example
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier: entrust
  Issued to: First Officer, Issued by:exampler
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)

Certificate identifier: entrust
  Issued to: First Officer, Issued by:exampler
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
```


Copyright © 2016, Juniper Networks, Inc.

```
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)
```

show security pki ca-certificate detail

```
user@host> show security pki ca-certificate detail
Certificate identifier: entrust
  Certificate version: 3
  Serial number: 4355 9235
  Issuer:
    Organization: example-org, Country: us
  Subject:
    Organization: example-org, Country: us
  Validity:
    Not before: 2005 Oct 18th, 23:54:22 GMT
    Not after: 2025 Oct 19th, 00:24:22 GMT
  Public key algorithm: rsaEncryption(1024 bits)
    cb:9e:2d:c0:70:f8:ea:3c:f2:b5:f0:02:48:87:dc:68:99:a3:57:4f
    0e:b9:98:0b:95:47:0d:1f:97:7c:53:17:dd:1a:f8:da:e5:08:d1:1c
    78:68:1f:2f:72:9f:a2:cf:81:e3:ce:c5:56:89:ce:f0:97:93:fa:36
    19:3e:18:7d:8c:9d:21:fe:1f:c3:87:8d:b3:5d:f3:03:66:9d:16:a7
    bf:18:3f:f0:7a:80:f0:62:50:43:83:4f:0e:d7:c6:42:48:c0:8a:b2
    c7:46:30:38:df:9b:dc:bc:b5:08:7a:f3:cd:64:db:2b:71:67:fe:d8
    04:47:08:07:de:17:23:13
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10 (sha1)
    71:6f:6a:76:17:9b:d6:2a:e7:5a:72:97:82:6d:26:86 (md5)
  Distribution CRL:
    C=us, O=example, CN=CRL1
    http://CA-1/CRL/example_us_crlfile.crl
  Use for key: CRL signing, Certificate signing
Certificate identifier: entrust
  Certificate version: 3
  Serial number: 4355 925c
  Issuer:
    Organization: example, Country: us
  Subject:
    Organization: example, Country: us, Common name: First Officer
  Validity:
    Not before: 2005 Oct 18th, 23:55:59 GMT
    Not after: 2008 Oct 19th, 00:25:59 GMT
  Public key algorithm: rsaEncryption(1024 bits)
    c0:a4:21:32:95:0a:cd:ec:12:03:d1:a2:89:71:8e:ce:4e:a6:f9:2f
    1a:9a:13:8c:f6:a0:3d:c9:bd:9d:c2:a0:41:77:99:1b:1e:ed:5b:80
    34:46:f8:5b:28:34:38:2e:91:7d:4e:ad:14:86:78:67:e7:02:1d:2e
    19:11:b7:fa:0d:ba:64:20:e1:28:4e:3e:bb:6e:64:dc:cd:b1:b4:7a
    ca:8f:47:dd:40:69:c2:35:95:ce:b8:85:56:d7:0f:2d:04:4d:5d:d8
    42:e1:4f:6b:bf:38:c0:45:1e:9e:f0:b4:7f:74:6f:e9:70:fd:4a:78
    da:eb:10:27:bd:46:34:33
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17 (sha1)
    23:79:40:c9:6d:a6:f0:ca:e0:13:30:d4:29:6f:86:79 (md5)
  Distribution CRL:
    C=us, O=example, CN=CRL1
    http://CA-1/CRL/example_us_crlfile.crl
  Use for key: Key encipherment
Certificate identifier: entrust
  Certificate version: 3
  Serial number: 4355 925b
```

Copyright © 2016, Juniper Networks, Inc.

```
Issuer:
  Organization: example, Country: us
Subject:
  Organization: example, Country: us, Common name: First Officer
Validity:
  Not before: 2005 Oct 18th, 23:55:59 GMT
  Not after: 2008 Oct 19th, 00:25:59 GMT
Public key algorithm: rsaEncryption(1024 bits)
  ea:75:c4:f3:58:08:ea:65:5c:7e:b3:de:63:0a:cf:cf:ec:9a:82:e2
  d7:e8:b9:2f:bd:4b:cd:86:2f:f1:dd:d8:a2:95:af:ab:51:a5:49:4e
  00:10:c6:25:ff:b5:49:6a:99:64:74:69:e5:8c:23:5b:b4:70:62:8e
  e4:f9:a2:28:d4:54:e2:0b:1f:50:a2:92:cf:6c:8f:ae:10:d4:69:3c
  90:e2:1f:04:ea:ac:05:9b:3a:93:74:d0:59:24:e9:d2:9d:c2:ef:22
  b9:32:c7:2c:29:4f:91:cb:5a:26:fe:1d:c0:36:dc:f4:9c:8b:f5:26
  af:44:bf:53:aa:d4:5f:67
Signature algorithm: sha1WithRSAEncryption
Fingerprint:
  46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f (sha1)
  ee:cc:c7:f4:5d:ac:65:33:0a:55:db:59:72:2c:dd:16 (md5)
Distribution CRL:
  C=us, O=example, CN=CRL1
  http://CA-1/CRL/example_us_crlfile.crl
Use for key: Digital signature
```


Copyright © 2016, Juniper Networks, Inc.

show security pki local-certificate (View)

Syntax

show security pki local-certificate
< **brief** | **detail** >
< certificate-id ***certificate-id-name*** >
<system-generated>

Release Information

Command modified in Junos OS Release 9.1. Subject string output field added in Junos OS Release 12.1X44-D10.

Description

Display information about the local digital certificates, corresponding public keys, and the automatically generated self-signed certificate configured on the device.

Options

- none—Display basic information about all configured local digital certificates, corresponding public keys, and the automatically generated self-signed certificate.
- **brief** | **detail**—(Optional) Display the specified level of output.
- certificate-id ***certificate-id-name*** —(Optional) Display information about only the specified local digital certificates and corresponding public keys.
- **system-generated**—Display information about the automatically generated self-signed certificate.

Required Privilege Level

view

Related Documentation

- clear security pki local-certificate (Device) on page 6990
- request security pki local-certificate generate-self-signed (Security) on page 7005

List of Sample Output

show security pki local-certificate certificate-id hello on page 799
show security pki local-certificate certificate-id hello detail on page 799
show security pki local-certificate system-generated on page 800
show security pki local-certificate system-generated detail on page 800
show security pki local-certificate certificate-id mycert - (local certificate enrolled online using SCEP) on page 800
show security pki local-certificate certificate-id mycert detail - (local certificate enrolled online using SCEP) on page 801

Output Fields

Table 50 lists the output fields for the **show security pki local-certificate** command. Output fields are listed in the approximate order in which they appear.

Table 50: show security pki local-certificate Output Fields

Field Name	Field Description
Certificate identifier	Name of the digital certificate.
Certificate version	Revision number of the digital certificate.
Serial number	Unique serial number of the digital certificate.

Copyright © 2016, Juniper Networks, Inc.

Table 50: show security pki local-certificate Output Fields *(continued)*

Field Name	Field Description
Issued to	Device that was issued the digital certificate.
Issued by	Authority that issued the digital certificate.
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority.
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority. • **Serial number**—Serial number of the device. If the certificate contains multiple subfield entries, all entries are displayed.
Subject string	Subject field as it appears in the certificate.
Alternate subject	Domain name or IP address of the device related to the digital certificate.
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption**(**1024 bits**).
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.
Fingerprint	Secure Hash Algorithm (**SHA1**) and Message Digest 5 (**MD5**) hashes used to identify the digital certificate.
Distribution CRL	Distinguished name information and URL for the certificate revocation list (**CRL**) server.
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Data encipherment**.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate certificate-id hello

```
user@host> show security pki local-certificate certificate-id hello
Certificate identifier: hello
  Issued to: cn1, Issued by: DC = local, DC = demo, CN = example-CN1
  Validity:
    Not before: 08- 8-2012 17:02
    Not after: 08- 8-2014 17:02
  Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki local-certificate certificate-id hello detail

```
user@host> show security pki local-certificate certificate-id hello detail
Certificate identifier: hello
  Certificate version: 3
  Serial number: 61ba9da000000000d72e
  Issuer:
    Common name: example-cn1,
    Domain component: local, Domain component: demo
  Subject:
    Organization: o1, Organization: o2,
    Organizational unit: ou1, Organizational unit: ou2, Country: US, State: CA,
    Locality: Sunnyvale, Common name: cn1, Common name: cn2,
    Domain component: dc1, Domain component: dc2
  Subject string:
    C=US, DC=dc1, DC=dc2, ST=CA, L=Sunnyvale, O=o1, O=o2, OU=ou1, OU=ou2, CN=cn1,
 CN=cn2
  Alternate subject: "user@example.net", user.example.net, 10.1.2.3
  Validity:
    Not before: 08- 8-2012 17:02
    Not after: 08- 8-2014 17:02
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:b4:14:01:d5:4f:79:87:d5:bb:e6:5e:c1:14
    97:da:b4:40:ad:1a:77:3e:ec:2e:68:8e:e4:93:a3:fe:7c:0b:58:af
    e1:20:27:82:ca:8d:6f:f0:97:d1:ad:fe:df:6c:cb:3c:b0:4f:cc:dd
    ac:d8:69:3f:3c:59:b5:2a:c6:83:e8:b3:94:5e:0a:2d:cd:e2:b0:15
    3e:97:a7:8a:4e:fb:59:f7:20:4c:ba:a8:80:3e:ba:be:69:ef:2b:32
    e4:1a:1c:24:53:1b:d5:c3:aa:d4:25:73:96:76:ea:49:d4:da:7e:3e
    0c:c6:6b:22:43:cb:04:84:0d:25:33:07:6b:49:41:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    ldap:///CN=example-cn1,CN=example,CN=CDP,CN=Public%20Key
%20Services,CN=Services,CN=Configuration,DC=demo,DC=local?certificateRevocationList?base?
objectClass=cRLDistributionPoint
    http://user.device.example.net/CertEnroll/device-user.crl
  Use for key: Key encipherment, Digital signature, 1.3.6.1.5.5.8.2.2,
  1.3.6.1.5.5.8.2.2
  Fingerprint:
    76:a8:5f:65:b4:bf:bd:10:d8:56:82:65:ff:0d:04:3a:a5:e9:41:dd (sha1)
    8f:99:a4:15:98:10:4b:b6:1a:3d:81:13:93:2a:ac:e7 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate system-generated

```
user@host> show security pki local-certificate system-generated
Certificate identifier: system-generated
  Issued to: JN1XXXXXX, Issued by: CN = JN10BXXXXX, CN = system generated, CN =
self-signed
  Validity:
    Not before: 10-30-2009 23:02
    Not after: 10-29-2014 23:02
  Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki local-certificate system-generated detail

```
user@host> show security pki local-certificate system-generated detail
Certificate identifier: system-generated
  Certificate version: 3
  Serial number: a1bc122bcaaaabbc1
  Issuer:
    Common name: JN1XXXXXXX, Common name: system generated, Common name:
self-signed
  Subject:
    Common name: JN1XXXXXXX, Common name: system generated, Common name:
self-signed
  Subject string:
    CN=JN10B9390AGB, CN=system generated, CN=self-signed
  Validity:
    Not before: 10-30-2009 23:02
    Not after: 10-29-2014 23:02
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:cb:c8:3f:e6:d3:e5:ca:9d:dc:2d:e9:ca:c7
    5f:b1:f5:3a:f0:1c:a7:55:43:0f:ef:fd:1c:fe:29:09:d5:37:d0:fa
    d6:ee:bc:b8:3f:58:d4:31:fb:96:4f:4f:cc:a9:1a:8f:2e:1b:50:6f
    2b:88:34:74:b2:6d:ad:94:b5:dd:3d:80:87:56:d0:42:50:4d:ac:d7
    8c:21:06:2d:07:1e:f4:d0:c7:85:2e:25:60:ad:1b:b5:b2:d2:1d:c8
    79:67:8c:56:06:04:75:6e:be:4e:99:b8:07:e6:9a:11:fe:b5:ec:c0
    1e:68:da:47:99:1b:b2:c8:07:ab:cd:6e:fe:c1:fd:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    be:1f:21:13:71:cd:9d:de:7a:41:d7:4c:52:8d:3e:d6:ba:db:75:96 (sha1)
    ba:fc:90:4b:5f:a8:66:a3:b9:64:89:9f:e2:45:b5:84 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Sample Output

show security pki local-certificate certificate-id mycert - (local certificate enrolled online using SCEP)

```
user@host> show security pki local-certificate certificate-id mycert
Certificate identifier: mycert
  Issued to: user1, Issued by: DC = local, DC = demo, CN = example-cn1
  Validity:
    Not before: 11-15-2012 18:58
    Not after: 11-15-2014 18:58
  Public key algorithm: rsaEncryption(1024 bits)
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate certificate-id mycert detail - (local certificate enrolled online using SCEP)

```
user@host> show security pki local-certificate certificate-id mycert detail
Certificate identifier: mycert
  Certificate version: 3
  Serial number: 1f00b50a0000000XXXX
  Issuer:
    Common name: example-cn1,
    Domain component: local, Domain component: demo
  Subject:
    Organization: example, Organizational unit: SSD, Country: US,
    Common name: user1, Serial number: SRX240-11152012
  Subject string:
    serialNumber=SRX240-11152012, C=US, O=example, OU=SSD, CN=user
  Alternate subject: "user@example.net", user.example.net, 10.150.1.2
  Validity:
    Not before: 11-15-2012 18:58
    Not after: 11-15-2014 18:58
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:e3:e5:ae:c0:82:af:db:94:01:2f:56:46:50
    7d:3d:0b:0c:f0:1f:1d:7d:c3:aa:d4:4c:a0:cd:23:8b:3f:47:05:ee
    7b:65:42:a0:dc:c4:ac:a7:b6:a6:9f:5c:ea:d8:22:b0:bf:03:75:09
    be:fa:77:cb:d6:67:19:e6:80:fa:a5:7c:93:af:96:66:9f:cc:45:d5
    eb:ab:c1:f0:32:a6:d9:27:1b:80:bb:57:ec:31:a2:e0:2b:e1:42:c0
    92:8a:9b:ed:a6:d2:ec:7c:84:5a:8a:d9:96:a7:7e:40:c3:80:0e:f4
    d6:a2:5d:78:93:3b:7d:d5:8a:f5:de:fb:bc:0d:6d:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    ldap:///CN=example-cn1,CN=example,CN=CDP,CN=Public%20Key%20Services,
CN=Services,CN=Configuration,DC=demo,DC=local?certificateRevocationList?
base?objectClass=cRLDistributionPoint
    http://user.device.example.net/CertEnroll/device-user.crl
  Use for key: Key encipherment, Digital signature, 1.3.6.1.5.5.8.2.2,
  1.3.6.1.5.5.8.2.2
  Fingerprint:
    1f:2f:a9:22:a8:d5:a9:36:cc:c4:bd:81:59:9d:9c:58:bb:40:15:72 (sha1)
    51:27:e4:d5:29:90:f7:85:9e:67:84:a1:75:d1:5b:16 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 805
show security policies policy-name p1 detail on page 806
show security policies (services-offload) on page 807
show security policies detail on page 807
show security policies detail (TCP Options) on page 808
show security policies policy-name p1 (Negated Address) on page 808
show security policies policy-name p1 detail (Negated Address) on page 809
show security policies global on page 809

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 51: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 51: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 51: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```

Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                 272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
```

```
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                 272 bps
    Input  packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Output packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Session rate       :                 108                   3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```

Copyright © 2016, Juniper Networks, Inc.

show services application-identification application

Syntax show services application-identification application [detail *application-name*| summary]

Release Information Command introduced in Junos OS Release 11.4.

Description Display detailed information about a specified application signature, all application signatures, or a summary of the existing application signatures and nested application signatures. Both custom and predefined application signatures and nested application signatures can be displayed.

Required Privilege Level view

Related Documentation
- request services application-identification application on page 758

List of Sample Output
show services application-identification application summary on page 811
show services application-identification application detail on page 811

Output Fields Table 52 lists the output fields for the **show services application-identification application** command. Output fields are listed in the approximate order in which they appear.

Table 52: show services application-identification application Output Fields

Field Name	Field Description
Application(s)	The number of application signatures present.
Nested Application(s)	The number of nested application signatures present.
Application Name	Name of the predefined or custom application signature. Must be a unique name with a maximum length of 32 characters.
Application Group	Name of the application signature group associated with this application signature/nested application signature. Must be a unique name with a maximum length of 32 characters.
Disabled	The status of the application signature or nested application signature and whether the signature method is currently used to identify this application. The default is No.
ID	The unique ID number of an application signature or nested application signature. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and nested application signatures; these IDs do not change. ID numbers for custom application signatures and nested application signatures use ID numbers 32,768 to 65,534.
Order	A unique number used to specify priority when multiple patterns are matched for the same session. The lowest order number takes the highest priority.

Copyright © 2016, Juniper Networks, Inc.

Table 52: show services application-identification application Output Fields *(continued)*

Field Name	Field Description
Application Tags	General information about this application type, for example, associated risk factors, technology, type of traffic, and so on. Support of application signature tags is dependent on the version of the loaded signature database. Please refer to the Juniper Networks security portal for further information.
Port Mapping: Default ports	The default port for this application type.
Signature: Port range	Default ranges: TCP/0 through 65,535; UDP/0 through 65,535 (optional).
Client-to-server: DFA Pattern	Pattern-matching scheme used for client-to-server traffic. Maximum length is 1023 (optional).
Client-to-server: Regex Pattern	A compatible regular expression used to match client-to-server traffic.
Server to-client: DFA Pattern	Pattern-matching scheme used for server-to-client traffic. Maximum length is 1023 (optional).
Server-to-client: Regex Pattern	A compatible regular expression used to match server-to-client traffic.
Minimum Data	The minimum number of bytes or packets to apply to the DFA pattern. Default is 10; range is 4 through 1024.

Sample Output

show services application-identification application summary

```
user@host> show services application-identification application summary
Application(s): 150
Nested Application(s): 600

Application              Disabled      ID        Order
...                      ...           ...       ...
junos:FTP                No            63        59
junos:HTTP               Yes           64        122
...                      ...           ...       ...
my:APPLICATION-A         No            700       730
my:APPLICATION-B         No            701       731
...                      ...           ...       ...
```

show services application-identification application detail

```
user@host> show services application-identification detail junos:FTP
Description: This signature detects the File Transfer Protocol (FTP), which
provides facilities for transferring files to and from remote computer systems.
            It usually runs on TCP port 21.
Application ID: 63
Disabled: No
Number of Parent Group(s): 1
Application Groups:
    junos:file-server
Application Tags:
    characteristic        : Supports File Transfer
```

Copyright © 2016, Juniper Networks, Inc.

```
    characteristic        : Known Vulnerabilities
    characteristic        : Capable of Tunneling
    risk                  : 3
    category              : FILE-SERVER
Port Mapping:
    Default ports: TCP/21
Signature:
     Port range: TCP/0-24,26-65535
     Client-to-server
         DFA Pattern:
\[(USER|STAT|PORT|CHMOD|ACCOUNT|BYE|ASCII|CLOB|HELP|AUTH|SYST|QUIT|STOR|PASV|CWD|PWD|MDTM|FEAT|OPTS)\](\s|\x0d
 0a\x|\x0a\x).*
         Regex Pattern: None
     Server-to-client
         DFA Pattern: (220|230|331|530)[\s\-].*
         Regex Pattern: None
     Minimum data client-to-server: 8
     Minimum data server-to-client: 8
     Order: 71
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification application-system-cache (View)

Syntax show services application-identification application-system-cache

Release Information Command introduced in Junos OS Release 10.2. Command updated in Junos OS Release 12.1X47-D10. Output updated in Junos OS Release 12.1X47-D15.

Description Display application ID from default port/protocol binding or from the application system cache.

Required Privilege Level view

Related Documentation
- clear services application-identification application-system-cache (Junos OS) on page 750

List of Sample Output show services application-identification application-system-cache on page 814

Output Fields Table 53 lists the output fields for the **show services application-identification application-system-cache** command. Output fields are listed in the approximate order in which they appear.

Table 53: show services application-identification application-system-cache Output Fields

Field Name	Field Description
application-cache	On or Off status of the application cache.
nested-application-cache	On or Off status of the nested application cache.
cache-unknown-result	On or Off status for caching unknown results.
cache-entry-timeout	The number of seconds the mapping information is saved.
pic	PIC number of the accumulated statistics. NOTE: The PIC number is always displayed as 0 for branch SRX Series devices.
Logical system name	Name of a specific logical system.
IP address	IP address.
Port	Port number.
Protocol	Type of protocol.
Application	Name of the application.
Encrypted	Yes or No to identify the traffic as encrypted or not.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show services application-identification application-system-cache

```
user@host> show services application-identification application-system-cache
Application System Cache Configurations:
  application-cache: on
  nested-application-cache: on
  cache-unknown-result: on
  cache-entry-timeout: 3600 seconds
  pic: 1/0
  Logical system name: root-logical-system
  IP address: 5.0.0.1                                  Port: 443    Protocol: TCP

  Application: SSL                                      Encrypted: Yes

  pic: 1/1
  Logical system name: root-logical-system
  IP address: 5.0.0.1                                  Port: 80     Protocol: TCP

  Application: HTTP                                     Encrypted: No
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification commit-status

Syntax show services application-identification commit-status]

Release Information Command introduced in Junos OS Release 15.1X49-D40.

Description Display information about the commit status. Because the custom signatures commit is performed asynchronously, the command output shows the current status of your configuration commit.

Required Privilege Level view

Related Documentation
- request services application-identification application on page 758

List of Sample Output
show services application-identification commit-status on page 815
show services application-identification commit-status on page 815
show services application-identification commit-status on page 815

Sample Output

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures commit is in progress
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures committed successfully
```

show services application-identification commit-status

```
user@host> show services application-identification commit-status
Custom signatures serialization failed
```

Copyright © 2016, Juniper Networks, Inc.

show services application-identification counter (AppSecure)

Syntax show services application-identification counter <ssl-encrypted-sessions>

Release Information Command introduced in Junos OS Release 10.2. Output updated in Junos OS Release 12.1X47-D10. Command and output updated in Junos OS Release 12.1X47-D15.

Description Display the status of all Junos OS application identification counter values per SPU.

Options **ssl-encrypted-sessions**—Display counters for SSL encrypted sessions.

Required Privilege Level view

Related Documentation
- clear services application-identification counter (Values) on page 751

List of Sample Output
show services application-identification counter on page 817
show services application-identification counter ssl-encrypted-sessions on page 818

Output Fields Table 54 lists the output fields for the **show services application-identification counter** command. Output fields are listed in an approximate order in which they appear.

Table 54: show services application-identification counter Output Fields

Field Name	Field Description
PIC	PIC number of the accumulated statistics. NOTE: The PIC number is always displayed as 0 for branch SRX Series devices.
Unknown applications	Number of unknown applications.
Encrypted unknown applications	Number of encrypted unknown applications.
Cache hits	Number of sessions that matched the application in the AI cache.
Cache misses	Number of sessions that did not find the application in the AI cache.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.
Client-to-server bytes processed	Number of client-to-server payload bytes processed.
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Client-to-server packets processed	Number of client-to-server packets processed.
Server-to-client packets processed	Number of server-to-client packets processed.

Copyright © 2016, Juniper Networks, Inc.

Table 54: show services application-identification counter Output Fields *(continued)*

Field Name	Field Description
Client-to-server bytes processed	Number of client-to-server payload bytes processed.
Server-to-client layer bytes processed	Number of server-to-client payload bytes processed.
Client-to-server encrypted packets processed	Number of client-to-server encrypted packets processed.
Server-to-client encrypted packets processed	Number of server-to-client encrypted packets processed.
Client-to-server encrypted bytes processed	Number of client-to-server encrypted payload bytes processed.
Server-to-client layer encrypted bytes processed	Number of server-to-client encrypted payload bytes processed.
Sessions bypassed due to resource allocation failure	Number of sessions bypassed due to resource allocation failure.
Segment case 1 - New segment to left	TCP segments contained before the previous segment.
Segment case 2 - New segment overlap right	TCP segments that start before the previous segment and are contained in it.
Segment case 3 - Old segment overlapped	TCP segments that start before the previous segment and extend beyond it.
Segment case 4 - New segment overlapped	TCP segments that start and end within the previous segment.
Segment case 5 - New segment overlap left	TCP segments that start within the previous segments and extend beyond it.
Segment case 6 - New segment overlap left	TCP segments that start after the previous segment. This is the normal case.

Sample Output

show services application-identification counter

```
user@host> show services application-identification counter

 pic: 6/0
  Counter type                                                    Value
 Unknown applications                                                5
 Encrpted unknown applications                                       0
 Cache hits                                                          0
 Cache misses                                                        8
 Client-to-server packets processed                                678
```

Copyright © 2016, Juniper Networks, Inc.

```
 Server-to-client packets processed                                  0
 Client-to-server bytes processed                                    83577
 Server-to-client bytes processed                                    0
 Client-to-server encrypted packets processed                        0
 Server-to-client encrypted packets processed                        0
 Client-to-server encrypted bytes processed                          0
 Server-to-client encrypted bytes processed                          0
 Sessions bypassed due to resource allocation failure                0
 Segment case 1 - New segment to left                                0
 Segment case 2 - New segment overlap right                          0
 Segment case 3 - Old segment overlapped                             0
 Segment case 4 - New segment overlapped                             0
 Segment case 5 - New segment overlap left                           0
 Segment case 6 - New segment to right                               0
```

Sample Output

show services application-identification counter ssl-encrypted-sessions

```
user@host> show services application-identification counter ssl-encrypted-sessions

 pic: 1/0
Counter type                                                      Value
AI cache hits                                                       0
AI cache hits by nested application                                 0
AI cache misses                                                     0
AI matches                                                          0
AI uni-matches                                                      0
AI no-matches                                                       0
AI partial matches                                                  0
AI no-partial matches                                               0
Sessions that triggered Appid create session API                    0
Sessions that do not incur signature match or decoding              0
Sessions that incur signature match or decoding                     0
Client-to-server packets processed                                  0
Server-to-client packets processed                                  0
Client-to-server layer-7 bytes processed                            0
Server-to-client layer-7 bytes processed                            0
Terminal first data packets on both direction                       0
pic: 1/1
Counter type                                                    Value
AI cache hits                                                       0
AI cache hits by nested application                                 0
AI cache misses                                                     0
AI matches                                                          0
AI uni-matches                                                      0
AI no-matches                                                       0
AI partial matches                                                  0
AI no-partial matches                                               0
Sessions that triggered Appid create session API                    0
Sessions that do not incur signature match or decoding              0
Sessions that incur signature match or decoding                     0
Client-to-server packets processed                                  0
Server-to-client packets processed                                  0
Client-to-server layer-7 bytes processed                            0
Server-to-client layer-7 bytes processed                            0
Terminal first data packets on both direction                       0
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification group

Syntax show services application-identification group [detail *application-group name*| summary]

Release Information Command introduced in Junos OS Release 11.4.

Description Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed.

Required Privilege Level view

Related Documentation
- request services application-identification group on page 762

List of Sample Output
show services application-identification group summary on page 819
show services application-identification group detail on page 820

Output Fields Table 55 lists the output fields for the **show services application-identification group** command. Output fields are listed in the approximate order in which they appear.

Table 55: show services application-identification group Output Fields

Field Name	Field Description
Description	Description of the specified application in the detailed display.
Group ID or **ID**	The unique ID number of an application signature or application signature group. ID numbers 1 through 32,767 are automatically generated for predefined application signatures and application signature groups; these IDs do not change. ID numbers for custom application signatures and application signature groups use ID numbers 32,768 to 65,534.
Disabled	The status of the application signature group and whether the signature method is currently used to identify this application. The default is No.
Application Group(s)	The application signature groups present.
Applications	The application signatures associated with this application signature group.

Sample Output

show services application-identification group summary

```
user@host> show services application-identification group summary
Application Group(s): 24
Application Groups                                Disabled  ID
  my:enterprise                                    No        32770
  junos:enterprise:voip                            No        25
  junos:peer-to-peer:voip                          No        24
  junos:peer-to-peer:chat                          No        23
  junos:peer-to-peer:file-sharing                  No        22
...
```

Copyright © 2016, Juniper Networks, Inc.

show services application-identification group detail

```
user@host> show services application-identification group detail junos:social-networking
Description: Detection for social networking sites such as Myspace, Facebook and
 similar.
Group ID: 8
Disabled: no
Application-groups:
    junos:social-networking:facebook;
    junos:social-networking:myspace;

Applications:
    junos:4CHAN;
    junos:ADULTFRIENDFINDER;
    junos:BADOO;
    junos:BEBO;
    junos:BLOGGER-POST;
    junos:BLOGSPOT-POST;
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification statistics applications

Syntax show services application-identification statistics applications <interval *interval-number*>

Release Information Command introduced in Junos OS Release 11.4. Command updated in Junos OS Release 12.1.

Description Display application usage statistics.

Options
- none—Display cumulative session and byte statistics per application. Statistics are displayed in alphabetical order.
- **interval** ***interval-number***—(Optional) Display interval statistics per application. Interval statistics are displayed in Top-N format, such that the first application displayed has the largest byte count. If this parameter is not specified, then the default is 1, which is the current interval. The previous interval is 2, and the least current (oldest) is 8.

Required Privilege Level view

Related Documentation
- statistics (Services) on page 728
- clear services application-identification application-statistics on page 747

List of Sample Output
show services application-identification statistics applications on page 821
show services application-identification statistics applications interval 3 on page 822

Output Fields Table 56 lists the output fields for the **show services application-identification statistics applications** command. Output fields are listed in the approximate order in which they appear.

Table 56: show services application-identification statistics applications Output Fields

Field Name	Field Description
Application	Name of the application.
Sessions	Number of sessions for the application.
Bytes	Size of the application in bytes.
Encrypted	Yes or No identifying the traffic as encrypted or not.

Sample Output

show services application-identification statistics applications

```
user@host> show services application-identification statistics applications

Last Reset: 2014-02-19 00:38:01 PST
                     Application           Sessions              Bytes
Encrypted
```

```
                                    SYSLOG                  2                18610
  No
```

show services application-identification statistics applications interval 3

```
user@host> show services application-identification statistics applications interval 8

Interval Start: 2014-02-19 21:10:29 PST
Elapsed time: 00:07:14
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification statistics application-groups

Syntax show services application-identification statistics application-groups <interval *interval-number*>

Release Information Command introduced in Junos OS Release 11.4.

Description Display application group usage statistics.

Options

- none—Display cumulative session and byte statistics per application group. Statistics are displayed in alphabetical order.
- **interval *interval-number*—** (Optional) Display interval statistics per application group. Interval statistics are displayed in Top-N format, such that the first application group displayed has the largest byte count. If this parameter is not specified, then the default is 1, which is the current interval. The previous interval is 2, and the least current (oldest) is 8.

Required Privilege Level view

Related Documentation

- statistics (Services) on page 728
- clear services application-identification application-statistics on page 747

List of Sample Output

show services application-identification statistics application-groups on page 823
show services application-identification statistics application-groups interval 8 on page 824

Output Fields Table 57 lists the output fields for the **show services application-identification statistics application-groups** command. Output fields are listed in the approximate order in which they appear.

Table 57: show services application-identification statistics application-groups Output Fields

Field Name	Field Description
Application Group	Displays the name of the application group.
Sessions	Displays the number of sessions for the application group.
Kilo Bytes	Displays the size of the application group in kilobytes.

Sample Output

show services application-identification statistics application-groups

```
user@host> show services application-identification statistics application-groups

Last Reset: 2014-02-19 00:38:01 PST
                Application Group          Sessions         Kilo Bytes
              junos:infrastructure                 2                 18
```

Copyright © 2016, Juniper Networks, Inc.

```
                junos:encryption                    1                 2
      junos:infrastructure:monitoring               2                 18
```

show services application-identification statistics application-groups interval 8

```
user@host> show services application-identification statistics application-groups interval 8

Interval Start: 2014-02-19 21:07:29 PST
Elapsed time: 00:07:15
```


Copyright © 2016, Juniper Networks, Inc.

show services application-identification status

Syntax show services application-identification status

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Display detailed information about application identification status.

Required Privilege Level view

Related Documentation
- request services application-identification application on page 758

List of Sample Output
show services application-identification status on page 826
show services application-identification status (DPI Performance Mode Enabled) on page 826

Output Fields Table 58 lists the output fields for the **show services application-identification status** command. Output fields are listed in the approximate order in which they appear.

Table 58: show services application-identification status Output Fields

Field Name	Field Description
Status	Status of application identification: **Enabled** or **Disabled**.
Sessions under app detection	Sessions undergoing application identification detection.
Engine Version	Application identification detector engine version.
Max TCP session packet memory	Maximum number of TCP sessions that application identification maintains.
Force packet plugin	Force packet plugin status: **Enabled** or **Disabled**.
Force stream plugin	Force stream plugin status: **Enabled** or **Disabled**.
DPI Performance mode	DPI performance mode status. This field is displayed only if the DPI performance mode is enabled.
Statistics collection interval	Frequency (in minutes) for collecting statistics.
Status	Status of application system cache: **Enabled** or **Disabled**.
Negative cache status	Status on the number of sessions that reach the Unknown cache entry: **Enabled** or **Disabled**.
Max Number of entries in cache	Maximum number of cache entries.

Copyright © 2016, Juniper Networks, Inc.

Table 58: show services application-identification status Output Fields *(continued)*

Field Name	Field Description
Cache timeout	Idle timeout after which the cache entries expires.
Download Server CGI	Name of the server from where protocol bundle was downloaded.
Auto Update	Status of auto update to receive protocol bundle updates from the server: **Enabled** or **Disabled**.
Status	Status of protocol bundle: **Active** or **Free**.
Version	Version of protocol bundle.
Session	The number of active sessions.

Sample Output

show services application-identification status

```
user@host> show services application-identification status
pic: 5/0

Application Identification
 Status                            Enabled
 Sessions under app detection      0
 Engine Version                    4.18.1-20 (build date Feb 15 2014)
 Max TCP session packet memory     30000
 Force packet plugin               Disabled
 Force stream plugin               Disabled
 Statistics collection interval    1 (in minutes)

Application System Cache
 Status                            Enabled
 Negative cache status             Disabled
 Max Number of entries in cache    131072
 Cache timeout                     3600 (in seconds)

Protocol Bundle
 Download Server                https://services.netscreen.com/cgi-bin/index.cgi

 AutoUpdate                        Disabled
Slot 1:
 Status                            Active
 Version                           1.30.4-22.005 (build date Jan 17 2014)
 Sessions                          0
Slot 2
 Status                            Free
```

Sample Output

show services application-identification status (DPI Performance Mode Enabled)

```
user@host> show services application-identification status
```


Copyright © 2016, Juniper Networks, Inc.

```
pic: 2/1

Application Identification
Status                            Enabled
Sessions under app detection      0
Engine Version                    4.18.2-24.006 (build date Jul 30 2014)
Max TCP session packet memory     30000
Force packet plugin               Disabled
Force stream plugin               Disabled
DPI Performance mode:             Enabled
Statistics collection interval    1 (in minutes)

Application System Cache
Status                            Enabled
Negative cache status             Disabled
Max Number of entries in cache    262144
Cache timeout                     3600 (in seconds)

Protocol Bundle
Download Server                 https://services.netscreen.com/cgi-bin/index.cgi
AutoUpdate                        Disabled
Slot 1:
Application package version       2399
Status                            Active
Version                           1.40.0-26.006 (build date May 1 2014)
Sessions                          0
Slot 2
Application package version       0
Status                            Free
Version
Sessions                          0
```

Copyright © 2016, Juniper Networks, Inc.

show services application-identification version

Syntax show services application-identification version

Release Information Command introduced in Junos OS Release 10.2.

Description Display the Junos OS application package version.

Required Privilege Level view

Related Documentation
- request services application-identification download on page 760

List of Sample Output show services application-identification version on page 828

Sample Output

show services application-identification version

The following output shows that the application package version is 1608.

```
user@host> show services application-identification version
Application package version: 1608
```


Copyright © 2016, Juniper Networks, Inc.

show services ssl proxy statistics

Syntax **show services ssl proxy statistics**

Release Information Command introduced in Junos OS Release 12.1.

Description Display information about the SSL proxy statistics.

Options **none**—Display summary information about SSL proxy.

Required Privilege Level view

Related Documentation
- clear services ssl proxy statistics on page 752

List of Sample Output show services ssl proxy statistics on page 829

Output Fields Table 59 describes the output fields for the **show services ssl proxy statistics** command. Output fields are listed in the approximate order in which they appear.

Table 59: show services ssl proxy statistics Output Fields

Field Name	Field Description
Sessions matched	The number of proxy sessions that are matched.
Sessions whitelisted	The number of sessions that are whitelisted. Whitelists comprise addresses or domain names that you want to exempt from the SSL proxy processing.
Sessions bypassed: non SSL	The number of proxy sessions that are bypassed because the non SSL sessions limit was exceeded
Sessions bypassed: memory overflow	The number of proxy sessions that are bypassed because the memory usage limit per session was reached.
Sessions created	The number of proxy sessions that are newly created.
Sessions ignored	The number of proxy sessions that are ignored.
Sessions active	The number of proxy sessions that are active.
Sessions dropped	The number of proxy sessions that are dropped.

Sample Output

show services ssl proxy statistics

```
user@host> show services ssl proxy statistics
PIC:fwdd0 fpc[0] pic[0] ------
        sessions matched                    30647
        sessions whitelisted                    0
```

```
sessions bypassed:non-ssl                0
sessions bypassed:mem overflow           0
sessions created                     25665
sessions ignored                         6
sessions active                          0
sessions dropped                         0
```


Copyright © 2016, Juniper Networks, Inc.

Attack Detection and Prevention Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 5

Overview

- Introduction to Attack Detection and Prevention on page 835

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 34

Introduction to Attack Detection and Prevention

- Attack Detection and Prevention Overview on page 835
- Understanding IPv6 Support for Screens on page 836
- Understanding Screens Options on SRX Series Devices on page 840
- Understanding Screen Options on the SRX5000 Module Port Concentrator on page 843
- Example: Configuring Multiple Screening Options on page 849
- Understanding Central Point Architecture Enhancements for Screens on page 854

Attack Detection and Prevention Overview

Attack detection and prevention, also known as stateful firewall, detects and prevents attacks in network traffic. An exploit can be either an information-gathering probe or an attack to compromise, disable, or harm a network or network resource. In some cases, the distinction between the two objectives of an exploit can be unclear. For example, a barrage of TCP SYN segments might be an IP address sweep with the intent of triggering responses from active hosts, or it might be a SYN flood attack with the intent of overwhelming a network so that it can no longer function properly. Furthermore, because an attacker usually precedes an attack by performing reconnaissance on the target, we can consider information-gathering efforts as a precursor to an impending attack—that is, they constitute the first stage of an attack. Thus, the term *exploit* encompasses both reconnaissance and attack activities, and the distinction between the two is not always clear.

Juniper Networks provides various detection and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution:

- Screen options at the zone level.
- Firewall policies at the inter-, intra-, and super-zone policy levels (*super-zone* here means in global policies, where no security zones are referenced).

To secure all connection attempts, Junos OS uses a dynamic packet-filtering method known as stateful inspection. Using this method, Junos OS identifies various components in the IP packet and TCP segment headers—source and destination IP addresses, source and destination port numbers, and packet sequence numbers—and maintains the state

Copyright © 2016, Juniper Networks, Inc.

of each TCP session and pseudo UDP session traversing the firewall. (Junos OS also modifies session states based on changing elements such as dynamic port changes or session termination.) When a responding TCP packet arrives, Junos OS compares the information reported in its header with the state of its associated session stored in the inspection table. If they match, the responding packet is allowed to pass the firewall. If the two do not match, the packet is dropped.

Junos OS screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone.

Related Documentation

- Understanding Security Zones on page 1050
- Understanding Network Reconnaissance Using IP Options on page 930

Understanding IPv6 Support for Screens

Juniper Networks provides various detection and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution. Screen options are at the zone level. Junos OS screen options secure a zone by inspecting it, and then allowing or denying all connection attempts that require crossing an interface bound to that zone.

By default, IPv6 packets bypass the screen module. However, you can configure screen options to check and filter packets based on IPv6 extension headers, packet headers, and ICMPv6 traffic. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IPv6 traffic.

- IPv6 Extension Header Checking and Filtering on page 836
- Maximum Number of Extension Headers on page 837
- Bad Option Extension Headers on page 838
- ICMPv6 Checking and Filtering on page 838
- IPv6 Packet Header Checking and Filtering on page 839

IPv6 Extension Header Checking and Filtering

You can use the **ipv6-extension-header** statement to selectively screen one or more extension headers. Table 60 lists common IPv6 extension headers and their type values.

Table 60: IPv6 Extension Headers and Type Values

Header Name	Header Type Value	Internet Standards
Authentication	51	RFC 2460
Encapsulating Security Payload	50	RFC 2460
Host Identify Protocol	139	RFC 5201

Copyright © 2016, Juniper Networks, Inc.

Table 60: IPv6 Extension Headers and Type Values *(continued)*

Header Name	Header Type Value	Internet Standards
Destination Options • ILNP nonce option • Home address option • Line identification option • Tunnel encapsulation limit option	60	RFC 2460
Fragment	44	RFC 2460
Hop-by-Hop Options • CALIPSO option • RPL option • SFM DPD option • Jumbo payload option • Quick start option • Router alert option	0	RFC 2460
Mobility	135	RFC 6275
No next	59	RFC 2460
Routing	43	RFC 2460
Shim6	140	RFC 5533

Maximum Number of Extension Headers

You can specify the maximum number of permitted extension headers in a packet by using the **ipv6-extension-header-limit** statement. Although the maximum number of extension headers in a packet is not explicitly specified, the order of extension headers is recommended in RFC 2460:

1. Hop-by-Hop Options header
2. Destination Options header
3. Routing header
4. Fragment extension header
5. Authentication header
6. Encapsulating Security Payload header
7. Destination Options header

Each extension header should occur at most once, except for the destination options header, which should occur at most twice (once before a routing header and once before the upper-layer protocol header).

Copyright © 2016, Juniper Networks, Inc.

The maximum extension header number based on RFC 2460 is 7. Other extension headers have been defined by subsequent RFCs. We recommend the maximum extension header number to be in the range of 0 through 32.

Bad Option Extension Headers

You can configure screens to detect and drop any packet with an incorrectly formatted IP option in the IP packet header (IPv4 or IPv6). The device records the event in the screen counters list for the ingress interface. Table 61 lists key criteria that the device uses to screen packets for bad options.

Table 61: Bad Option Extension Header Screening Criteria

Screening Criteria	Internet Standards	Description
Routing extension header is after fragment header	RFC 2460	The order of extension headers in a packet is defined; accordingly, the fragment extension header must be after the routing header.
Wrong router alert parameter	RFC 2711	This option is located in the hop-by-hop header and in the Junos OS implementation: • There can be only one option of this type per hop-by-hop header • The header length must be 2. • There can be only one router alert option in one extension header.
More than one back-to-back pad option	draft-krishnan-ipv6-hopbyhop-00	This type of traffic is screened as error packets.
Non-zero payload in PadN option	RFC 4942	The system checks that the PadN only has zero octets in its payload.
Padding beyond the next eight-octet boundary	RFC 4942	The system checks for padding beyond the next eight octet boundary. There is no legitimate reason for padding beyond the next eight octet boundary.
Jumbo payload with non-zero IPv6 header payload	RFC 2675	The payload length field in the IPv6 header must be set to zero in every packet that carries the jumbo payload option.

ICMPv6 Checking and Filtering

You can enable ICMPv6 checking and filtering. The system then checks whether the ICMPv6 packet received matches the defined criteria and performs the specified action on matching packets. Some of the key defined criteria are as follows:

- Information message of unknown type—Many types of ICMPv6 information messages are defined, such as echo request (value 128), echo reply (value 129), and router solicitation (value 133). The maximum type definition is 149. Any value higher than 149 is treated as an unknown type and screened accordingly.

Copyright © 2016, Juniper Networks, Inc.

- Does not meet the ICMPv6 ND packet format rules (RFC 4861)—There are standard rules, such as the IP Hop limit field has a value of 255, ICMP checksum must be valid, the ICMP code must be 0, and so on.
- Malformed ICMPv6 packet filtering—For instance, the ICMPv6 packet is too big (message type 2), the next header is set to routing (43), and routing header is set to hop-by hop.

IPv6 Packet Header Checking and Filtering

You can enable the checking and filtering of IPv6 packet headers using the **ipv6-malformed-header** statement. Once enabled, the system verifies any incoming IPv6 packet to check if it matches any of the defined criteria. The system then performs the specified action (drop or alarm-without-drop) on matching packets. Table 62 lists key criteria that the device uses to screen packets.

Table 62: IPv6 Packet Header Screening Criteria

Screening Criteria	Internet Standards	Description
Deprecated site-local source and destination addresses	RFC 3879	The IPv6 site-local unicast prefix (1111111011 binary or FEC0::/10) is not supported.
Illegal multicast address scope values	RFC 4291	The unassigned multicast address scope values are treated as illegal.
Documentation-only prefix (2001:DB8::/32)	RFC 3849	IANA is to record the allocation of the IPv6 global unicast address prefix (2001:DB8::/32) as a documentation-only prefix in the IPv6 address registry. No end party is to be assigned this address.
Deprecated IPv4-compatible IPv6 source and destination addresses (::/96)	RFC 4291	The IPv4-compatible IPv6 address has been deprecated and is not supported.
ORCHID source and destination addresses (2001:10::/28)	RFC 5156	Addresses of the Overlay Routable Cryptographic Hash Identifiers (2001:10::/28) are used as identifiers and cannot be used for routing at the IP layer. Addresses within this block must not appear on the public Internet.
An IPv4 address embedded inside the IPv6 address (64:ff9b::/96) is an illegal, unacceptable IPv4 address	RFC 6052	The IPv6 address, 64:ff9b::/96, is reserved as “Well-known Prefix” for use in algorithmic mapping.

Related Documentation

- Example: Configuring Multiple Screening Options on page 849
- ipv6-extension-header on page 1862
- ipv6-extension-header-limit on page 1863
- icmpv6-malformed on page 1860
- ipv6-malformed-header on page 1863

Copyright © 2016, Juniper Networks, Inc.

Understanding Screens Options on SRX Series Devices

On all SRX Series devices, the screens are divided into two categories:

- Statistics-based screens
- Signature-based screens

Table 63 lists all the statistics-based screen options.

Table 63: Statistics-Based Screen Options

Screen Option Name	Description
ICMP flood	Use the ICMP flood IDS option to protect against ICMP flood attacks. An ICMP flood attack typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed. The threshold value defines the number of ICMP packets per second (pps) allowed to ping the same destination address before the device rejects further ICMP packets.
UDP flood	Use the UDP flood IDS option to protect against UDP flood attacks. A UDP flood attack occurs when an attacker sends IP packets containing a UDP datagram with the purpose of slowing down the resources, such that valid connections can no longer be handled. The threshold value defines the number of UDP packets per second allowed to ping the same destination IP address. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.
TCP SYN flood source	Use the TCP SYN flood source IDS option to set the source threshold value. The threshold value defines the number of SYN segments to be received per second before the device begins dropping connection requests. The applicable range is 4 through 500,000 SYN pps.
TCP SYN flood destination	Use the SYN flood destination IDS option to set the destination threshold value. The threshold value defines the number of SYN segments received per second before the device begins dropping connection requests. The applicable range is 4 through 500,000 SYN pps.
TCP SYN flood	Use the TCP SYN flood IDS option to detect and prevent SYN flood attacks. Such attacks occur when the connecting host continuously sends TCP SYN requests without replying to the corresponding ACK responses.
TCP port scan	Use the TCP port scan IDS option to prevent the port scan attacks. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
TCP SYN-ACK-ACK proxy	Use the TCP SYN-ACK-ACK proxy screen option to prevent SYN-ACK-ACK attack. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, SRX Series devices running Junos OS reject further connection requests from that IP address.

Copyright © 2016, Juniper Networks, Inc.

Table 63: Statistics-Based Screen Options *(continued)*

Screen Option Name	Description
ICMP IP sweep	Use the ICMP IP sweep IDS option to detect and prevent an IP sweep attack. An IP sweep attack occurs when an attacker sends ICMP echo requests (pings) to multiple destination addresses. If a target host replies, the reply reveals the target's IP address to the attacker. If the device receives 10 ICMP echo requests within the number of microseconds specified in this statement, it flags this as an IP sweep attack, and rejects the eleventh and all further ICMP packets from that host for the remainder of the second. The threshold value defines the maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device.
TCP SYN flood alarm	Use the TCP SYN flood alarm IDS option to set the alarm threshold value. The threshold value defines the number of half-complete proxy connections per second at which the device makes entries in the event alarm log. The range is 1 through 500,000 requests per second.
TCP SYN flood attack	Use the TCP SYN flood attack IDS option to set the attack threshold value. The threshold value defines the number of SYN packets per second required to trigger the SYN proxy response. The range is 1 through 500,000 proxied pps.
UDP udp sweep	Use the UDP udp sweep IDS option to detect and prevent UDP sweep attacks. In a UDP sweep attack, an attacker sends UDP packets to the target device. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends UDP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as a UDP sweep attack. If the **alarm-without-drop** option is not set, the device rejects the eleventh and all further UDP packets from that host for the remainder of the specified threshold period. The threshold value defines the number of microseconds for which the device accepts 10 UDP packets from the same remote source to different destination addresses.

Table 64 lists all the signature-based screen options.

Table 64: Signature-Based Screen Options

Screen Option Name	Description
TCP Winnuke	Enable or disable the TCP WinNuke attacks IDS option. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP SYN fragment	Use the TCP SYN fragment attack IDS option to drop any packet fragments used for the attack. A SYN fragment attack floods the target host with SYN packet fragments. The host caches these fragments, waiting for the remaining fragments to arrive so it can reassemble them. The flood of connections that cannot be completed eventually fills the host's memory buffer. No further connections are possible, and damage to the host's operating system can occur.
TCP no flag	Use the TCP tcp no flag IDS option to drop illegal TCP packets with a missing or malformed flag field. The threshold value defines the number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
TCP SYN FIN	Use the TCP SYN FIN IDS option to detect an illegal combination of flags that attackers can use to consume sessions on the target device, thus resulting in a denial-of-service (DoS) condition.

Copyright © 2016, Juniper Networks, Inc.

Table 64: Signature-Based Screen Options *(continued)*

Screen Option Name	Description
TCP land	Enable or disable the TCP land attack IDS option. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and the source IP address.
TCP FIN no ACK	Use the FIN bit with no ACK bit IDS option to detect an illegal combination of flags, and reject packets that have this combination.
ICMP ping of death	Use the ping of death IDS option to detect and reject oversized and irregular ICMP packets. Although the TCP/IP specification requires a specific packet size, many ping implementations allow larger packet sizes. Larger packets can trigger a range of adverse system reactions, including crashing, freezing, and restarting. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
ICMP fragment	Use the ICMP fragment IDS option to detect and drop any ICMP frame with the More Fragments flag set or with an offset indicated in the **offset** field.
ICMP large	Use the ICMP large IDS option to detect and drop any ICMP frame with an IP length greater than 1024 bytes.
IP unknown protocol	Use the IP unknown protocol IDS option to discard all received IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6. Such protocol numbers are undefined or reserved.
IP bad option	Use the IP bad IDS option to detect and drop any packet with an incorrectly formatted IP option in the IP packet header. The device records the event in the screen counters list for the ingress interface. This screen option is applicable to IPv4 and IPv6.
IP strict source route option	Use the IP strict source route IDS option to detect packets where the IP option is 9 (strict source routing), and record the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable only to IPv4.
IP loose source route option	Use the IP loose source route IDS option to detect packets where the IP option is 3 (loose source routing), and record the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified. The type 0 routing header of the loose source route option is the only related header defined in IPv6.
IP source route option	Use the IP source route IDS option to detect packets and record the event in the screen counters list for the ingress interface.
IP stream option	Use the IP stream IDS option to detect packets where the IP option is 8 (stream ID), and record the event in the screen counters list for the ingress interface. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams. Currently, this screen option is applicable only to IPv4.
IP block fragment	Enable or disable the IP packet fragmentation blocking. When this feature is enabled, Junos OS denies IP fragments on a security zone and blocks all IP packet fragments that are received at interfaces bound to that zone.

Copyright © 2016, Juniper Networks, Inc.

Table 64: Signature-Based Screen Options *(continued)*

Screen Option Name	Description
IP record route option	Use the IP record route IDS option to detect packets where the IP option is 7 (record route), and record the event in the screen counters list for the ingress interface. This option records the IP addresses of the network devices along the path that the IP packet travels. Currently, this screen option is applicable only to IPv4.
IP timestamp option	Use the IP timestamp IDS option to detect packets where the IP option list includes option 4 (Internet timestamp), and record the event in the screen counters list for the ingress interface. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination. Currently, this screen option is applicable only to IPv4.
IP security option	Use the IP security IDS option to detect packets where the IP option is 2 (security), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
IP spoofing	Use the IP address spoofing IDS option to prevent spoofing attacks. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
IP tear drop	Use the IP tear drop IDS option to block teardrop attacks. Teardrop attacks occur when fragmented IP packets overlap and cause the host attempting to reassemble the packets to crash. The tear drop option directs the device to drop any packets that have such a discrepancy. Teardrop attacks exploit the reassembly of fragmented IP packets.

Related Documentation

- Understanding Screen Options on the SRX5000 Module Port Concentrator on page 843
- Example: Configuring Multiple Screening Options on page 849

Understanding Screen Options on the SRX5000 Module Port Concentrator

The SRX5000 line Module Port Concentrator (SRX5K-MPC) supports Junos OS screen options. Screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone.

Using screen options, your security device can protect against different internal and external attacks, including SYN flood attacks, UDP flood attacks, and port scan attacks. Junos OS applies screen checks to traffic prior to the security policy processing, resulting in less resource utilization.

The screen options are divided into the following two categories:

- Statistics-based screens
- Signature-based screens

Copyright © 2016, Juniper Networks, Inc.

Statistics-Based Screens

All screen features implemented on an SRX5K-MPC are independent of Layer 2 or Layer 3 mode. The flood protections are used to defend against SYN flood attacks, session table flood attacks, firewall denial-of-service (DoS) attacks, and network DoS attacks.

The following four types of threshold-based flood protection are performed on each processor for both IPv4 and IPv6:

- UDP-based flood protection
- ICMP-based flood protection
- TCP source-based SYN flood protection
- TCP destination-based SYN flood protection

NOTE: If one of the two types of TCP SYN flood protections is configured on a zone, the second type of TCP SYN flood protection is automatically enabled on the same zone. These two types of protections always work together.

Each type of flood protection is threshold-based, and the threshold is calculated per zone on each microprocessor. If the flood is detected on a microprocessor chip, that particular microprocessor takes action against the offending packets based on the configuration:

- Default action (report and drop)—Screen logging and reporting are done on an SPU, so offending packets need to be forwarded to the central point or SPU for this purpose. To protect SPUs from flooding, only the first offending packet for each screen in a zone is sent to the SPU for logging and reporting in each second. The rest of the offending packets are counted and dropped in a microprocessor.

  For example, assume UDP flooding is configured at a logical interface with a threshold set to 5000 packets per second. If UDP packets come in at the rate of 20,000 per second, then about 5000 UDP packets are forwarded to the central point or SPU each second, and the remaining packets are detected as flooding. However, only one UDP flooding packet is sent to the SPU for logging and reporting in each second. The remaining packets are dropped in the microprocessor.

- Alarm only (alarm-without-drop)—An offending packet detected by screen protection is not dropped. It skips the rest of the screen checks and is forwarded to the central point or SPU with the screen result copied to its meta-header. It is not counted as a dropped packet.

Differences Between IOC1 and IOC2

The behavior of screens is the same whether the device has either IOC1 or an IOC2 card. However, there are differences in the threshold values for the statistics-based screens. Table 65 lists the statistics-based screen options and the behavior of the screens depending on whether the device has either IOC1 or an IOC2 card.

Copyright © 2016, Juniper Networks, Inc.

Table 65: Statistics-Based Screen Options

Screen Option Name	Description	IOC1	IOC2
ICMP flood	Sets the ICMP flood threshold value. The ICMP flood screen option is used to protect against ICMP flood attacks. An ICMP flood attack typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed. The threshold value defines the number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps. On SRX5000 line devices, when the IOC2 card is in start of file (SOF) mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.

Copyright © 2016, Juniper Networks, Inc.

Table 65: Statistics-Based Screen Options *(continued)*

Screen Option Name	Description	IOC1	IOC2
UDP flood	Sets the UDP flood threshold value. The UDP flood screen option is used to protect against UDP flood attacks. UDP flood attack occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled. The threshold value defines the number of UDP pps allowed to ping the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps. On SRX5000 line devices, when the IOC2 card is in start of file (SOF) mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.

Copyright © 2016, Juniper Networks, Inc.

Table 65: Statistics-Based Screen Options *(continued)*

Screen Option Name	Description	IOC1	IOC2
TCP SYN flood source	Sets the TCP SYN flood source threshold value. The threshold value defines the number of SYN segments to be received per second before the device begins dropping connection requests. The applicable range is 4 through 500,000 SYN pps.	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised..	On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps. On SRX5000 line devices, when the IOC2 card is in start of file (SOF) mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.

Copyright © 2016, Juniper Networks, Inc.

Table 65: Statistics-Based Screen Options *(continued)*

Screen Option Name	Description	IOC1	IOC2
TCP SYN flood destination	Sets the TCP SYN flood destination threshold value. The threshold value defines the number of SYN segments received per second before the device begins dropping connection requests. The applicable range is 4 through 500,000 SYN pps.	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps. On SRX5000 line devices, when the IOC2 card is in start of file (SOF) mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.

NOTE: On SRX5400, SRX5600, and SRX5800 line devices, the screen threshold value is set for each IOC in the DUT for the LAG/LACP and RLAG/RETH child links. When you have cross-IOC child interfaces as a part of LAG/LACP or RETH/RLAG interfaces and the ingress traffic is also traversing multiple child links across IOCs, set the threshold value to match the total number of packets passed by the screen from multiple IOCs with the expected total number of packets per second (pps) at the egress interface.

Signature-Based Screens

The SRX5K-MPC provides signature-based screen options along with sanity checks on the received packet.

Sometimes packets received by the device are malformed or invalid, and they might cause damage to the device and network. These packets must be dropped during initial stages of processing.

For both signature-based screen options and sanity checks, the packet contents, including packet header, status and control bits, and extension headers (for IPv6), are examined.

Copyright © 2016, Juniper Networks, Inc.

You can configure the screens according to your requirements, whereas packet sanity checks are performed by default.

The packet sanity checks and screen options are performed on packets received on ingress interfaces.

The processor does sanity checks and runs some screen features to detect the malformed and malicious ingress packets received from physical interfaces. Packets that fail a sanity check are counted and dropped.

The following packet sanity checks are supported:

- IPv4 sanity check
- IPv6 sanity check

The following screen features are supported:

- IP-based screen
- UDP-based screen
- TCP-based screen
- ICMP-based screen

The screen features are applicable to both IPv4 and IPv6 packets, with the exception of IP options screens, which only apply to IPv4 packets. If a packet is detected by one screen option, it skips the rest of the screen checks and is forwarded to the central point or Services Processing Unit (SPU) for logging and statistics collection.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Example: Configuring Multiple Screening Options

This example shows how to create one intrusion detection service (IDS) profile for multiple screening options.

- Requirements on page 849
- Overview on page 850
- Configuration on page 850
- Verification on page 852

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Copyright © 2016, Juniper Networks, Inc.

Overview

In a security zone, you can apply one IDS profile to multiple screening options. In this example we are configuring the following screening options:

- ICMP screening
- IP screening
- TCP screening
- UDP screening

These screening options are assigned to an untrust zone.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option screen-config icmp ip-sweep threshold 1000
set security screen ids-option screen-config icmp fragment
set security screen ids-option screen-config icmp large
set security screen ids-option screen-config icmp flood threshold 200
set security screen ids-option screen-config icmp ping-death
set security screen ids-option screen-config ip bad-option
set security screen ids-option screen-config ip stream-option
set security screen ids-option screen-config ip spoofing
set security screen ids-option screen-config ip strict-source-route-option
set security screen ids-option screen-config ip unknown-protocol
set security screen ids-option screen-config ip tear-drop
set security screen ids-option screen-config tcp syn-fin
set security screen ids-option screen-config tcp tcp-no-flag
set security screen ids-option screen-config tcp syn-frag
set security screen ids-option screen-config tcp port-scan threshold 1000
set security screen ids-option screen-config tcp syn-ack-ack-proxy threshold 500
set security screen ids-option screen-config tcp syn-flood alarm-threshold 500
set security screen ids-option screen-config tcp syn-flood attack-threshold 500
set security screen ids-option screen-config tcp syn-flood source-threshold 50
set security screen ids-option screen-config tcp syn-flood destination-threshold 1000
set security screen ids-option screen-config tcp syn-flood timeout 10
set security screen ids-option screen-config tcp land
set security screen ids-option screen-config tcp winnuke
set security screen ids-option screen-config tcp tcp-sweep threshold 1000
set security screen ids-option screen-config udp flood threshold 500
set security screen ids-option screen-config udp udp-sweep threshold 1000
set security zones security-zone untrust screen screen-config
```

Enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure an IDS profile for multiple screening options:

1. Configure the ICMP screening options.

```
[edit security screen ids-option screen-config]
user@host# set icmp ip-sweep threshold 1000
user@host# set icmp fragment
user@host# set icmp large
user@host# set icmp flood threshold 200
user@host# set icmp ping-death
```

2. Configure the IP screening options.

```
[edit security screen ids-option screen-config]
user@host# set ip bad-option
user@host# set ip stream-option
user@host# set ip spoofing
user@host# set ip strict-source-route-option
user@host# set ip unknown-protocol
user@host# set ip tear-drop
```

3. Configure the TCP screening options.

```
[edit security screen ids-option screen-config]
user@host# set tcp syn-fin
user@host# set tcp tcp-no-flag
user@host# set tcp syn-frag
user@host# set tcp port-scan threshold 1000
user@host# set tcp syn-ack-ack-proxy threshold 500
user@host# set tcp syn-flood alarm-threshold 500
user@host# set tcp syn-flood attack-threshold 500
user@host# set tcp syn-flood source-threshold 50
user@host# set tcp syn-flood destination-threshold 1000
user@host# set tcp syn-flood timeout 10
user@host# set tcp land
user@host# set tcp winnuke
user@host# set tcp tcp-sweep threshold 1000
```

4. Configure the UDP screening options.

```
[edit security screen ids-option screen-config]
user@host# set udp flood threshold 500
user@host# set udp udp-sweep threshold 1000
```

5. Attach the IDS profile to the zone.

```
[edit]
user@host# set security zones security-zone untrust screen screen-config
```

Results

From configuration mode, confirm your configuration by entering the **show security screen ids-option screen-config** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security screen ids-option screen-config
icmp {
    ip-sweep threshold 1000;
    fragment;
    large;
    flood threshold 200;
    ping-death;
}
    ip {
        bad-option;
        stream-option;
        spoofing;
        strict-source-route-option;
        unknown-protocol;
        tear-drop;
    }
    tcp {
        syn-fin;
        tcp-no-flag;
        syn-frag;
        port-scan threshold 1000;
        syn-ack-ack-proxy threshold 500;
        syn-flood {
            alarm-threshold 500;
            attack-threshold 500;
            source-threshold 50;
            destination-threshold 1000;
            timeout 10;
        }
        land;
        winnuke;
        tcp-sweep threshold 1000;
    }
    udp {
        flood threshold 500;
        udp-sweep threshold 1000;
    }
```

```
[edit]
user@host# show security zones
security-zone untrust {
    screen screen-config;
}
```

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

Verification

Verifying the IDS Profile for Multiple Screening Options

Purpose Verify that the IDS profile for multiple screening options is configured properly.

Copyright © 2016, Juniper Networks, Inc.

Action Enter the **show security screen ids-option screen-config Screen object status** and **show security zones** command from operational mode.

```
user@host> show security screen ids-option screen-config
Screen object status:

Name                                          Value
  ICMP flood threshold                        200
  UDP flood threshold                         500
  TCP winnuke                                 enabled
  TCP port scan threshold                     1000
  ICMP address sweep threshold                1000
  TCP sweep threshold                         1000
  UDP sweep threshold                         1000
  IP tear drop                                enabled
  TCP SYN flood attack threshold              500
  TCP SYN flood alarm threshold               500
  TCP SYN flood source threshold              50
  TCP SYN flood destination threshold         1000
  TCP SYN flood timeout                       10
  IP spoofing                                 enabled
  ICMP ping of death                          enabled
  TCP land attack                             enabled
  TCP SYN fragment                            enabled
  TCP no flag                                 enabled
  IP unknown protocol                         enabled
  IP bad options                              enabled
  IP strict source route option               enabled
  IP stream option                            enabled
  ICMP fragmentation                          enabled
  ICMP large packet                           enabled
  TCP SYN FIN                                 enabled
  TCP SYN-ACK-ACK proxy threshold             500
```

```
user@host> show security zones

Security zone: untrust
Send reset for non-SYN session TCP packets: Off
           Policy configurable: Yes
           Screen: screen-config
           Interfaces bound: 0
           Interfaces:
```

NOTE: On all SRX Series devices, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.

The synchronization cookie or proxy never drops packets; therefore the **alarm-without-drop** (not **drop**) action is shown in the system log.

Copyright © 2016, Juniper Networks, Inc.

Understanding Central Point Architecture Enhancements for Screens

The central point architecture is enhanced to achieve a higher number of connections per second (CPS) and due to the enhancements, the central point session and central point packet processing have been moved from the central point to the Services Processing Unit (SPU).

Previously, the central point had a session limit and if no resources (session limit entries) were available, then the packet was always permitted by the session limit. Now, both the central point and the SPU have session limits. If there are no resources available in the central point, but resources are available in the SPU, then the central point cannot limit the sessions but the SPU can limit the sessions. So, the SPU drops the packet SPU session limit checking and only if both the central point and the SPU have no resources will the packet *always* be perrmitted.

The following scenarios describe when the central point and the SPU determine whether to permit or drop a packet.

- When the central point has no session limit entry and the SPU has a session limit entry:
  a. If the session limit counter of the SPU is larger than the threshold value, the packet is dropped.
  b. If the session limit counter of the SPU is not larger than the threshold value, the packet is permitted.
- When the SPU does not have a session limit entry:
  a. If the session limit counter of the SPU is larger than the threshold value, the packet is permitted.
  b. If the session limit counter of the SPU is not larger than ththreshold, the packet is permitted.

NOTE: **An extra message is sent to the central point to maintain accurate session counts might impact the number of connections per second (CPS) for screens. This impacts the source or destination session limit.**

Global traffic statistics lacking a central point might impact some global view screens. Only the SYN cookie has no global view, and the global traffic statistics are handled by the SPU, so the counter might be not accurate as before. For other statistics-based screens, handled by both the central point and the SPU, the counters are accurate.

Previously, statistics-based screens were handled only by the central point and the log and the SNMP trap could be rate-limited strictly. Now both the central point and the SPU can generate the log and the SNMP trap independently. Therefore, the log and the SNMP trap might be larger than before.

Related Documentation

- Understanding Screens Options on SRX Series Devices on page 840

Copyright © 2016, Juniper Networks, Inc.

- SRX5000 Line Devices Processing Overview on page 1671

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 6

Configuring Screen Options to Protect Against Denial-of-Service Attacks

- Understanding DoS Attacks on page 859
- Protecting Against Firewall DoS Attacks on page 861
- Protecting Against Network DoS Attacks on page 873
- Protecting Against OS-Specific DoS Attacks on page 913

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 35

Understanding DoS Attacks

- DoS Attack Overview on page 859

DoS Attack Overview

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed. The target can be the firewall, the network resources to which the firewall controls access, or the specific hardware platform or operating system of an individual host.

If a DoS attack originates from multiple source addresses, it is known as a distributed denial-of-service (DDoS) attack. Typically, the source address of a DoS attack is spoofed. The source addresses in a DDoS attack might be spoofed, or the actual addresses of compromised hosts might be used as “zombie agents” to launch the attack.

The device can defend itself and the resources it protects from DoS and DDoS attacks.

Related Documentation

- Firewall DoS Attacks Overview on page 861
- Network DoS Attacks Overview on page 873
- OS-Specific DoS Attacks Overview on page 913

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 36

Protecting Against Firewall DoS Attacks

- Firewall DoS Attacks Overview on page 861
- Understanding Firewall Filters on the SRX5000 Module Port Concentrator on page 861
- Understanding Session Table Flood Attacks on page 862
- Understanding Source-Based Session Limits on page 863
- Example: Setting Source-Based Session Limits on page 864
- Understanding Destination-Based Session Limits on page 866
- Example: Setting Destination-Based Session Limits on page 867
- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 869
- Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack on page 870

Firewall DoS Attacks Overview

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed.

If attackers discover the presence of the Juniper Networks firewall, they might launch a DoS attack against it instead of the network behind it. A successful DoS attack against a firewall amounts to a successful DoS attack against the protected network in that it thwarts attempts of legitimate traffic to traverse the firewall.

An attacker might use session table floods and SYN-ACK-ACK proxy floods to fill up the session table of Junos OS and thereby produce a DoS.

Related Documentation

- DoS Attack Overview on page 859
- Network DoS Attacks Overview on page 873
- OS-Specific DoS Attacks Overview on page 913
- Understanding Session Table Flood Attacks on page 862

Understanding Firewall Filters on the SRX5000 Module Port Concentrator

The SRX5000 line Module Port Concentrator (SRX5K-MPC) for the SRX5400, SRX5600, and SRX5800 supports a firewall filter to provide filter based forwarding and packet

Copyright © 2016, Juniper Networks, Inc.

filtering at logical interfaces including the chassis loopback interface. A firewall filter is used to secure networks, to protect Routing Engines and Packet Forwarding Engines, and to ensure class of service (CoS).

The firewall filter provides:

- Filter-based forwarding at logical interfaces
- Protection of a Routing Engine from DoS attacks
- Blocking of certain types of packets to reach a Routing Engine and packet counter

The firewall filter examines packets and performs actions according to the configured filter policy. The policy is composed of match conditions and actions. The match conditions cover various fields of Layer 3 packet and Layer 4 header information. In association with the match conditions, various actions are defined in the firewall filter policy, and these actions include **accept**, **discard**, **log** counter, and so on.

After configuring the firewall filter, you can apply a logical interface to the firewall filter in the ingress or egress, or in both directions. All packets passing through the logical interface are checked by the firewall filter. As part of the firewall filter configuration, a policer is defined and applied to the logical interface. A policer restricts the traffic bandwidth at the logical interface.

NOTE: **Firewall filtering on an SRX5K-MPC does not support aggregated Ethernet interfaces.**

NOTE: **On SRX5400, SRX5600 and SRX5800 devices with an SRX5K-MPC, applying a policer at the loopback (lo0) interface ensures that the Packet Forwarding Engine discards certain types of packets and prevents them from reaching the Routing Engine.**

Related Documentation

- Firewall DoS Attacks Overview on page 861

Understanding Session Table Flood Attacks

A successful DoS attack overwhelms its victim with such a massive barrage of false simulated traffic that it becomes unable to process legitimate connection requests. DoS attacks can take many forms—SYN flood, SYN-ACK-ACK flood, UDP flood, ICMP flood, and so on—but they all seek the same objective, which is to fill up their victim's session table.

When the session table is full, that host cannot create any new sessions and begins rejecting new connection requests. The source-based session limits screen option and the destination-based session limit screen option help mitigate such attacks.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- DoS Attack Overview on page 859
- Understanding Source-Based Session Limits on page 863
- Understanding SYN Flood Attacks on page 874
- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 869

Understanding Source-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. One benefit of setting a source-based session limit is that it can stem an attack such as the Nimda virus (which is actually both a virus and a worm) that infects a server and then begins generating massive amounts of traffic from that server. Because all the virus-generated traffic originates from the same IP address, a source-based session limit ensures that the firewall can curb such excessive amounts of traffic. See Figure 33.

Figure 33: Limiting Sessions Based on Source IP Address

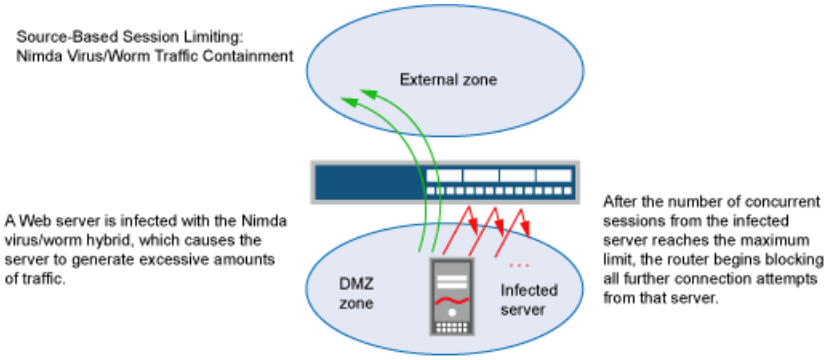

Another benefit of source-based session limiting is that it can mitigate attempts to fill up the firewall's session table if all the connection attempts originate from the same source IP address.

Determining what constitutes an acceptable number of connection requests requires a period of observation and analysis to establish a baseline for typical traffic flows. You also need to consider the maximum number of concurrent sessions required to fill up the session table of the particular Juniper Networks platform you are using. To see the maximum number of sessions that your session table supports, use the CLI command **show security flow session summary**, and then look at the last line in the output, which lists the number of current (allocated) sessions, the maximum number of sessions, and the number of failed session allocations:

```
Maximum-sessions: 32768
```

The default maximum for source-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Junos OS supports source-based session limits for both IPv4 and IPv6 traffic.

Related Documentation

- DoS Attack Overview on page 859
- Example: Setting Source-Based Session Limits on page 864

Example: Setting Source-Based Session Limits

This example shows how to limit the amount of sessions based on source IP.

- Requirements on page 864
- Overview on page 864
- Configuration on page 864
- Verification on page 865

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

The following example shows how to limit the number of sessions that any one server in the DMZ and in zone a can initiate. Because the DMZ contains only webservers, none of which should initiate traffic, you set the source-session limit at the lowest possible value, which is one session. On the other hand, zone a contains personal computers, servers, printers, and so on, many of which do initiate traffic. For zone a, you set the source-session limit to a maximum of 80 concurrent sessions.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 1-limit-session limit-session source-ip-based 1
set security zones security-zone dmz screen 1-limit-session
set security screen ids-option 80-limit-session limit-session source-ip-based 80
set security zones security-zone zone_a screen 80-limit-session
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To configure the source-based session limits:

1. Specify the number of concurrent sessions based on source IP for the DMZ zone.

```
[edit security]
user@host# set screen ids-option 1-limit-session limit-session source-ip-based 1
```


Copyright © 2016, Juniper Networks, Inc.

2. Set the security zone for the DMZ to the configuration limit.

```
[edit security]
user@host# set zones security-zone dmz screen 1-limit-session
```

3. Specify the number of concurrent sessions based on source IP for the zone a zone.

```
[edit security]
user@host# set screen ids-option 80-limit-session limit-session source-ip-based
  80
```

4. Set the security zone for zone a to the configuration limit.

```
[edit security]
user@host# set zones security-zone zone_a screen 80-limit-session
```

Results From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option 1-limit-session {
    limit-session {
    source-ip-based 1;
    }
  }
  ids-option 80-limit-session {
    limit-session {
    source-ip-based 1;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone dmz {
    screen 1-limit-session;
  }
  security-zone zone_a {
    screen 80-limit-session;
  }
```

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

Verification

Verifying Source-Based Session Limits

Purpose Verify source-based session limits.

Action Enter the **show security screen ids-option 1-limit-session** , **show security screen ids-option 80-limit-session**, and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option 1-limit-session
Screen object status:
Name                                        Value
  Session source limit threshold            1
```

```
user@host> show security screen ids-option 80-limit-session
Screen object status:
Name                                        Value
  Session source limit threshold            80

user@host> show security zones
Security zone: dmz
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 1-limit-session
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows the source session limit values for DMZ zone and zone a.

Related Documentation

- Understanding Session Table Flood Attacks on page 862
- Understanding Source-Based Session Limits on page 863
- Example: Setting Destination-Based Session Limits on page 867

Understanding Destination-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. A wily attacker can launch a distributed denial-of-service (DDoS) attack. In a DDoS attack, the malicious traffic can come from hundreds of hosts, known as "zombie agents," that are surreptitiously under the control of an attacker. In addition to the SYN, UDP, and ICMP flood detection and prevention screen options, setting a destination-based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. See Figure 34.

Copyright © 2016, Juniper Networks, Inc.

Figure 34: Distributed DOS Attack

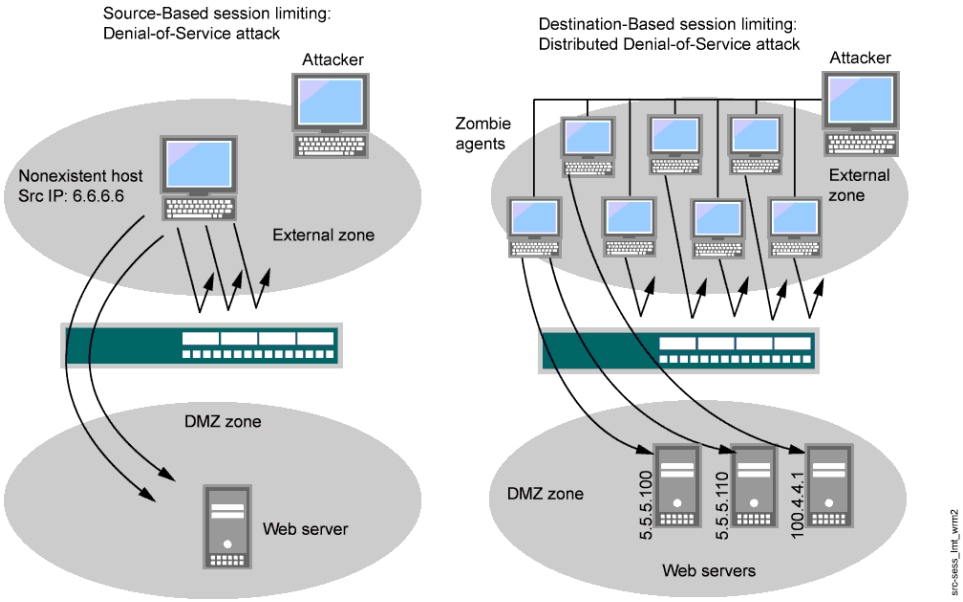

The default maximum for destination-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Related Documentation

- DoS Attack Overview on page 859
- Example: Setting Destination-Based Session Limits on page 867
- Understanding Source-Based Session Limits on page 863

Example: Setting Destination-Based Session Limits

This example shows how to set the destination-based session limits.

- Requirements on page 867
- Overview on page 868
- Configuration on page 868
- Verification on page 869

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you limit the amount of traffic to a webserver at 1.2.2.5. The server is in the DMZ. The example assumes that after observing the traffic flow from the external zone to this server for a month, you have determined that the average number of concurrent sessions it receives is 2000. Also, you set the new session limit at 2000 concurrent sessions. Although traffic spikes might sometimes exceed that limit, the example assumes that you are opting for firewall security over occasional server inaccessibility.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 2000-limit-session limit-session destination-ip-based
  2000
set security zones security-zone external_zone screen 2000-limit-session
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To set the destination-based session limits:

1. Specify the number of concurrent sessions.

```
[edit]
user@host# set security screen ids-option 2000-limit-session limit-session
  destination-ip-based 2000
```

2. Set the security zone for the external zone.

```
[edit]
user@host# set security zones security-zone external_zone screen
  2000-limit-session
```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option 2000-limit-session {
    limit-session {
    destination-ip-based 2000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone external_zone {
    screen 2000-limit-session;
  }
```


Copyright © 2016, Juniper Networks, Inc.

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

Verification

Verifying Destination-Based Session Limits

Purpose Verify destination-based session limits.

Action Enter the **show security screen ids-option 2000-limit-session** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option 2000-limit-session
node0:
--------------------------------------------------------------------------
Screen object status:

Name                                          Value
  Session destination limit threshold         2000
            Value

user@host> show security zones
Security zone: external_zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 2000-limit-session
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows the destination session limit values for external zone.

Related Documentation

- Understanding Destination-Based Session Limits on page 866
- DoS Attack Overview on page 859

Understanding SYN-ACK-ACK Proxy Flood Attacks

When an authentication user initiates a Telnet or an FTP connection, the user sends a SYN segment to the Telnet or FTP server. Junos OS intercepts the SYN segment, creates an entry in its session table, and proxies a SYN-ACK segment to the user. The user then replies with an ACK segment. At this point, the initial three-way handshake is complete. Junos OS sends a login prompt to the user. If the user, with malicious intent, does not log in but instead continues initiating SYN-ACK-ACK sessions, the firewall session table can fill up to the point where the device begins rejecting legitimate connection requests.

To prevent such an attack, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, Junos OS rejects further connection requests from that IP address. By default, the threshold is 512 connections from any single IP address. You can change this threshold (to any number between 1 and 250,000) to better suit the requirements of your network environment.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Junos OS supports SYN-ACK-ACK proxy protection for both IPv4 and IPv6 addresses.

Related Documentation

- DoS Attack Overview on page 859
- Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack on page 870

Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack

This example shows how to protect your network against a SYN-ACK-ACK proxy flood attack.

- Requirements on page 870
- Overview on page 870
- Configuration on page 870
- Verification on page 871

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a SYN-ACK-ACK proxy flood. The value unit is connections per source address. The default value is 512 connections from any single address.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 1000-syn-ack-ack-proxy tcp syn-ack-ack-proxy threshold
    1000
set security zones security-zone zone screen 1000-syn-ack-ack-proxy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To protect against a SYN-ACK-ACK proxy flood attack:

1. Specify the source session limits.

```
[edit]
user@host# set security screen ids-option 1000-syn-ack-ack-proxy tcp
    syn-ack-ack-proxy threshold 1000
```

2. Set the security zone for zone screen.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security zones security-zone zone screen 1000-syn-ack-ack-proxy
```

Results From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option 1000-syn-ack-ack-proxy {
    tcp {
    syn-ack-ack-proxy threshold 1000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone zone {
    screen 1000-syn-ack-ack-proxy;
  }
```

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

Verification

Verifying SYN-ACK-ACK Proxy Flood Attack

Purpose Verify SYN-ACK-ACK proxy flood attack.

Action Enter the **show security screen ids-option 1000-syn-ack-ack-proxy** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option 1000-syn-ack-ack-proxy
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  TCP SYN-ACK-ACK proxy threshold           1000
```

```
user@host> show security zones
Security zone: zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 1000-syn-ack-ack-proxy
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows that there is no attack from SYN-ACK-ACK-proxy flood.

Related Documentation

- Understanding SYN-ACK-ACK Proxy Flood Attacks on page 869
- DoS Attack Overview on page 859

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 37

Protecting Against Network DoS Attacks

- Network DoS Attacks Overview on page 873
- Understanding SYN Flood Attacks on page 874
- Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection on page 877
- Example: Enabling SYN Flood Protection for Webservers in the DMZ on page 879
- Understanding Whitelists for SYN Flood Screens on page 885
- Example: Configuring Whitelists for SYN Flood Screens on page 886
- Understanding SYN Cookie Protection on page 887
- Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection on page 890
- Understanding ICMP Flood Attacks on page 893
- Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection on page 894
- Understanding UDP Flood Attacks on page 896
- Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection on page 897
- Understanding Land Attacks on page 899
- Protecting Your Network Against Land Attacks by Enabling Land Attack on page 900
- Understanding Screen IPv6 Tunneling Control on page 902
- Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options on page 905

Network DoS Attacks Overview

A denial-of-service (DoS) attack directed against one or more network resources floods the target with an overwhelming number of SYN, ICMP, or UDP packets or with an overwhelming number of SYN fragments.

Depending on the attackers' purpose and the extent and success of previous intelligence gathering efforts, the attackers might single out a specific host, such as a device or server or they might aim at random hosts across the targeted network. Either approach has the potential of upsetting service to a single host or to the entire network, depending on how critical the role of the victim is to the rest of the network.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- DoS Attack Overview on page 859
- Firewall DoS Attacks Overview on page 861
- OS-Specific DoS Attacks Overview on page 913

Understanding SYN Flood Attacks

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating incomplete connection requests that it can no longer process legitimate connection requests.

Two hosts establish a TCP connection with a triple exchange of packets known as a *three-way handshake*: A sends a SYN segment to B; B responds with a SYN/ACK segment; and A responds with an ACK segment. A SYN flood attack inundates a site with SYN segments containing forged (spoofed) IP source addresses with nonexistent or unreachable addresses. B responds with SYN/ACK segments to these addresses and then waits for responding ACK segments. Because the SYN/ACK segments are sent to nonexistent or unreachable IP addresses, they never elicit responses and eventually time out. See Figure 35.

Figure 35: SYN Flood Attack

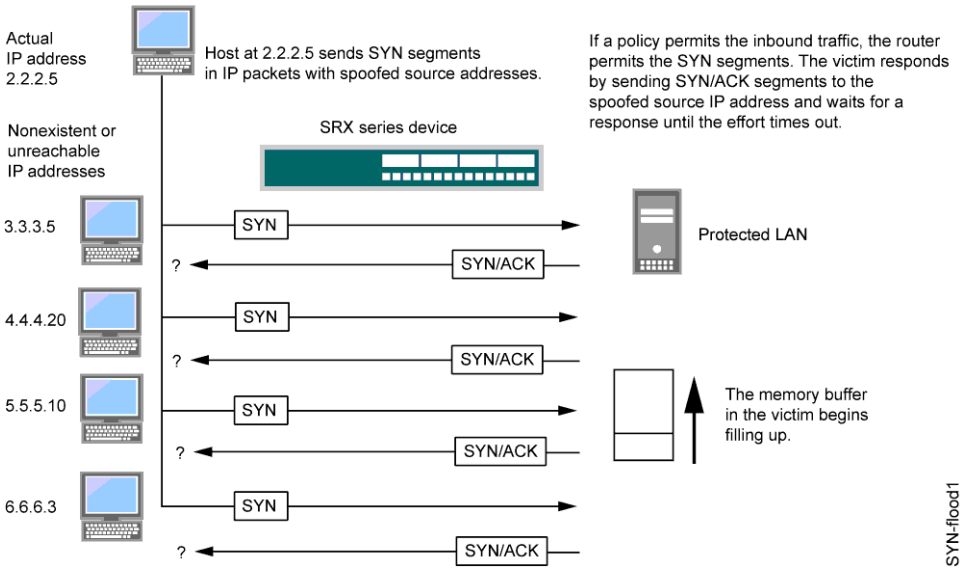

By flooding a host with incomplete TCP connections, the attacker eventually fills the memory buffer of the victim. Once this buffer is full, the host can no longer process new TCP connection requests. The flood might even damage the victim's operating system. Either way, the attack disables the victim and its normal operations.

This topic includes the following sections:

- SYN Flood Protection on page 875
- SYN Flood Options on page 876

Copyright © 2016, Juniper Networks, Inc.

SYN Flood Protection

Junos OS can impose a limit on the number of SYN segments permitted to pass through the firewall per second. You can base the attack threshold on the destination address and ingress interface port, the destination address only, or the source address only. When the number of SYN segments per second exceeds the set threshold, Junos OS will either start proxying incoming SYN segments, replying with SYN/ACK segments and storing the incomplete connection requests in a connection queue, or it will drop the packets.

SYN proxying only happens when a destination address and ingress interface port attack threshold is exceeded. If a destination address or source address threshold is exceeded, additional packets are simply dropped.

In Figure 36, the SYN attack threshold for a destination address and ingress interface port has been exceeded and Junos OS has started proxying incoming SYN segments. The incomplete connection requests remain in the queue until the connection is completed or the request times out.

Figure 36: Proxying SYN Segments

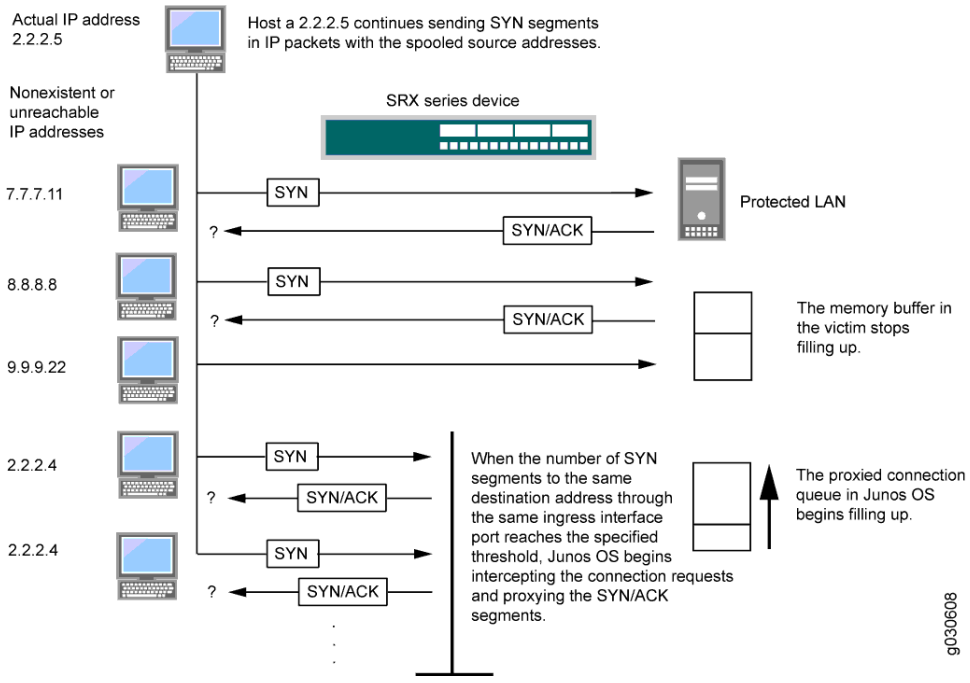

SYN Flood Options

You can set the following parameters for proxying uncompleted TCP connection requests:

- Attack Threshold—This option allows you to set the number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address per second required to activate the SYN proxying mechanism. Although you can set the threshold to any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, if it is an e-business site that normally gets 20,000 SYN segments per second, you might want to set the threshold to 30,000 per second. If a smaller site normally gets 20 SYN segments per second, you might consider setting the threshold to 40.
- Alarm Threshold—This option allows you to set the number of proxied, half-complete TCP connection requests per second after which Junos OS enters an alarm in the event log. The value you set for an alarm threshold triggers an alarm when the number of proxied, half-completed connection requests to the same destination address per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3000 SYN segments to the same destination address per second is required to trigger an alarm entry in the log.

  For each SYN segment to the same destination address in excess of the alarm threshold, the attack detection module generates a message. At the end of the second, the logging module compresses all similar messages into a single log entry that indicates how many SYN segments to the same destination address and port number arrived after exceeding the alarm threshold. If the attack persists beyond the first second, the event log enters an alarm every second until the attack stops.
- Source Threshold—This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address—before Junos OS begins dropping connection requests from that source.

  Tracking a SYN flood by source address uses different detection parameters from tracking a SYN flood by destination address. When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect.
- Destination Threshold—This option allows you to specify the number of SYN segments received per second for a single destination IP address before Junos OS begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number.

  When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

  Consider a case where Junos OS has policies permitting FTP requests and HTTP requests to the same IP address. If the SYN flood attack threshold is 1000 packets per second (pps) and an attacker sends 999 FTP packets and 999 HTTP pps, Junos OS

Copyright © 2016, Juniper Networks, Inc.

treats both FTP and HTTP packets with the same destination address as members of a single set and rejects the 1001st packet—FTP or HTTP—to that destination.

- Timeout—This option allows you to set the maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 1–50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see any dropped connections during normal traffic conditions. Twenty seconds is a very conservative timeout for a three-way handshake ACK response.

NOTE: Junos OS supports SYN flood protection for both IPv4 and IPv6 traffic.

Related Documentation

- Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection on page 877
- *Configuring SYN Flood Protection Options (CLI Procedure)*
- Example: Enabling SYN Flood Protection for Webservers in the DMZ on page 879
- Understanding Whitelists for SYN Flood Screens on page 885
- Example: Configuring Whitelists for SYN Flood Screens on page 886

Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection

This example shows how to protect your network against SYN flood attacks by enabling SYN flood protection.

- Requirements on page 877
- Overview on page 877
- Configuration on page 877
- Verification on page 878

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable the zone-syn-flood protection screen option and set the timeout value to 20. You also specify the zone where the flood might originate.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option zone-syn-flood tcp syn-flood source-threshold 10000
```

Copyright © 2016, Juniper Networks, Inc.

```
set security screen ids-option zone-syn-flood tcp syn-flood destination-threshold 10000
set security zones security-zone untrust screen zone-syn-flood
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To enable SYN flood protection:

1. Specify the screen object name.

```
[edit]
user@host# set security screen ids-option zone-syn-flood tcp syn-flood
  source-threshold 10000
user@host# set security screen ids-option zone-syn-flood tcp syn-flood
  destination-threshold 10000
```

2. Set the security zone for the zone screen.

```
[edit]
user@host# set security zones security-zone untrust screen zone-syn-flood
```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option zone-syn-flood {
    tcp {
    syn-flood {
      source-threshold 10000;
        destination-threshold 10000;
        timeout 20;
      }
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone untrust {
    screen zone-syn-flood;
    interfaces {
      ge-0/0/1.0;
    }
  }
```

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

Verification

Verifying SYN Flood Protection

Purpose

Verify SYN flood protection.

Copyright © 2016, Juniper Networks, Inc.

Action Enter the **show security screen ids-option zone-syn-flood** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option zone-syn-flood
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  TCP SYN flood attack threshold            200
  TCP SYN flood alarm threshold             512
  TCP SYN flood source threshold            10000
  TCP SYN flood destination threshold       10000
  TCP SYN flood timeout                     20
```

```
user@host> show security zones
Security zone: untrust
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: zone-syn-flood
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.0
```

Meaning The sample output shows that SYN flood protection is enabled with source and destination threshold.

Related Documentation

- Understanding SYN Flood Attacks on page 874
- Example: Enabling SYN Flood Protection for Webservers in the DMZ on page 879
- Understanding Whitelists for SYN Flood Screens on page 885
- Example: Configuring Whitelists for SYN Flood Screens on page 886

Example: Enabling SYN Flood Protection for Webservers in the DMZ

This example shows how to enable SYN flood protection for webservers in the DMZ.

- Requirements on page 879
- Overview on page 879
- Configuration on page 882
- Verification on page 885

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

This example shows how to protect four webservers in the DMZ from SYN flood attacks originating in the external zone, by enabling the SYN flood protection screen option for the external zone. See Figure 37.

Copyright © 2016, Juniper Networks, Inc.

NOTE: We recommend that you augment the SYN flood protection that Junos OS provides with device-level SYN flood protection on each webserver. In this example, the webservers are running UNIX, which also provides some SYN flood defenses, such as adjusting the length of the connection request queue and changing the timeout period for incomplete connection requests.

Figure 37: Device-Level SYN Flood Protection

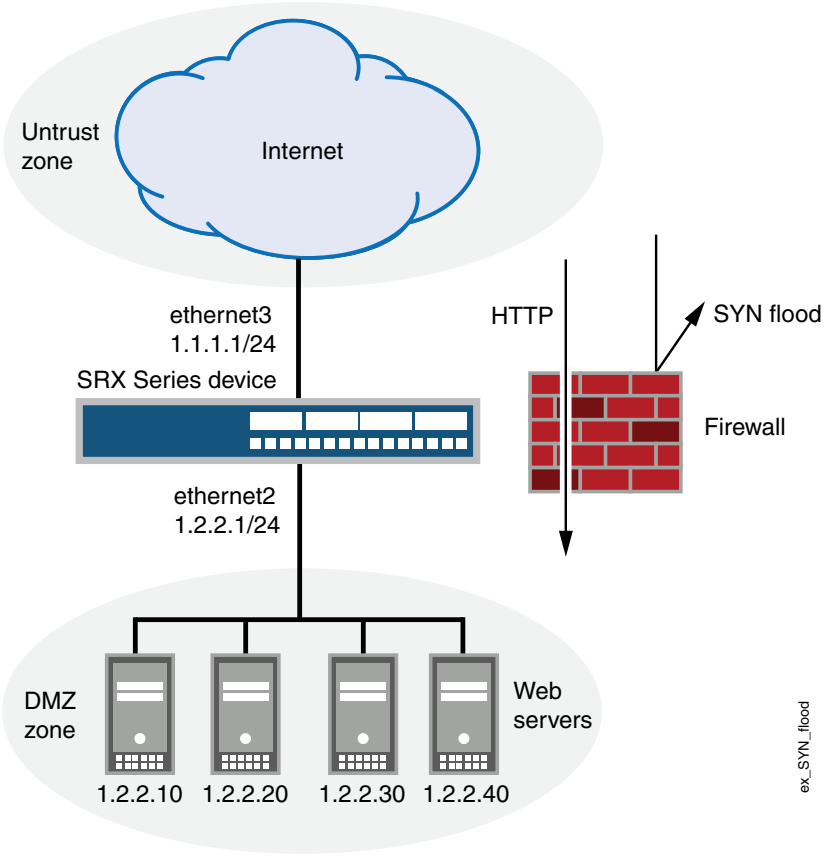

To configure the SYN flood protection parameters with appropriate values for your network, you must first establish a baseline of typical traffic flows. For example, for one week, you run a sniffer on ethernet3—the interface bound to zone_external—to monitor the number of new TCP connection requests arriving every second for the four webservers in the DMZ. Your analysis of the data accumulated from one week of monitoring produces the following statistics:

- Average number of new connection requests per server: 250 per second
- Average peak number of new connection requests per server: 500 per second

Copyright © 2016, Juniper Networks, Inc.

NOTE: A sniffer is a network-analyzing device that captures packets on the network segment to which you attach it. Most sniffers allow you to define filters to collect only the type of traffic that interests you. Later, you can view and evaluate the accumulated information. In this example, you want the sniffer to collect all TCP packets with the SYN flag set arriving at ethernet3 and destined for one of the four webservers in the DMZ. You might want to continue running the sniffer at regular intervals to see whether there are traffic patterns based on the time of day, day of the week, time of the month, or season of the year. For example, in some organizations, traffic might increase dramatically during a critical event. Significant changes probably warrant adjusting the various thresholds.

Based on this information, you set the following SYN flood protection parameters for zone_external as shown in Table 66.

Table 66: SYN Flood Protection Parameters

Parameter	Value	Reason for Each Value
Attack threshold	625 pps	This is 25% higher than the average peak number of new connection requests per second per server, which is unusual for this network environment. When the number of SYN packets per second for any one of the four webservers exceeds this number, the device begins proxying new connection requests to that server. (In other words, beginning with the 626th SYN packet to the same destination address in one second, the device begins proxying connection requests to that address.)
Alarm threshold	250 pps	When the device proxies 251 new connection requests in one second, it makes an alarm entry in the event log. By setting the alarm threshold somewhat higher than the attack threshold, you can avoid alarm entries for traffic spikes that only slightly exceed the attack threshold.
Source threshold	25 pps	When you set a source threshold, the device tracks the source IP address of SYN packets, regardless of the destination address. (Note that this source-based tracking is separate from the tracking of SYN packets based on destination address, which constitutes the basic SYN flood protection mechanism.) In the one week of monitoring activity, you observed that no more than 1/25 of new connection requests for all servers came from any one source within a one-second interval. Therefore, connection requests exceeding this threshold are unusual and provide sufficient cause for the device to execute its proxying mechanism. (Note that 25 pps is 1/25 of the attack threshold, which is 625 pps.) If the device tracks 25 SYN packets from the same source IP address, then, beginning with the 26th packet, it rejects all further SYN packets from that source for the remainder of that second and for the next second as well.
Destination threshold	4000 pps	When you set a destination threshold, the device runs a separate tracking of only the destination IP address, regardless of the destination port number. Because the four webservers receive only HTTP traffic (destination port 80)—no traffic to any other destination port number reaches them—setting another destination threshold offers no additional advantage.

Copyright © 2016, Juniper Networks, Inc.

Table 66: SYN Flood Protection Parameters *(continued)*

Parameter	Value	Reason for Each Value
Timeout	20 seconds	The default value of 20 seconds is a reasonable length of time to hold incomplete connection requests.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 1.2.2.1/24
set interfaces fe-1/0/0 unit 0 family inet address 1.1.1.1/24
set security zones security-zone zone_dmz interfaces ge-0/0/0.0
set security zones security-zone zone_external interfaces fe-1/0/0.0
set security zones security-zone zone_dmz address-book address ws1 1.2.2.10/32
set security zones security-zone zone_dmz address-book address ws2 1.2.2.20/32
set security zones security-zone zone_dmz address-book address ws3 1.2.2.30/32
set security zones security-zone zone_dmz address-book address ws4 1.2.2.40/32
set security zones security-zone zone_dmz address-book address-set web_servers address
  ws1
set security zones security-zone zone_dmz address-book address-set web_servers address
  ws2
set security zones security-zone zone_dmz address-book address-set web_servers address
  ws3
set security zones security-zone zone_dmz address-book address-set web_servers address
  ws4
set security policies from-zone zone_external to-zone zone_dmz policy id_1 match
  source-address any destination-address web_servers application junos-http
set security policies from-zone zone_external to-zone zone_dmz policy id_1 then permit
set security screen ids-option zone_external-syn-flood tcp syn-flood alarm-threshold
  250 attack-threshold 625 source-threshold 25 timeout 20
set security zones security-zone zone_external screen zone_external-syn-flood
```

Step-by-Step Procedure

To configure SYN flood protection parameters:

1. Set interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 1.2.2.1/24
user@host# set interfaces fe-1/0/0 unit 0 family inet address 1.1.1.1/24
user@host# set security zones security-zone zone_dmz interfaces ge-0/0/0.0
user@host# set security zones security-zone zone_external interfaces fe-1/0/0.0
```

2. Define addresses.

```
[edit]
user@host# set security zones security-zone zone_dmz address-book address ws1
  1.2.2.10/32
user@host# set security zones security-zone zone_dmz address-book address ws2
  1.2.2.20/32
user@host# set security zones security-zone zone_dmz address-book address ws3
  1.2.2.30/32
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security zones security-zone zone_dmz address-book address ws4
  1.2.2.40/32
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws1
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws2
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws3
user@host# set security zones security-zone zone_dmz address-book address-set
  web_servers address ws4
```

3. Configure the policy.

```
[edit]
user@host# set security policies from-zone zone_external to-zone zone_dmz policy
  id_1 match source-address any
user@host# set security policies from-zone zone_external to-zone zone_dmz policy
  id_1 match destination-address web_servers
user@host# set security policies from-zone zone_external to-zone zone_dmz policy
  id_1 match application junos-http
user@host# set security policies from-zone zone_external to-zone zone_dmz policy
  id_1 then permit
```

4. Configure the screen options.

```
[edit]
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
  alarm-threshold 250
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
  attack-threshold 625
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
  source-threshold 25
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood
  timeout 20
user@host# set security zones security-zone zone_external screen
  zone_external-syn-flood
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security policies**, and **show security screen** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
  user@host# show interfaces
  ge-0/0/0 {
  unit 0 {
    family inet {
    address 1.2.2.1/24;
    }
  }
}
fe-1/0/0 {
```

Copyright © 2016, Juniper Networks, Inc.

```
unit 0 {
   family inet {
   address 1.1.1.1/24;
   }
}
}
...
[edit]
user@host# show security zones
...
   security-zone zone_dmz {
address-book {
address ws1 1.2.2.10/32;
   address ws2 1.2.2.20/32;
   address ws3 1.2.2.30/32;
   address ws4 1.2.2.40/32;
   address-set web_servers {
   address ws1;
   address ws2;
   address ws3;
   address ws4;
   }
}
interfaces {
   ge-0/0/0.0;
   }
}
security-zone zone_external {
   screen zone_external-syn-flood;
   interfaces {
   fe-1/0/0.0;
}
}
[edit]
user@host# show security policies
from-zone zone_external to-zone zone_dmz {
   policy id_1 {
match {
source-address any;
   destination-address web_servers;
   application junos-http;
   }
then {
permit;
   }
      }
   }
[edit]
   user@host# show security screen
...
ids-option zone_external-syn-flood {
   tcp {
syn-flood {
alarm-threshold 250;
   attack-threshold 625;
   source-threshold 25;
```


Copyright © 2016, Juniper Networks, Inc.

```
    timeout 20;
}
  }
}
```

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

Verification

Verifying SYN Flood Protection for Webservers in the DMZ

Purpose Verify SYN flood protection for webservers in the DMZ.

Action From operational mode, enter the **show interfaces**, **show security zones**, **show security policies**, and **show security screen ids-option zone_external-syn-flood** commands.

Related Documentation

- Understanding SYN Flood Attacks on page 874
- Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection on page 877
- *Configuring SYN Flood Protection Options (CLI Procedure)*

Understanding Whitelists for SYN Flood Screens

Junos OS provides the administrative option to configure a whitelist of trusted IP addresses to which the SYN flood screen will not reply with a SYN/ACK. Instead, the SYN packets from the source addresses or to the destination addresses in the list are allowed to bypass the SYN cookie and SYN proxy mechanisms. This feature is needed when you have a service in your network that cannot tolerate proxied SYN/ACK replies under any condition, including a SYN flood event.

Both IP version 4 (IPv4) and IP version 6 (IPv6) whitelists are supported. Addresses in a whitelist should be all IPv4 or all IPv6. In each whitelist, there can be up to 32 IP address prefixes. You can specify multiple addresses or address prefixes as a sequence of addresses separated by spaces and enclosed in square brackets.

A whitelist can cause high CPU usage on a central point depending on the traffic level. For example, when no screen is enabled, the connections per second (cps) is 492K; when the screen is enabled and the whitelist is disabled, the cps is 373K; and when both the screen and the whitelist are enabled, the cps is 194K. After enabling the whitelist, the cps drops by 40 percent.

Related Documentation

- Understanding SYN Flood Attacks on page 874
- Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection on page 877
- Example: Configuring Whitelists for SYN Flood Screens on page 886

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Whitelists for SYN Flood Screens

This example shows how to configure whitelists of IP addresses to be exempted from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process.

- Requirements on page 886
- Overview on page 886
- Configuration on page 886
- Verification on page 887

Requirements

Before you begin, configure a security screen and enable the screen in the security zone. See "Example: Enabling SYN Flood Protection for Webservers in the DMZ" on page 879.

Overview

In this example, you configure whitelists named **wlipv4** and **wlipv6**. All addresses are IP version 4 (IPv4) for **wlipv4**, and all addresses are IP version 6 (IPv6) for **wlipv6**. Both whitelists include destination and source IP addresses.

Multiple addresses or address prefixes can be configured as a sequence of addresses separated by spaces and enclosed in square brackets, as shown in the configuration of the destination addresses for **wlipv4**.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option js1 tcp syn-flood white-list wlipv4 source-address 1 .1.1.0/24
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
  2.2.2.2/32
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
  3.3.3.3/32
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
  4.4.4.4/32
set security screen ids-option js1 tcp syn-flood white-list wlipv6 source-address 2001::1/64
set security screen ids-option js1 tcp syn-flood white-list wlipv6 destination-address
  2002::1/64
```

Step-by-Step Procedure

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

To configure the whitelists:

1. Specify the name of the whitelist and the IP addresses to be exempted from the SYN/ACK.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security screen ids-option js1 tcp syn-flood]
user@host# set white-list wlipv4 source-address 1.1.1.0/24
user@host# set white-list wlipv4 destination-address [2.2.2.2 3.3.3.3 4.4.4.4]
user@host# set white-list wlipv6 source-address 2001::1/64
user@host# set white-list wlipv6 destination-address 2002::1/64
```

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

```
[edit]
user@host# show security screen
ids-option js1 {
  tcp {
    syn-flood {
      white-list wlipv4 {
        source-address 1.1.1.0/24;
        destination-address [ 2.2.2.2/32 3.3.3.3/32 4.4.4.4/32 ];
      }
      white-list wlipv6 {
        source-address 2001::1/64;
        destination-address 2002::1/64;
      }
    }
  }
}
```

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

Verification

Verifying Whitelist Configuration

Purpose Verify that the whitelist is configured properly.

Action From operational mode, enter the **show security screen ids-option** command.

Related Documentation

- Understanding SYN Flood Attacks on page 874
- Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection on page 877
- Understanding Whitelists for SYN Flood Screens on page 885

Understanding SYN Cookie Protection

SYN cookie is a stateless SYN proxy mechanism you can use in conjunction with other defenses against a SYN flood attack.

As with traditional SYN proxying, SYN cookie is activated when the SYN flood attack threshold is exceeded. However, because SYN cookie is stateless, it does not set up a session or policy and route lookups upon receipt of a SYN segment, and it maintains no

Copyright © 2016, Juniper Networks, Inc.

connection request queues. This dramatically reduces CPU and memory usage and is the primary advantage of using SYN cookie over the traditional SYN proxying mechanism.

When SYN cookie is enabled on Junos OS and becomes the TCP-negotiating proxy for the destination server, it replies to each incoming SYN segment with a SYN/ACK containing an encrypted cookie as its initial sequence number (ISN). The cookie is an MD5 hash of the original source address and port number, destination address and port number, and ISN from the original SYN packet. After sending the cookie, Junos OS drops the original SYN packet and deletes the calculated cookie from memory. If there is no response to the packet containing the cookie, the attack is noted as an active SYN attack and is effectively stopped.

If the initiating host responds with a TCP packet containing the cookie +1 in the TCP ACK field, Junos OS extracts the cookie, subtracts 1 from the value, and recomputes the cookie to validate that it is a legitimate ACK. If it is legitimate, Junos OS starts the TCP proxy process by setting up a session and sending a SYN to the server containing the source information from the original SYN. When Junos OS receives a SYN/ACK from the server, it sends ACKs to the server and to the initiation host. At this point the connection is established and the host and server are able to communicate directly.

NOTE: **The use of SYN cookie or SYN proxy enables the SRX Series device to protect the TCP servers behind it from SYN flood attacks in IPv6 flows.**

Figure 38 shows how a connection is established between an initiating host and a server when SYN cookie is active on Junos OS.

Copyright © 2016, Juniper Networks, Inc.

Figure 38: Establishing a Connection with SYN Cookie Active

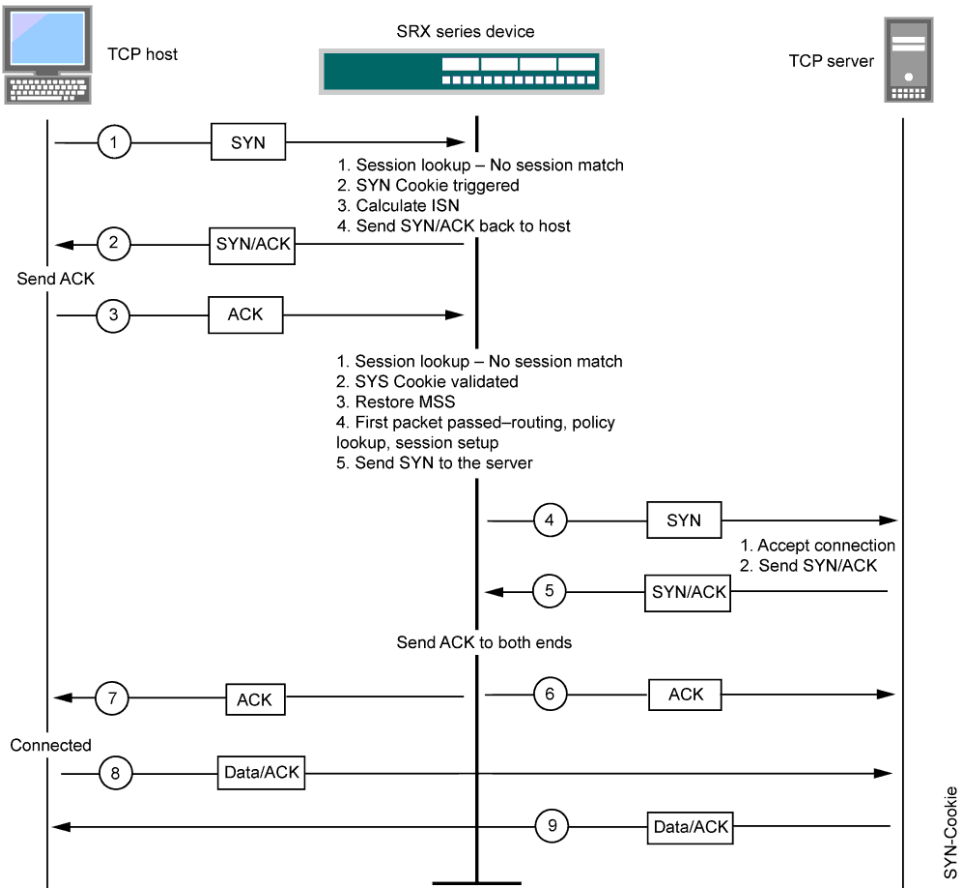

SYN Cookie Options

You can set the following parameters for incomplete TCP proxy connection requests:

- Attack Threshold—This option allows you to set the number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address and port number per second required to activate the SYN proxy mechanism. Although you can set the threshold to any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, for an e-business site that normally gets 2000 SYN segments per second, you might want to set the threshold to 30,000 SYN segments per second. The valid threshold range is 1 to 1,000,000. For a smaller site that normally gets 20 SYN segments per second, you might consider setting the threshold to 40 SYN segments per second.
- Alarm Threshold—This option allows you to set the number of proxied, half-complete TCP connection requests per second after which Junos OS enters an alarm in the event log. The alarm threshold value you set triggers an alarm when the number of proxied, half-completed connection requests to the same destination address and port number per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3001 SYN segments to the same destination address and port number per second is required to

Copyright © 2016, Juniper Networks, Inc.

trigger an alarm entry in the log. The valid threshold range is 1 to 1,000,000 and the default alarm threshold value is 512.

- Source Threshold—This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address and port number—before Junos OS begins dropping connection requests from that source.

  When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect. The valid threshold range is 4 to 1,000,000 and the default alarm threshold value is 4000.

- Destination Threshold—This option allows you to specify the number of SYN segments received per second for a single destination IP address before Junos OS begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number. The valid threshold range is 4 to 1,000,000 and the default alarm threshold value is 4000.

  When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

- Timeout—This option allows you to set the maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 0 to 50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see dropped connections during normal traffic conditions.

  When either a source or destination threshold is not configured, the system will use the default threshold value. The default source and destination threshold value is 4000 pps.

Related Documentation

- Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection on page 890
- DoS Attack Overview on page 859

Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection

This example shows how to detect and protect your network against SYN flood attacks by enabling the SYN cookie protection.

- Requirements on page 891
- Overview on page 891
- Configuration on page 891
- Verification on page 892

Copyright © 2016, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set the external-syn-flood timeout value to 20 and set the security zone for external screen to external-syn-flood. Also, you set the protection mode to syn-cookie.

NOTE: The SYN cookie feature can detect and protect only against spoofed SYN flood attacks, minimizing the negative impact on hosts that are secured by Junos OS. If an attacker uses a legitimate IP source address, rather than a spoofed IP source, then the SYN cookie mechanism does not stop the attack.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option external-syn-flood tcp syn-flood timeout 20
set security zones security-zone external screen external-syn-flood
set security flow syn-flood-protection-mode syn-cookie
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To enable the SYN cookie protection:

1. Specify the external-syn-flood timeout value.

   ```
   [edit]
   user@host# set security screen ids-option external-syn-flood tcp syn-flood timeout
     20
   ```

2. Set the security-zone for external screen.

   ```
   [edit]
   user@host# set security zones security-zone external screen external-syn-flood
   ```

3. Set the protection mode.

   ```
   [edit]
   user@host# set security flow syn-flood-protection-mode syn-cookie
   ```

Results

From configuration mode, confirm your configuration by entering the **show security screen**, **show security zones**, and **show security flow** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security screen
  set security flow syn-flood-protection-mode syn-cookie {
    tcp {
    syn-flood {
      source-ip-based 1;
      }
    }
  }

[edit]
user@host# show security zones
  security-zone external {
    screen external-syn-flood;
  }
[edit]
user@host# show security flow
  syn-flood-protection-mode syn-cookie;
```

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

Verification

Verifying SYN Cookie Protection

Purpose Verifying SYN cookie protection.

Action Enter the **show security screen ids-option external-syn-flood** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option external-syn-flood
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  TCP SYN flood attack threshold            200
  TCP SYN flood alarm threshold             512
  TCP SYN flood source threshold            4000
  TCP SYN flood destination threshold       4000
  TCP SYN flood timeout                     20

user@host> show security zones
Security zone: external
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: external-syn-flood
  Interfaces bound: 0
  Interfaces:

user@host> show security zones
Security zone: external
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: external-syn-flood
  Interfaces bound: 0
  Interfaces:
```


Copyright © 2016, Juniper Networks, Inc.

Meaning The sample output shows that SYN cookie protection is enabled with a source and destination threshold.

Related Documentation
- Understanding SYN Cookie Protection on page 887
- DoS Attack Overview on page 859

Understanding ICMP Flood Attacks

An ICMP flood typically occurs when ICMP echo requests overload the target of the attack with so many requests that the target expends all its resources responding until it can no longer process valid network traffic.

NOTE: ICMP messages generated in flow mode are limited to 12 messages every 10 seconds. This rate limit is calculated on a per-CPU basis. Once the threshold is reached, no further acknowledgement messages are sent to the device.

When enabling the ICMP flood protection feature, you can set a threshold that, once exceeded, invokes the ICMP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, Junos OS ignores further ICMP echo requests for the remainder of that second plus the next second as well. See Figure 39.

NOTE: An ICMP flood can consist of any type of ICMP message. Therefore, Junos OS monitors all ICMP message types, not just echo requests.

Copyright © 2016, Juniper Networks, Inc.

Figure 39: ICMP Flooding

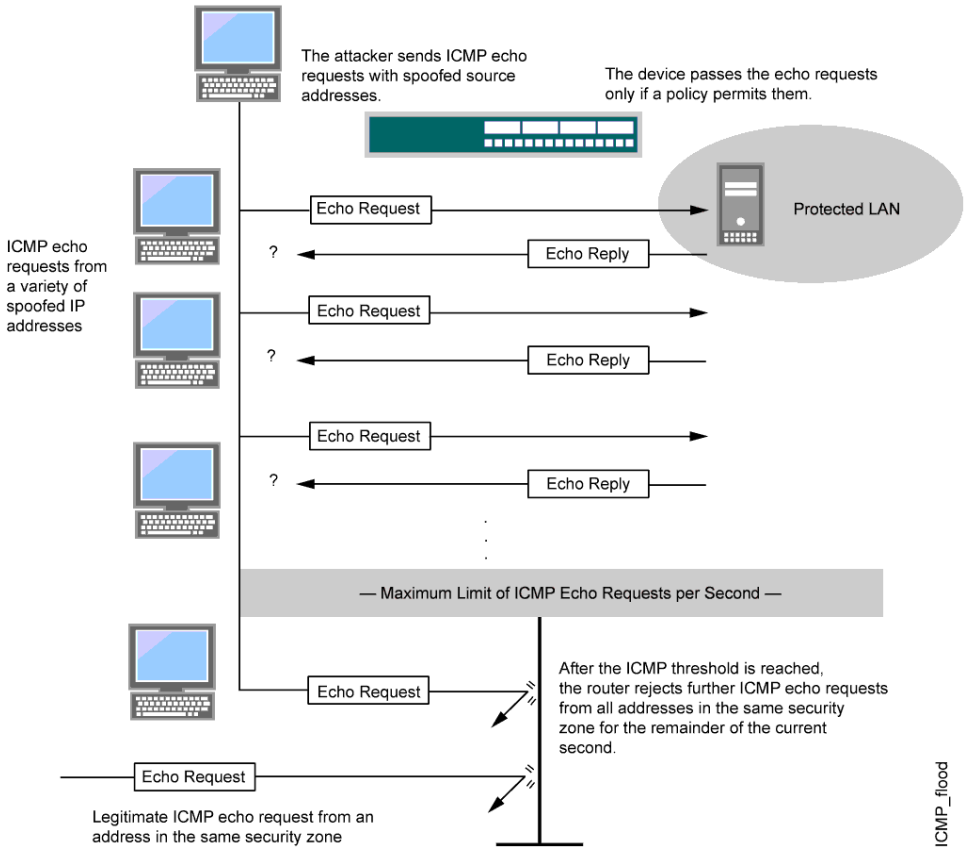

NOTE: Junos OS supports ICMP flood protection for both IPv4 and IPv6 traffic.

Related Documentation

- Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection on page 894
- DoS Attack Overview on page 859

Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection

This example shows how to protect your network against ICMP flood attacks by enabling ICMP flood protection.

- Requirements on page 895
- Overview on page 895
- Configuration on page 895
- Verification on page 896

Copyright © 2016, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before enabling ICMP flood protection.

Overview

In this example, you enable ICMP flood protection. The value unit is ICMP packets per second, or pps. The default value is 1000 pps. You specify the zone where a flood might originate.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 1000-icmp-flood icmp flood threshold 1000
set security zones security-zone zone screen 1000-icmp-flood
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To enable ICMP flood protection:

1. Specify the ICMP flood threshold value.

   ```
   [edit]
   user@host# set security screen ids-option 1000-icmp-flood icmp flood threshold
     1000
   ```

2. Set the security zone for zone screen.

   ```
   [edit]
   user@host# set security zones security-zone zone screen 1000-icmp-flood
   ```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option 1000-icmp-flood {
    icmp {
    flood threshold 1000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone zone {
    screen 1000-icmp-flood;
  }
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying ICMP Flood Protection

Purpose Verify ICMP flood protection

Action Enter the **show security screen ids-option 1000-icmp-flood** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option 1000-icmp-flood
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                          Value
  ICMP flood threshold                        1000
```

```
user@host> show security zones
Security zone: zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 1000-icmp-flood
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows that ICMP flood protection is enabled and threshold is set.

Related Documentation

- Understanding ICMP Flood Attacks on page 893
- DoS Attack Overview on page 859

Understanding UDP Flood Attacks

Similar to an ICMP flood, a UDP flood occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the victim to the point that the victim can no longer handle valid connections.

After enabling the UDP flood protection feature, you can set a threshold that, once exceeded, invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second, or pps.) If the number of UDP datagrams from one or more sources to a single destination exceeds this threshold, Junos OS ignores further UDP datagrams to that destination for the remainder of that second plus the next second as well. See Figure 40.

NOTE: The high-end SRX Series devices do not drop the packet in the next second.

Copyright © 2016, Juniper Networks, Inc.

Figure 40: UDP Flooding

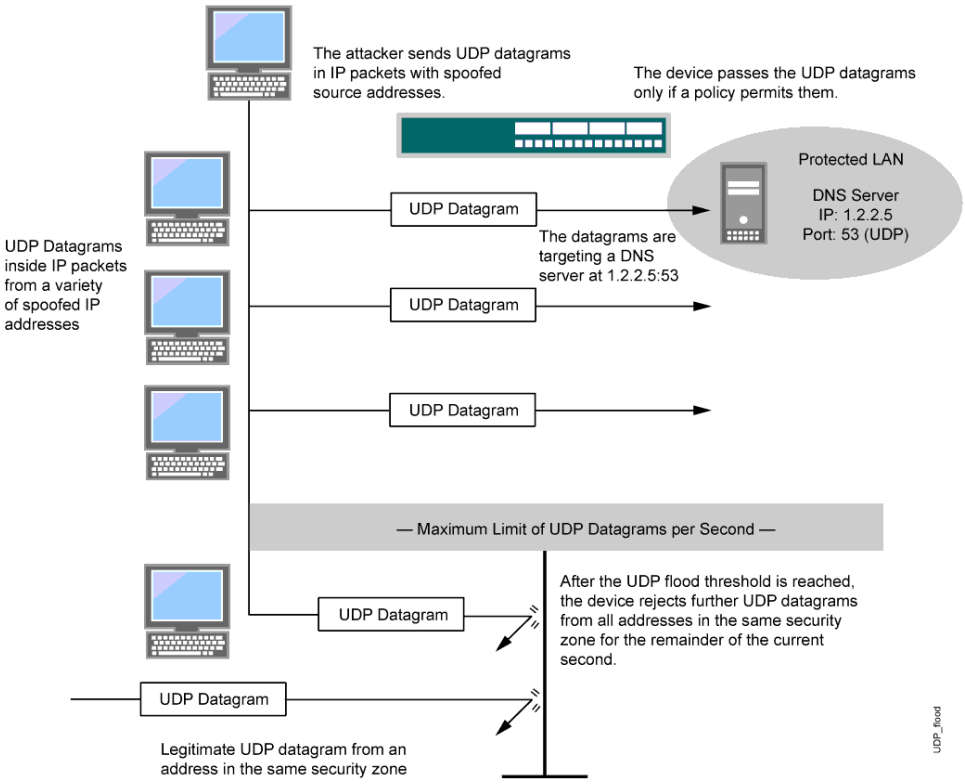

NOTE: Junos OS supports UDP flood protection for IPV4 and IPv6 packets.

Related Documentation

- Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection on page 897
- DoS Attack Overview on page 859

Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection

This example shows how to protect your network against UDP flood attacks by enabling UDP flood protection.

- Requirements on page 897
- Overview on page 898
- Configuration on page 898
- Verification on page 899

Requirements

No special configuration beyond device initialization is required before enabling UDP flood protection.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you enable UDP flood protection. The value unit is UDP packets per second, or pps. The default value is 1000 pps. You specify the zone where a flood might originate.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 1000-udp-flood udp flood threshold 1000
set security zones security-zone external screen 1000-udp-flood
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To enable UDP flood protection:

1. Specify the UDP flood threshold value.

```
[edit]
user@host# set security screen ids-option 1000-udp-flood udp flood threshold
  1000
```

2. Set the security zone for external screen.

```
[edit]
user@host# set security zones security-zone external screen 1000-udp-flood
```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option 1000-udp-flood {
    udp {
    flood threshold 1000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone external {
    screen 1000-udp-flood;
  }
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying UDP Flood Protection

Purpose Verify UDP flood protection.

Action Enter the **show security screen ids-option 1000-udp-flood** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option 1000-udp-flood
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  UDP flood threshold                       1000
```

```
user@host> show security zones
Security zone: external
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 1000-udp-flood
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows that UDP flood protection is enabled and threshold is set.

Related Documentation

- Understanding UDP Flood Attacks on page 896
- DoS Attack Overview on page 859

Understanding Land Attacks

Combining a SYN attack with IP spoofing, a land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and the source IP address.

The receiving system responds by sending the SYN-ACK packet to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, causing a denial of service (DoS). See Figure 41.

Copyright © 2016, Juniper Networks, Inc.

Figure 41: Land Attack

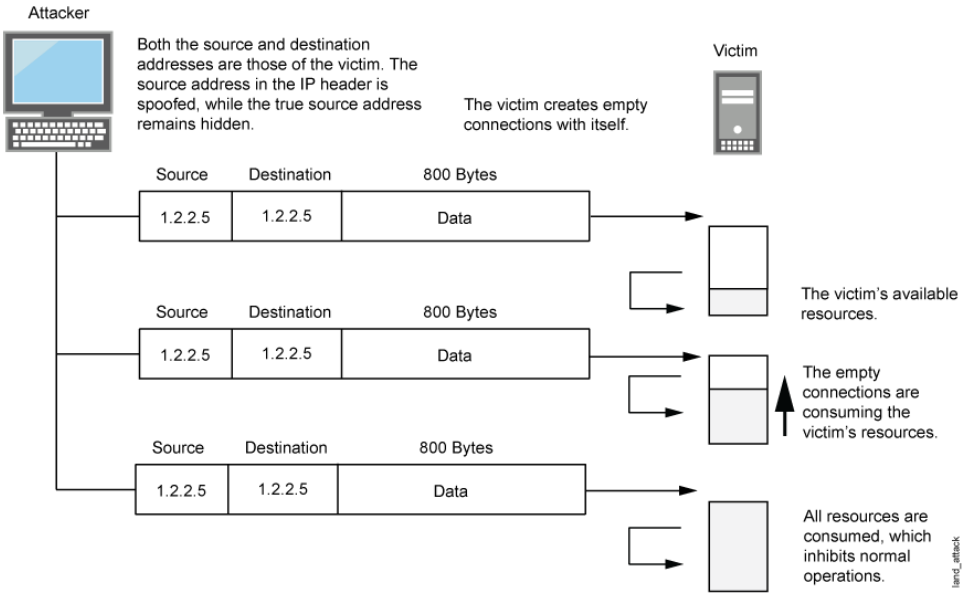

When you enable the screen option to block land attacks, Junos OS combines elements of the SYN flood defense and IP spoofing protection to detect and block any attempts of this nature.

NOTE: Junos OS supports land attack protection for both IPv4 and IPv6 packets.

Related Documentation

- Protecting Your Network Against Land Attacks by Enabling Land Attack on page 900
- DoS Attack Overview on page 859

Protecting Your Network Against Land Attacks by Enabling Land Attack

This example shows how to protect your network against attacks by enabling land attack.

- Requirements on page 900
- Overview on page 901
- Configuration on page 901
- Verification on page 901

Requirements

No special configuration beyond device initialization is required before enabling land attack.

Copyright © 2016, Juniper Networks, Inc.

Overview

This example shows how to enable protection against a land attack. In this example, you set the security screen object name as land and set the security zone as zone.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option land tcp land
set security zones security-zone zone screen land
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To enable protection against a land attack:

1. Specify the screen object name.

   ```
   [edit]
   user@host# set security screen ids-option land tcp land
   ```

2. Set the security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone screen land
   ```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen
  ids-option land {
    tcp {
    land;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone zone {
    screen land;
  }
```

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

Verification

Verifying Protection Against a Land Attack

Purpose

Verify protection against a land attack.

Copyright © 2016, Juniper Networks, Inc.

Action Enter the **show security screen ids-option land** and **show security zones** commands from operational mode.

```
user@host> show security screen ids-option land
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  TCP land attack                           enabled
```

```
user@host> show security zones
Security zone: zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: land
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows that protection against a land attack is enabled.

Related Documentation

- Understanding Land Attacks on page 899
- DoS Attack Overview on page 859

Understanding Screen IPv6 Tunneling Control

Several IPv6 transition methodologies are provided to utilize the tunneling of IPv6 packets over IPv4 networks that do not support IPv6. For this reason, these methods use public gateways and bypass the policies of the operator.

The security of tunneled packets is a major concern for service providers, because tunneled packets are easily accessed by attackers. Numerous IPv6 transition methodologies have evolved for sending tunneled packets through a network; however, because some of them operate on public gateways, they bypass the policies of the operator. This means that packet transmission is exposed to attackers. To overcome and secure transfer of packets, the IPv6 end nodes are required to de-capsulate the encapsulated data packets. Screen is one of the latest available technologies for blocking or allowing tunneling traffic based on user preferences.

You can configure the following screen options to check and filter packets based on IPv6 extension headers, packet headers, and Bad-Inner-Header IPv6 or IPv4 address validation. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IP tunneling.

- **GRE 4in4 Tunnel**: The GRE 4in4 Tunnel screen matches the following signature: **| IPv4 outer header | GRE header | IPv4 inner header**

  An outer IPv4 header must be **Protocol 47 GRE Encapsulation**. A GRE header must have **protocol E-type 0x0800 IPv4**. If these conditions are met, this packet is classified as GRE 4in4 tunnel signature.

- **GRE 4in6 Tunnel**: The GRE 4in6 Tunnel screen matches the following signature: **IPv6 outer main header | IPv6 extension header(s) | GRE header | IPv4 inner header**

Copyright © 2016, Juniper Networks, Inc.

An outer IPv6 main header or an IPv6 extension header must have a **Next Header of value 47 for GRE**. A GRE header must have **protocol E-type 0x0800 IPv4**. If these conditions are met, this packet is classified as GRE 4in6 tunnel signature.

- **GRE 6in4 Tunnel**: The GRE 6in4 Tunnel screen matches the following signature: **IPv4 outer header | GRE header | IPv6 inner header**

  An outer IPv4 header must be **Protocol 47 GRE Encapsulation**. A GRE header must have **protocol E-type 0x086DD IPv6** . If these conditions are met, this packet is classified as GRE 6in4 tunnel signature.

- **GRE 6in6 Tunnel**: The GRE 6in6 Tunnel screen matches the following signature: **IPv6 outer main header | IPv6 extension header(s) | GRE header | IPv6 inner header**

  An outer IPv6 main header or an IPv6 extension header must have a Next Header of **value 47 for GRE**. A GRE header must have **protocol E-type 0x086DD` IPv6**. If these conditions are met, this packet is classified as GRE 6in6 tunnel signature.

- **IPinIP 6to4relay Tunnel** : The IPinIP 6to4relay Tunnel screen matches the following signature: **| IPv4 outer header | IPv6 inner header**

  An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. An outer header source address or destination address must be in network **192.88.99.0/24**. An inner IPv6 header source address or destination address must be in network **2002: /16**. If these conditions are met, this packet is classified as IPinIP 6to4relay tunnel signature.

- **IPinIP 6in4 Tunnel** : The IPinIP 6in4 Tunnel screen matches the following signature: **| IPv4 outer header | IPv6 inner header**

  An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. If this condition is met, this packet is classified as IPinIP 6in4 tunnel signature.

  NOTE: **Typically, when IPv6 packets need to be transported in a complete IPv4 network, the IPv6 packets utilizes a point-to-point 6in4 tunnel.**

- **IPinIP 6over4 Tunnel** : The IPinIP 6over4 Tunnel screen matches the following signature: **| IPv4 outer header | IPv6 inner header**

  An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation:W**. An inner header source address or destination address must be in **fe80::/64** network. If these conditions are met, this packet is classified as IPinIP 6over4 tunnel signature.

- **IPinIP 4in6 Tunnel** : The IPinIP 4in6 Tunnel screen matches the following signature: **| IPv6 outer main header | IPv6 extension header(s) | IPv4 inner header**

  An outer IPv6 header or an IPv6 extension header must have a Next Header of **value 04 for IPv4**. If these conditions are met, this packet is classified as IPinIP 4in6 tunnel signature.

- **IPinIP ISATAP Tunnel**: The IPinIP ISATAP Tunnel screen matches the following signature: **| IPv6 outer main header | IPv6 inner header**

  An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. An inner IPv6 header source address or destination address must be in **fe80::200:5efe/96 or fe80::5efe/96**

Copyright © 2016, Juniper Networks, Inc.

network. If these conditions are met, this packet is classified as IPinIP ISATAP tunnel signature.

- **IPinIP DS-Lite Tunnel**: The IPinIP DS-Lite Tunnel screen matches the following signature: **| IPv6 outer main header | IPv6 extension header(s) | IPv4 inner header**

  An outer IPv6 header or an IPv6 extension header must have a Next Header of **value 04 for IPv4**. An inner IPv4 source address or destination address must be in **192.0.0.0/29** network. If these conditions are met, this packet is classified as IPinIP DS-Lite tunnel signature.

- **IPinIP 6in6 Tunnel**: The IPinIP 6in6 Tunnel screen matches the following signature: **| IPv6 outer main header | IPv6 extension header(s) | IPv6 inner main header**

  An outer IPv6 main header or an IPv6 extension header must have a Next Header of **value 41 for IPv6**. An inner IPv6 main header must be **Version 6**. If these two conditions are met, this packet is classified as IPinIP 6in6 tunnel signature.

- **IPinIP 4in4 Tunnel**: The IPinIP 4in4 Tunnel screen matches the following signature: **| IPv6 outer header | IPv4 inner header** . An outer IPv4 header must have a Protocol of **value 04 for IPv4**. An inner IPv4 header must be **Version 4**.

- **IPinUDP Teredo Tunnel**: The IPinUDP Teredo Tunnel matches the following signature: **IPv4 outer header | UDP header | IPv6 inner header**

  An outer IPv4 header must have a **Protocol of 17 for UDP payload**. A UDP header source or destination port must be **3544**. An inner IPv6 header source address or destination address must be in network **2001:0000:/32**.

- **IP Tunnel Bad Inner-Header Check**: The Bad Inner Header Tunnel screen checks the tunnel traffic inner header information for consistency. The packet drops when any of the following is detected:
  - Inner header does not match outer header.
  - Inner header TTL or Hop Limit must not be 0 or 255.
  - Inner header IPv6 address checking.
  - Inner header IPv4 address checking.
  - Outer and Inner header length checks:
  - Inner header IPv4 and IPv6 TCP/UDP/ICMP header length check:

    TCP/UDP/ICMP header length must fit inside of inner IPv4/IPv6/EH6 header length when inner IP(v4/v6) is not a first, next, or last fragment.
  - TCP: The minimum TCP header size must fit in the previous encapsulation length.
  - ICMP: The minimum ICMP header size must fit in the previous encapsulation length.
  - Fragmented packets: For fragmented packets, if the tunnel information needs to be checked for a screen and is not in the first fragment, then checking is not performed except the parts of the tunnel encapsulation that are included in the first fragment. Length checks are performed on first fragment packets using the actual packet buffer length, but the length checks are ignored because the inner header is larger than the outer header.

Copyright © 2016, Juniper Networks, Inc.

- When the outer header is first fragment, do not examine the past physical packet length of the fragment.
- When the inner header is a first fragment, do not examine the past length of the fragment.

For non-first fragment packets, checking is not performed in Bad Inner Header Tunnel screen.

- When outer header is a non-first fragment, examine the packet for screens that only use IP header signatures, because the payload cannot be examined.
- When inner header is a non-first fragment, do not examine the next packet.
- The IPv4 inner header checks that IPv4 header is from 20 to 50 bytes.

Starting with Junos OS Release 12.3X48-D10, the syslog messages **RT_SCREEN_IP** and **RT_SCREEN_IP_LS** for the IP tunneling screen have been updated to include the tunnel screen attacks and log-without-drop criteria. The following list illustrates some examples of these new system log messages for each of the tunnel types:

- **RT_SCREEN_IP: Tunnel GRE 6in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: alarm-without-drop**
- **RT_SCREEN_IP: Tunnel GRE 6in6! source: 1212::12, destination: 1111::11, zone name: untrust, interface name: ge-0/0/1.0, action: drop**
- **RT_SCREEN_IP: Tunnel GRE 4in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: drop**
- **RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 6in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: alarm-without-drop**
- **RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 6in6! source: 1212::12, destination: 1111::11, zone name: untrust, interface name: ge-0/0/1.0, action: drop**
- **RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 4in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: drop**

Related Documentation

- Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options on page 905
- Network DoS Attacks Overview on page 873
- Understanding Land Attacks on page 899

Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options

This example shows how to configure the tunnel screens to enable the screens to control, allow, or block the transit of tunneled traffic.

- Requirements on page 906
- Overview on page 906

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 907
- Verification on page 910

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Junos OS Release 12.3X48-D10 and later

Before you begin:

- Understand the IPv6 Tunneling control. See "Understanding Screen IPv6 Tunneling Control" on page 902.

Overview

You can configure the following IP tunneling screen options to check and filter packets, based on IPv6 extension headers, packet headers, and bad-inner-header IPv6 or IPv4 address validation. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IP tunneling. The following tunneling screen options are assigned to an untrust zone.

- GRE 4in4 Tunnel
- GRE 4in6 Tunnel
- GRE 6in4 Tunnel
- GRE 6in6 Tunnel
- IPinUDP Teredo Tunnel
- IPinIP 4in4 Tunnel
- IPinIP 4in6 Tunnel
- IPinIP 6in4 Tunnel
- IPinIP 6in6 Tunnel
- IPinIP 6over4 Tunnel
- IPinIP 6to4relay Tunnel
- IPinIP ISATAP Tunnel
- IPinIP DS-Lite Tunnel
- Bad Inner Header Tunnel

Copyright © 2016, Juniper Networks, Inc.

Configuration

To configure the IP tunneling screen options, perform these tasks:

- Configuring GRE Tunnel Screens on page 907
- Configuring an IPinUDP Teredo Tunnel Screen on page 907
- Configuring an IPinIP Tunnel Screen on page 908
- Configuring a Bad-Inner-Header Tunnel Screen on page 909
- Results on page 909

Configuring GRE Tunnel Screens

CLI Quick Configuration

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

```
set security screen ids-option screen1 ip tunnel gre gre-4in4
set security screen ids-option screen1 ip tunnel gre gre-4in6
set security screen ids-option screen1 ip tunnel gre gre-6in4
set security screen ids-option screen1 ip tunnel gre gre-6in6
set security zones security-zone untrust screen screen1
```

Step-by-Step Procedure

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

To configure a GRE tunnel screen:

1. Configure a GRE tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

   ```
   [edit security screen ids-option screen1 ip tunnel gre]
   user@host# set gre-4in4
   user@host# set gre-4in6
   user@host# set gre-6in4
   user@host# set gre gre-6in6
   ```

2. Configure the screens in the security zones.

   ```
   user@host#set security zones security-zone untrust screen screen1
   ```

Configuring an IPinUDP Teredo Tunnel Screen

CLI Quick Configuration

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

```
set security screen ids-option screen1 ip tunnel ip-in-udp teredo
set security zones security-zone untrust screen screen1
```

Step-by-Step Procedure

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

To configure an IPinUDP Teredo tunnel screen:

1. Configure an IPinUDP Teredo tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

```
[edit security screen ids-option screen1 ip tunnel]
user@host# set ip-in-udp teredo
```

2. Configure the screens in the security zones.

```
user@host# set security zones security-zone untrust screen screen1
```

Configuring an IPinIP Tunnel Screen

CLI Quick Configuration

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

```
set security screen ids-option screen1 ip tunnel ipip dslite
set security screen ids-option screen1 ip tunnel ipip ipip-4in4
set security screen ids-option screen1 ip tunnel ipip ipip-4in6
set security screen ids-option screen1 ip tunnel ipip ipip-6in4
set security screen ids-option screen1 ip tunnel ipip ipip-6in6
set security screen ids-option screen1 ip tunnel ipip ipip-6over4
set security screen ids-option screen1 ip tunnel ipip ipip-6to4relay
set security screen ids-option screen1 ip tunnel ipip isatap
set security zones security-zone untrust screen screen1
```

Step-by-Step Procedure

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

To configure an IPinIP tunnel screen:

1. Configure an IPinIP tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

```
[edit security screen ids-option screen1 ip tunnel ipip]
user@host# set dslite
user@host# set ipip-4in4
user@host# set ipip-4in6
user@host# set ipip-6in4
user@host# set ipip-6in6
user@host# set ipip-6over4
user@host# set ipip-6to4relay
user@host# set ipip-isatap
```

2. Configure the screens in the security zones.

```
user@host# set security zones security-zone untrust screen screen1
```


Copyright © 2016, Juniper Networks, Inc.

Configuring a Bad-Inner-Header Tunnel Screen

CLI Quick Configuration

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

```
set security screen ids-option screen1 ip tunnel bad-inner-header
set security zones security-zone untrust screen screen1
```

Step-by-Step Procedure

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

To configure a bad-inner-header tunnel screen:

1. Configure a bad-inner-header tunnel screen to check the tunnel traffic inner header information for consistency.

   ```
   [edit security screen ids-option screen1 ip tunnel]
   user@host# set bad-inner-header
   ```

2. Configure the screens in the security zones.

   ```
   user@host# set security zones security-zone untrust screen screen1
   ```

Results

From configuration mode, confirm your configuration by entering the **show security screen** and **show security screen statistics zone untrust ip tunnel** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
  user@host# show security screen
  ...
  ids-option screen1 {
  ip{
  tunnel {
    gre {
      gre-4in4;
      gre-4in6;
      gre-6in4;
      gre-6in6;
    }
    ip-in-udp {
      teredo;
    }
    ipip {
      ipip-4in4;
      ipip-4in6;
```

Copyright © 2016, Juniper Networks, Inc.

```
            ipip-6in4;
            ipip-6in6;
            ipip-6over4;
            ipip-6to4relay;
            isatap;
            dslite;
          }
          bad-inner-header;
        }
        }
      }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Security Screen Configuration on page 910
- Verifying IP Tunnel Screens in the Security Zones on page 910

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option screen1** command.

```
user@host> show security screen ids-option screen1
show security screen ids-option screen1:
   Name                                Value
  IP Tunnel Bad Inner Header          enabled
  IP Tunnel GRE 6in4                  enabled
  IP Tunnel GRE 4in6                  enabled
  IP Tunnel GRE 6in6                  enabled
  IP Tunnel GRE 4in4                  enabled
  IP Tunnel IPinUDP Teredo            enabled
  IP Tunnel IPIP 6to4 Relay           enabled
  IP Tunnel IPIP 6in4                 enabled
  IP Tunnel IPIP 6over4               enabled
  IP Tunnel IPIP 4in6                 enabled
  IP Tunnel IPIP 4in4                 enabled
  IP Tunnel IPIP 6in6                 enabled
  IP Tunnel IPIP ISATAP               enabled
  IP Tunnel IPIP DS-Lite              enabled
```

Meaning The **show security screen ids-option screen1** command displays screen object status as enabled.

Verifying IP Tunnel Screens in the Security Zones

Purpose Verify that the IP tunneling screen options are configured properly in the security zones.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security screen statistics zone untrust ip tunnel** command.

```
user@host> show security screen statistics zone untrust ip tunnel

IP Tunnel Screen statistics:

  IDS attack type                    Statistics
  IP tunnel GRE 6in4                         0
  IP tunnel GRE 4in6                         0
  IP tunnel GRE 6in6                         0
  IP tunnel GRE 4in4                         0
  IP tunnel IPIP 6to4 relay                  0
  IP tunnel IPIP 6in4                        0
  IP tunnel IPIP 6over4                      0
  IP tunnel IPIP 4in6                        0
  IP tunnel IPIP 4in4                        0
  IP tunnel IPIP 6in6                        0
  IP tunnel IPIP ISATAP                      0
  IP tunnel IPIP DS-Lite                     0
  IP tunnel IPinUDP Teredo                   0
  IP tunnel bad inner header                 0
```

Meaning The **show security screen statistics zone untrust ip tunnel** command displays the IP tunnel screen statistics summary.

Related Documentation

- Understanding Screen IPv6 Tunneling Control on page 902
- Network DoS Attacks Overview on page 873
- Understanding Land Attacks on page 899

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 38

Protecting Against OS-Specific DoS Attacks

- OS-Specific DoS Attacks Overview on page 913
- Understanding Ping of Death Attacks on page 913
- Example: Protecting Against a Ping of Death Attack on page 914
- Understanding Teardrop Attacks on page 915
- Example: Protecting Against a Teardrop Attack on page 917
- Understanding WinNuke Attacks on page 917
- Example: Protecting Against a WinNuke Attack on page 919

OS-Specific DoS Attacks Overview

If an attacker not only identifies the IP address and responsive port numbers of an active host but also its operating system (OS), instead of resorting to brute-force attacks, the attacker can launch more elegant attacks that can produce one-packet or two-packet "kills."

OS-specific denial-of-service (DoS) attacks, including ping of death attacks, teardrop attacks, and WinNuke attacks, can cripple a system with minimal effort. If Junos OS is protecting hosts susceptible to these attacks, you can configure Junos OS to detect these attacks and block them before they reach their target.

Related Documentation

- Understanding Ping of Death Attacks on page 913
- DoS Attack Overview on page 859

Understanding Ping of Death Attacks

OS-specific DoS attacks, such as ping of death attacks, can cripple a system with minimal effort.

The maximum allowable IP packet size is 65,535 bytes, including the packet header, which is typically 20 bytes. An ICMP echo request is an IP packet with a pseudo header, which is 8 bytes. Therefore, the maximum allowable size of the data area of an ICMP echo request is 65,507 bytes (65,535 - 20 - 8 = 65,507).

Copyright © 2016, Juniper Networks, Inc.

However, many ping implementations allow the user to specify a packet size larger than 65,507 bytes. A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting.

When you enable the ping of death screen option, Junos OS detects and rejects such oversized and irregular packet sizes even when the attacker hides the total packet size by fragmenting it. See Figure 42.

NOTE: For information about IP specifications, see RFC 791, *Internet Protocol*. For information about ICMP specifications, see RFC 792, *Internet Control Message Protocol*. For information about ping of death attacks, see http://www.insecure.org/sploits/ping-o-death.html.

Figure 42: Ping of Death

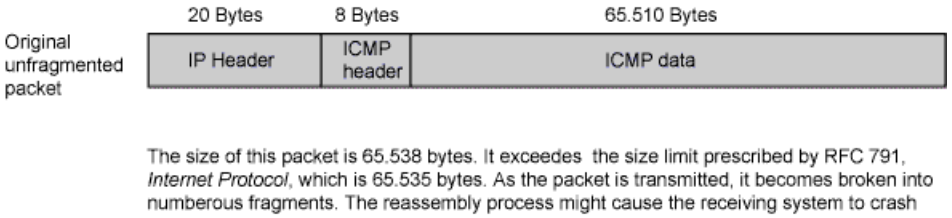

NOTE: Junos OS supports ping of death protection for both IPv4 and IPv6 packets.

Related Documentation

- Example: Protecting Against a Ping of Death Attack on page 914
- DoS Attack Overview on page 859

Example: Protecting Against a Ping of Death Attack

This example shows how to protect against a ping-of-death attack.

- Requirements on page 914
- Overview on page 914
- Configuration on page 915
- Verification on page 915

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a ping-of-death attack and specify the zone where the attack originates.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To enable protection against a ping of death:

1. Specify the screen object name.

   ```
   [edit]
   user@host# set security screen ids-option ping-death icmp ping-death
   ```

2. Set the security zone for zone screen.

   ```
   [edit]
   user@host# set security zones security-zone zone screen ping-death
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen ids-option ping-death** and **show security zones** commands in operational mode.

Related Documentation

- Understanding Ping of Death Attacks on page 913
- DoS Attack Overview on page 859

Understanding Teardrop Attacks

OS-specific denial-of-service (DoS) attacks, such as teardrop attacks, can cripple a system with minimal effort.

Teardrop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the fields is the fragment offset field, which indicates the starting position, or offset, of the data contained in a fragmented packet relative to the data of the original unfragmented packet. See Figure 43.

Copyright © 2016, Juniper Networks, Inc.

Figure 43: Teardrop Attacks

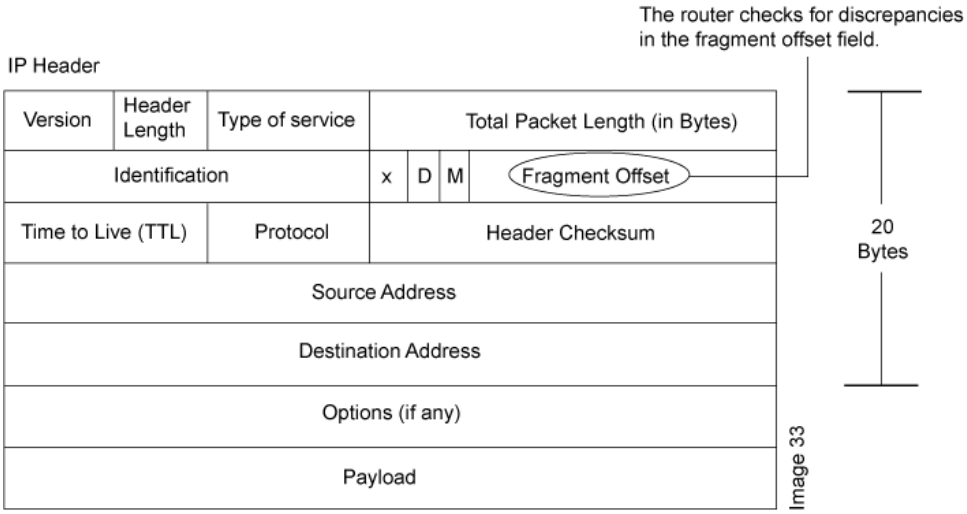

When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash, especially if it is running an older OS that has this vulnerability. See Figure 44.

Figure 44: Fragment Discrepancy

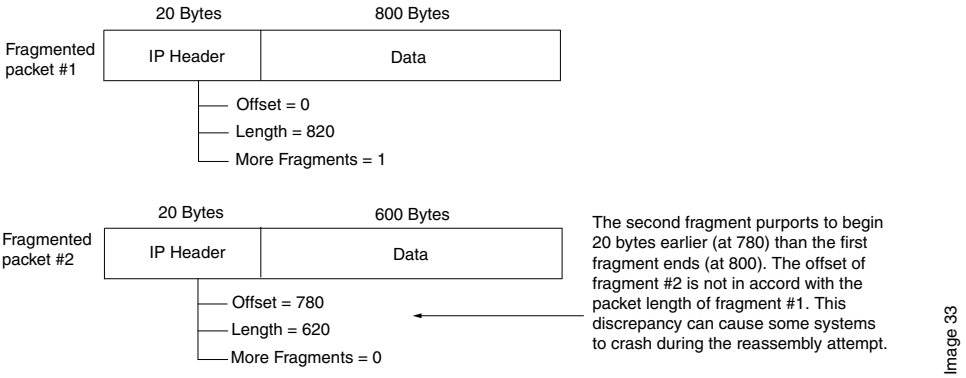

After you enable the teardrop attack screen option, whenever Junos OS detects this discrepancy in a fragmented packet, it drops it.

NOTE: Junos OS supports teardrop attack prevention for both IPv4 and IPv6 packets.

Related Documentation

- Example: Protecting Against a Teardrop Attack on page 917
- DoS Attack Overview on page 859

Copyright © 2016, Juniper Networks, Inc.

Example: Protecting Against a Teardrop Attack

This example shows how to protect against a teardrop attack.

- Requirements on page 917
- Overview on page 917
- Configuration on page 917
- Verification on page 917

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a teardrop attack and also specify the zone where the attack originates.

Configuration

Step-by-Step Procedure

To enable protection against teardrop attack:

1. Specify the screen name.

   ```
   [edit]
   user@host# set security screen ids-option tear-drop ip tear-drop
   ```

2. Associate the screen with a security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone screen tear-drop
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen ids-option tear-drop** and **show security zones** commands in operational mode.

Related Documentation

- Understanding Teardrop Attacks on page 915
- DoS Attack Overview on page 859

Understanding WinNuke Attacks

OS-specific denial-of-service (DoS) attacks, such as WinNuke attacks, can cripple a system with minimal effort.

Copyright © 2016, Juniper Networks, Inc.

WinNuke is a DoS attack targeting any computer on the Internet running Windows. The attacker sends a TCP segment—usually to NetBIOS port 139 with the urgent (URG) flag set—to a host with an established connection (see Figure 45). This introduces a NetBIOS fragment overlap, which causes many machines running Windows to crash. After the attacked machine is rebooted, the following message appears, indicating that an attack has occurred:

```
An exception OE has occurred at 0028:[address] in VxD MSTCP(01) +
000041AE. This was called from 0028:[address] in VxD NDIS(01) +
00008660. It may be possible to continue normally.
Press any key to attempt to continue.
Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information in
   all applications.
Press any key to continue.
```

Figure 45: WinNuke Attack Indicators

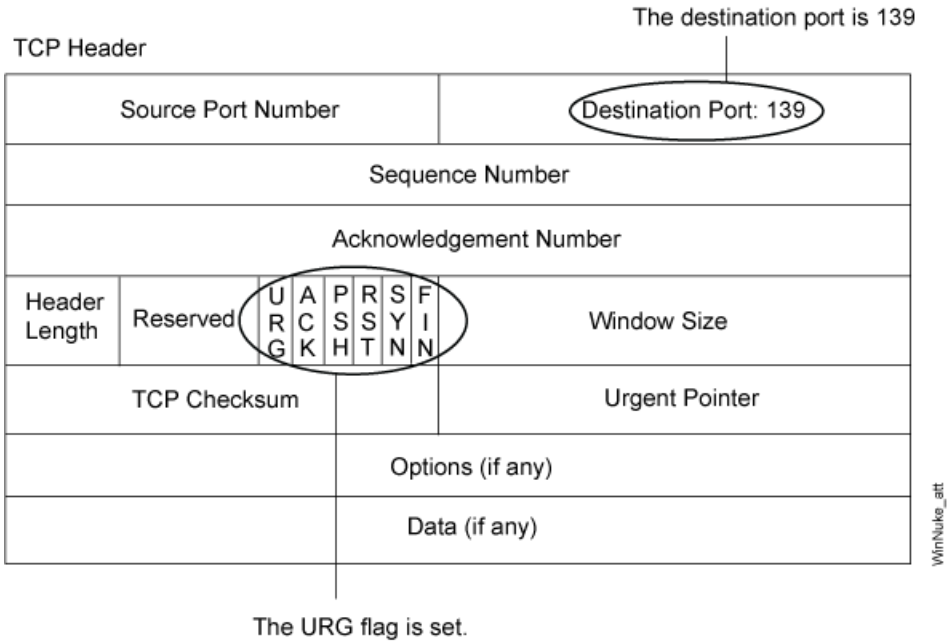

If you enable the WinNuke attack defense screen option, Junos OS scans any incoming Microsoft NetBIOS session service (port 139) packets. If Junos OS observes that the URG flag is set in one of those packets, it unsets the URG flag, clears the URG pointer, forwards the modified packet, and makes an entry in the event log noting that it has blocked an attempted WinNuke attack.

NOTE: Junos OS supports WinNuke attack protection for both IPv4 and IPv6 traffic.

Related Documentation

- Example: Protecting Against a WinNuke Attack on page 919
- DoS Attack Overview on page 859

Copyright © 2016, Juniper Networks, Inc.

Example: Protecting Against a WinNuke Attack

This example shows how to protect against a WinNuke attack.

- Requirements on page 919
- Overview on page 919
- Configuration on page 919
- Verification on page 919

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a WinNuke attack and specify the zone where the attack originates.

Configuration

Step-by-Step Procedure

To enable protection against WinNuke attack:

1. Specify the screen name.

   ```
   [edit]
   user@host# set security screen ids-option winnuke tcp winnuke
   ```

2. Associate the screen with a security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone screen winnuke
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen ids-option winnuke** and **show security zones** commands in operational mode.

Related Documentation

- Understanding WinNuke Attacks on page 917
- DoS Attack Overview on page 859

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 7

Configuring Reconnaissance Deterrence for Security Devices

- Protecting Against IP Sweep and Port Options on page 923
- Protecting Against System Probes and Flag Set on page 937
- Protecting Against Attacker Evasion Techniques on page 947

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 39

Protecting Against IP Sweep and Port Options

- Reconnaissance Deterrence Overview on page 923
- Understanding IP Address Sweeps on page 923
- Example: Blocking IP Address Sweeps on page 924
- Understanding TCP Port Scanning on page 926
- Enhancing Traffic Management by Blocking Port Scans on page 927
- Understanding UDP Port Scanning on page 930
- Understanding Network Reconnaissance Using IP Options on page 930
- Example: Detecting Packets That Use IP Screen Options for Reconnaissance on page 933

Reconnaissance Deterrence Overview

Attackers can better plan their attack when they first know the layout of the targeted network (which IP addresses have active hosts), the possible entry points (which port numbers are active on the active hosts), and the constitution of their victims (which operating system the active hosts are running). To gain this information, attackers must perform reconnaissance.

Juniper Networks provides several screen options for deterring attackers' reconnaissance efforts and thereby hindering them from obtaining valuable information about the protected network and network resources.

Related Documentation

- Understanding Network Reconnaissance Using IP Options on page 930

Understanding IP Address Sweeps

An address sweep occurs when one source IP address sends a defined number of ICMP packets sent to different hosts within a defined interval (5000 microseconds is the default). The purpose of this attack is to send ICMP packets—typically echo requests—to various hosts in the hopes that at least one replies, thus uncovering an address to target.

Junos OS internally logs the number of ICMP packets to different addresses from one remote source. Using the default settings, if a remote host sends ICMP traffic to 10

Copyright © 2016, Juniper Networks, Inc.

addresses in 0.005 seconds (5000 microseconds), then the device flags this as an address sweep attack and rejects all further ICMP packets from that host for the remainder of the specified threshold time period. See Figure 46.

Figure 46: Address Sweep

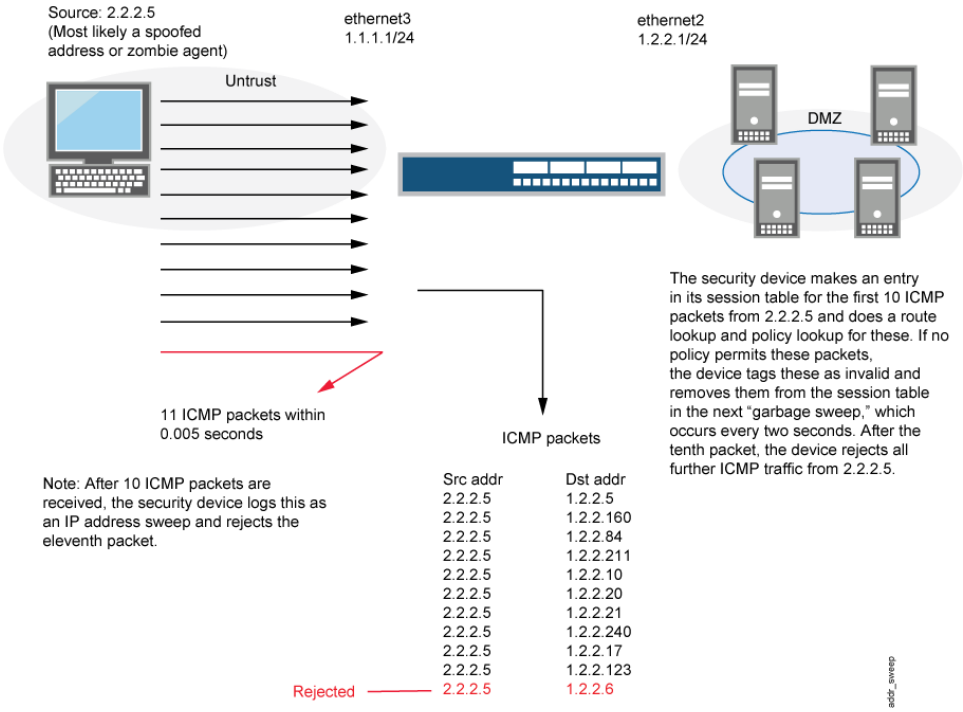

Consider enabling this screen option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable the screen option. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.

NOTE: Junos OS supports this screen option for ICMPv6 trafffic also.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking IP Address Sweeps

This example describes how to configure a screen to block an IP address sweep originating from a security zone.

- Requirements on page 925
- Overview on page 925
- Configuration on page 925
- Verification on page 925

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Understand how IP address sweeps work. See "Understanding IP Address Sweeps" on page 923.
- Configure security zones. See "Security Zones and Interfaces Overview" on page 1049.

Overview

You need to enable a screen for a security zone if you have configured a policy that permits ICMP traffic from that zone. If you have not configured such a policy, then your system denies all ICMP traffic from that zone, and the attacker cannot perform an IP address sweep successfully anyway.

In this example you configure a **5000-ip-sweep** screen to block IP address sweeps originating in the zone-1 security zone.

Configuration

Step-by-Step Procedure

To configure a screen to block IP address sweeps:

1. Configure a screen.

   ```
   [edit]
   user@host# set security screen ids-option 5000-ip-sweep icmp ip-sweep threshold
     5000
   ```

2. Enable the screen in the security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone-1 screen 5000-ip-sweep
   ```

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

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 925
- Verifying the Security Screen Configuration on page 926

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
```

Copyright © 2016, Juniper Networks, Inc.

```
Policy configurable: Yes
Screen: 5000-ip-sweep
 Interfaces bound: 1
Interfaces:
  ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** ***screen-name*** command.

```
[edit]
user@host> show security screen ids-option 5000-ip-sweep
Screen object status:

Name                                    Value
ICMP address sweep threshold            5000
```

Understanding TCP Port Scanning

A port scan occurs when one source IP address sends IP packets containing TCP SYN segments to 10 different destination ports within a defined interval (5000 microseconds is the default). The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.

Junos OS internally logs the number of different ports scanned from one remote source. Using the default settings, if a remote host scans 10 ports in 0.005 seconds (5000 microseconds), then the device flags this as a port scan attack and rejects all further packets from the remote source, regardless of the destination IP address, for the remainder of the specified timeout period. See Figure 47.

Copyright © 2016, Juniper Networks, Inc.

Figure 47: Port Scan

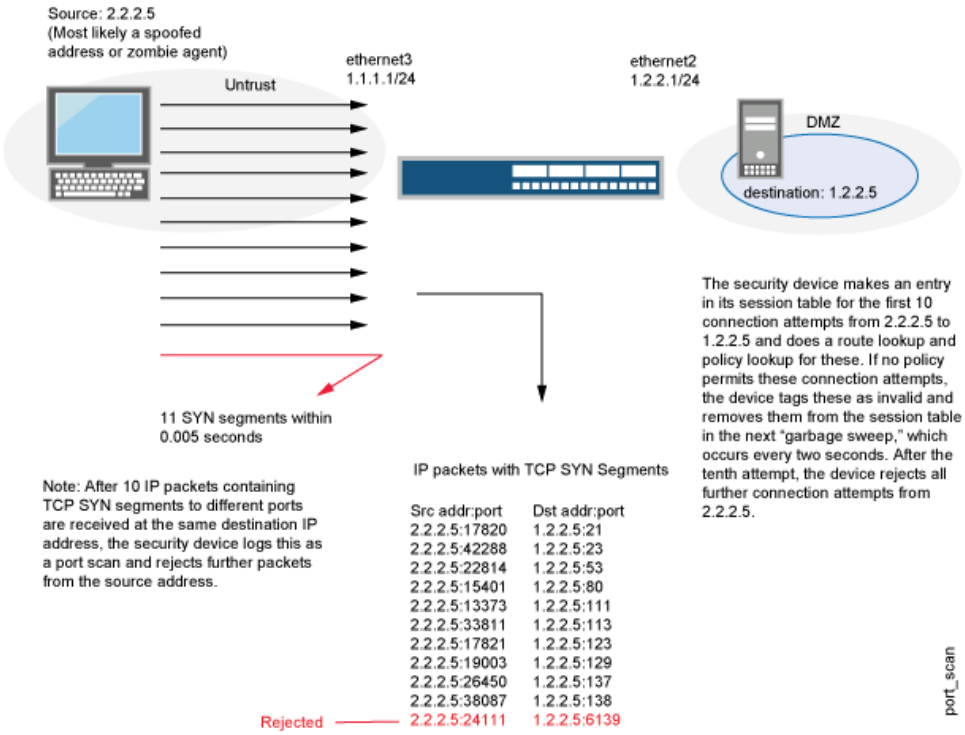

NOTE: Junos OS supports port scanning for both IPv4 and IPv6 traffic.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Enhancing Traffic Management by Blocking Port Scans

This example shows how to enhance traffic management by configuring a screen to block port scans originating from a particular security zone.

- Requirements on page 928
- Overview on page 928
- Configuration on page 928
- Verification on page 929

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, understand how port scanning works. See "Understanding TCP Port Scanning" on page 926.

Overview

You can use a port scan to block IP packets containing TCP SYN segments or UDP segments sent to different ports from the same source address within a defined interval. The purpose of this attack is to scan the available services in the hopes that at least one port will respond. Once a port responds, it is identified as a service to target.

In this example, you configure a 5000 port-scan screen to block port scans originating from a particular security zone and then assign the screen to the zone called zone-1.

Configuration

CLI Quick Configuration

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

```
set security screen ids-option 5000-port-scan tcp port-scan threshold 5000
set security screen ids-option 10000-port-scan udp port-scan threshold 10000
set security zones security-zone zone-1 screen 5000-port-scan
```

Step-by-Step Procedure

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

To configure a screen to block port scans:

1. Configure the screen.

```
[edit security]
user@host# set security screen ids-option 5000-port-scan tcp port-scan threshold
  5000
user@host#set security screen ids-option 10000-port-scan udp port-scan threshold
  10000
```

2. Enable the screen in the security zone.

```
[edit security]
user@host# set security zones security-zone zone-1 screen 5000-port-scan
```

Results

From configuration mode, confirm your configuration by entering the **show security screen ids-option 5000-port-scan** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security screen ids-option 5000-port-scan
tcp {
```


Copyright © 2016, Juniper Networks, Inc.

```
    port-scan threshold 5000;
  }
  udp {
    port-scan threshold 10000;
  }

  [edit]
  user@host# show security zones
  security-zone zone-1 {
    screen 5000-port-scan;
  }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 929
- Verifying the Security Screen Configuration on page 929

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 5000-port-scan
  Interfaces bound: 0
  Interfaces:
```

Meaning The sample output shows that the screen for zone-1 is enabled for port scan blocking.

Verifying the Security Screen Configuration

Purpose Verify the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** *screen-name* command.

```
[edit]
user@host> show security screen ids-option 5000-port-scan
Screen object status:
Name                                        Value
  TCP port scan threshold                   5000
  UDP port scan threshold                   10000
```

Meaning The sample output shows that the port scan blocking is operational with TCP and UDP threshold.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding TCP Port Scanning on page 926

Understanding UDP Port Scanning

UDP port scan gives statistical information on a session threshold. As the incoming packets traverse the screen, the sessions are established. The number of sessions threshold enforced is based on zone, source IP, and the threshold period and does not allow more than 10 new sessions in the configured threshold period, for each zone and source IP address. The UDP port scan is disabled by default. When the UDP port scan is enabled, the default threshold period is 5000 microseconds. This value can be manually set to a range of 1000-1,000,000 microseconds. This feature protects some exposed public UDP services against DDoS attacks. See Figure 48.

Figure 48: UDP Port Scan

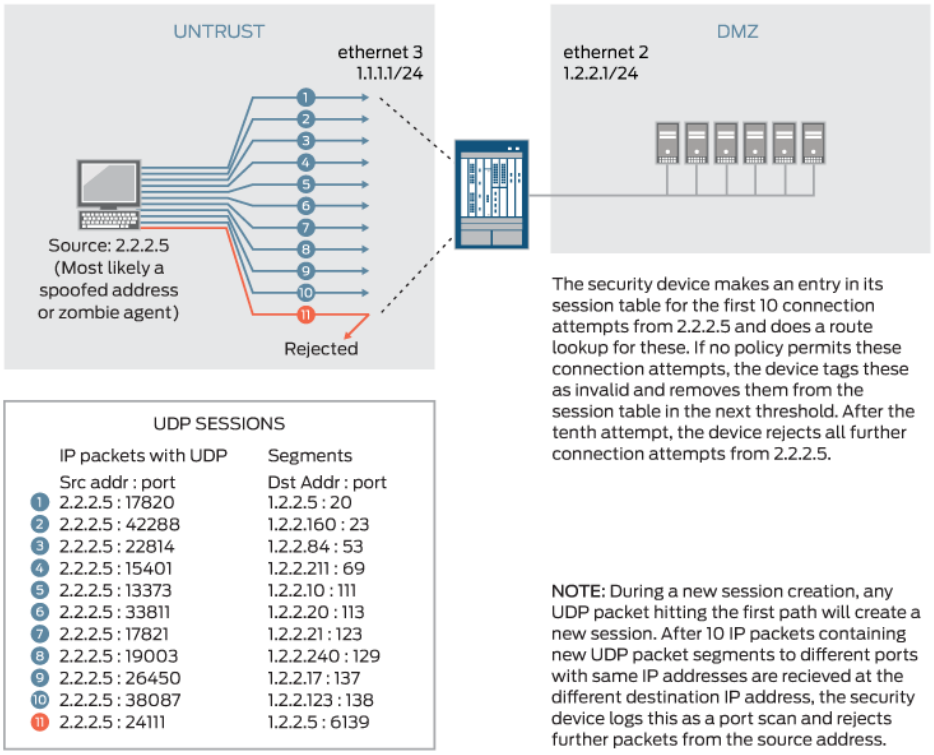

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Understanding Network Reconnaissance Using IP Options

The IP standard RFC 791, *Internet Protocol*, specifies a set of options for providing special routing controls, diagnostic tools, and security.

Copyright © 2016, Juniper Networks, Inc.

RFC 791 states that these options are “unnecessary for the most common communications” and, in reality, they rarely appear in IP packet headers. These options appear after the destination address in an IP packet header, as shown in Figure 49. When they do appear, they are frequently being put to some illegitimate use.

Figure 49: Routing Options

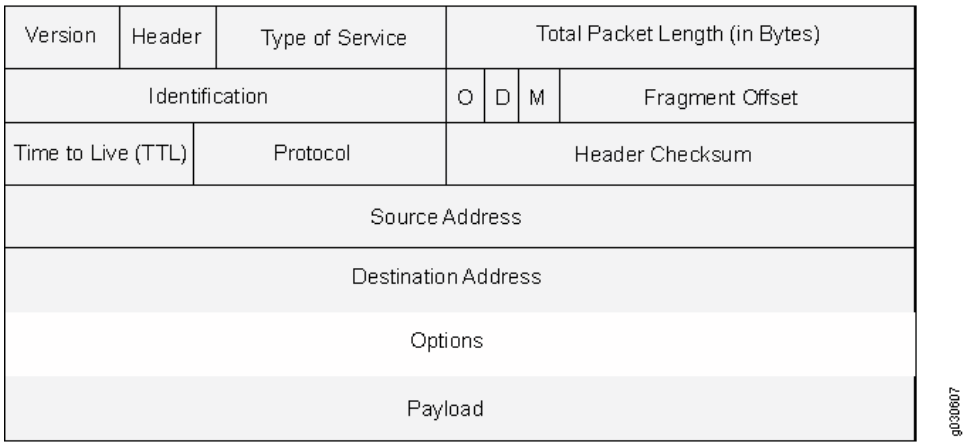

This topic contains the following sections:

- Uses for IP Packet Header Options on page 931
- Screen Options for Detecting IP Options Used for Reconnaissance on page 933

Uses for IP Packet Header Options

Table 67 lists the IP options and their accompanying attributes.

Table 67: IP Options and Attributes

Type	Class	Number	Length	Intended Use	Nefarious Use
End of Options	0*	0	0	Indicates the end of one or more IP options.	None.
No Options	0	1	0	Indicates there are no IP options in the header.	None.
Security	0	2	11 bits	Provides a way for hosts to send security, TCC (closed user group) parameters, and Handling Restriction Codes compatible with Department of Defense (DoD) requirements. (This option, as specified in RFC 791, *Internet Protocol,* and RFC 1038, *Revised IP Security Option,* is obsolete.) Currently, this screen option is applicable only to IPv4.	Unknown. However, because it is obsolete, its presence in an IP header is suspect.

Copyright © 2016, Juniper Networks, Inc.

Table 67: IP Options and Attributes *(continued)*

Type	Class	Number	Length	Intended Use	Nefarious Use
Loose Source Route	0	3	Varies	Specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.	Evasion. The attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network.
Record Route	0	7	Varies	Records the IP addresses of the network devices along the path that the IP packet travels. The destination machine can then extract and process the route information. (Due to the size limitation of 40 bytes for both the option and storage space, this can only record up to 9 IP addresses.) Currently, this screen option is applicable only to IPv4.	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.
Stream ID	0	8	4 bits	(Obsolete) Provided a way for the 16-bit SATNET stream identifier to be carried through networks that did not support the stream concept. Currently, this screen option is applicable only to IPv4.	Unknown. However, because it is obsolete, its presence in an IP header is suspect.
Strict Source Route	0	9	Varies	Specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable only to IPv4.	Evasion. An attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network.
Timestamp	2**	4		Records the time (in coordinated universal time [UTC]***) when each network device receives the packet during its trip from the point of origin to its destination. The network devices are identified by IP address. This option develops a list of IP addresses of the devices along the path of the packet and the duration of transmission between each one. Currently, this screen option is applicable only to IPv4.	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet has passed.

Copyright © 2016, Juniper Networks, Inc.

Table 67: IP Options and Attributes *(continued)*

Type	Class	Number	Length	Intended Use	Nefarious Use

* The class of options identified as 0 was designed to provide extra packet or network control.

** The class of options identified as 2 was designed for diagnostics, debugging, and measurement.

*** The timestamp uses the number of milliseconds since midnight UTC. UTC is also known as Greenwich Mean Time (GMT), which is the basis for the international time standard.

Screen Options for Detecting IP Options Used for Reconnaissance

The following screen options detect IP options that an attacker can use for reconnaissance or for some unknown but suspect purpose:

- Record Route—Junos OS detects packets where the IP option is 7 (record route) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- Timestamp—Junos OS detects packets where the IP option list includes option 4 (Internet timestamp) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- Security—Junos OS detects packets where the IP option is 2 (security) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- Stream ID—Junos OS detects packets where the IP option is 8 (stream ID) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.

If a packet with any of the previous IP options is received, Junos OS flags this as a network reconnaissance attack and records the event for the ingress interface.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Detecting Packets That Use IP Screen Options for Reconnaissance

This example shows how to detect packets that use IP screen options for reconnaissance.

- Requirements on page 933
- Overview on page 934
- Configuration on page 934
- Verification on page 935

Requirements

Before you begin, understand how network reconnaissance works. See "Understanding Network Reconnaissance Using IP Options" on page 930.

Copyright © 2016, Juniper Networks, Inc.

Overview

RFC 791, *Internet Protocol,* specifies a set of options for providing special routing controls, diagnostic tools, and security. The screen options detect IP options that an attacker can use for reconnaissance, including record route, timestamp, security, and stream ID.

In this example, you configure an IP screen screen-1 and enable it in a security zone called zone-1.

NOTE: You can enable only one screen in one security zone.

Configuration

CLI Quick Configuration

To quickly detect packets with the record route, timestamp, security, and stream ID IP screen options, copy the following commands and paste them into the CLI.

```
[edit]
set security screen ids-option screen-1 ip record-route-option
set security screen ids-option screen-1 ip timestamp-option
set security screen ids-option screen-1 ip security-option
set security screen ids-option screen-1 ip stream-option
set security zones security-zone zone-1 screen screen-1
```

Step-by-Step Procedure

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

To detect packets that use IP screen options for reconnaissance:

1. Configure IP screen options.

   NOTE: Currently, these screen options support IPv4 only.

   ```
   [edit security screen]
   user@host# set ids-option screen-1 ip record-route-option
   user@host# set ids-option screen-1 ip timestamp-option
   user@host# set ids-option screen-1 ip security-option
   user@host# set ids-option screen-1 ip stream-option
   ```

2. Enable the screen in the security zone.

   ```
   [edit security zones ]
   user@host# set security-zone zone-1 screen screen-1
   ```

Results

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

```
[edit]
```


Copyright © 2016, Juniper Networks, Inc.

```
[user@host]show security screen
  ids-option screen-1 {
    ip {
      record-route-option;
      timestamp-option;
      security-option;
      stream-option;
    }
  }
[edit]
[user@host]show security zones
  zones {
    security-zone zone-1 {
      screen screen-1;
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 935
- Verifying the Security Screen Configuration on page 935

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** *screen-name* command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                                    Value
  IP record route option                  enabled
  IP timestamp option                     enabled
```

Copyright © 2016, Juniper Networks, Inc.

```
IP security option                          enabled
IP stream option                            enabled
```


Copyright © 2016, Juniper Networks, Inc.

CHAPTER 40

Protecting Against System Probes and Flag Set

- Understanding Operating System Probes on page 937
- Understanding Domain Name System Resolve on page 937
- Understanding TCP Headers with SYN and FIN Flags Set on page 938
- Example: Blocking Packets with SYN and FIN Flags Set on page 939
- Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set on page 940
- Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set on page 941
- Understanding TCP Header with No Flags Set on page 943
- Example: Blocking Packets with No Flags Set on page 944

Understanding Operating System Probes

Before launching an exploit, attackers might try to probe the targeted host to learn its operating system (OS). With that knowledge, they can better decide which attack to launch and which vulnerabilities to exploit. Junos OS can block reconnaissance probes commonly used to gather information about OS types.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Understanding Domain Name System Resolve

Prior to Junos OS Release 12.1X47, DNS resolution was performed with only UDP as a transport. Messages carried by UDP are restricted to 512 bytes; longer messages are truncated and the traffic class (TC) bit is set in the header. The maximum length of UDP DNS response messages is 512 bytes, but the maximum length of TCP DNS response messages is 65,535 bytes. A DNS resolver knows whether the response is complete if the TC bit is set in the header. Hence, a TCP DNS response can carry more information than a UDP DNS response.

There are three types of DNS resolve behaviors:

- UDP DNS resolve
- TCP DNS resolve

Copyright © 2016, Juniper Networks, Inc.

- UDP/TCP DNS resolve

NOTE: A policy uses UDP/TCP DNS resolve to resolve IP addresses. In UDP/TCP DNS resolve, UDP DNS resolve is first used, and when it gets truncated TCP DNS resolve is used.

NOTE: A Routing Engine policy supports a maximum of 1024 IPv4 address prefixes and 256 IPv6 address prefixes that can be sent to the PFE. If the maximum number of IPv4 or IPv6 address prefixes exceeds the limits, the addresses over the limitations will not be sent to the PFE and a syslog message is generated. The maximum number of addresses in a TCP DNS response is 4094 for IPv4 addresses and 2340 for IPv6 addresses, but only 1024 IPv4 addresses and 256 IPv6 addresses are loaded to the PFE.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Understanding TCP Headers with SYN and FIN Flags Set

Both the SYN and FIN control flags are not normally set in the same TCP segment header. The SYN flag synchronizes sequence numbers to initiate a TCP connection. The FIN flag indicates the end of data transmission to finish a TCP connection. Their purposes are mutually exclusive. A TCP header with the SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. See Figure 50.

Figure 50: TCP Header with SYN and FIN Flags Set

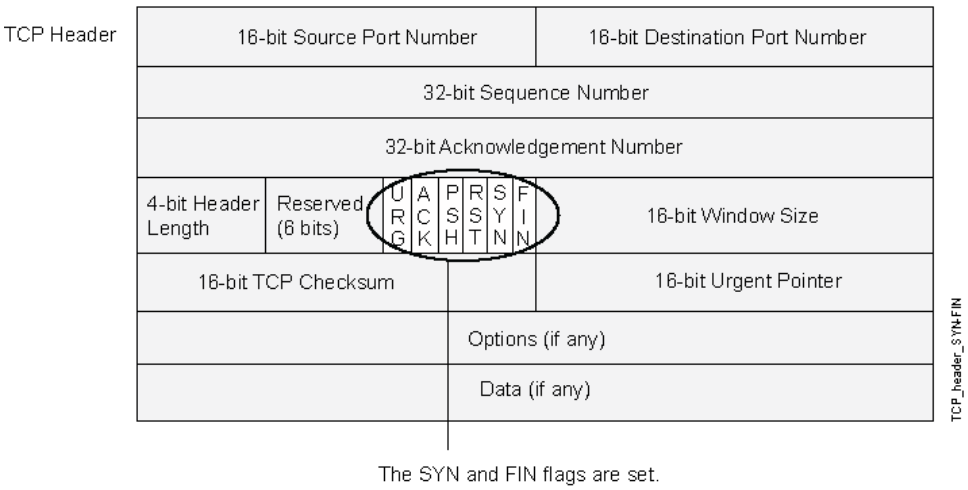

An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks.

Copyright © 2016, Juniper Networks, Inc.

When you enable this screen option, Junos OS checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.

NOTE: Junos OS supports TCP header with SYN and FIN flags set protection for both IPv4 and IPv6 traffic.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking Packets with SYN and FIN Flags Set

This example shows how to create a screen to block packets with the SYN and FIN flags set.

- Requirements on page 939
- Overview on page 939
- Configuration on page 939
- Verification on page 940

Requirements

Before you begin, understand how TCP headers with SYN and FIN flags work. See "Understanding TCP Headers with SYN and FIN Flags Set" on page 938.

Overview

The TCP header with the SYN and FIN flags set cause different responses from a targeted device depending on the OS it is running. The syn-fin screen is enabled for the security zone.

In this example, you create a screen called screen-1 in a security zone to block packets with the SYN and FIN flags set.

Configuration

Step-by-Step Procedure

To block packets with both the SYN and FIN flags set:

1. Configure the screen.

   ```
   [edit]
   user@host# set security screen ids-option screen-1 tcp syn-fin
   ```

2. Enable the screen in the security zone.

   ```
   [edit ]
   user@host# set security zones security-zone zone-1 screen screen-1
   ```

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

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 940
- Verifying the Security Screen Configuration on page 940

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** ***screen-name*** command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                          Value
  TCP SYN FIN                   enabled
```

Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set

Figure 51 shows TCP segments with the FIN control flag set (to signal the conclusion of a session and terminate the connection). Normally, TCP segments with the FIN flag set also have the ACK flag set (to acknowledge the previous packet received). Because a TCP header with the FIN flag set but not the ACK flag is anomalous TCP behavior, there is no uniform response to this. The OS might respond by sending a TCP segment with the RST flag set. Another might completely ignore it. The victim's response can provide the attacker with a clue as to its OS. (Other purposes for sending a TCP segment with the FIN flag set are to evade detection while performing address and port scans and to evade defenses on guard for a SYN flood by performing a FIN flood instead.)

Copyright © 2016, Juniper Networks, Inc.

NOTE: Vendors have interpreted RFC 793, *Transmission Control Protocol*, variously when designing their TCP/IP implementations. When a TCP segment arrives with the FIN flag set but not the ACK flag, some implementations send RST segments, while others drop the packet without sending an RST.

Figure 51: TCP Header with FIN Flag Set

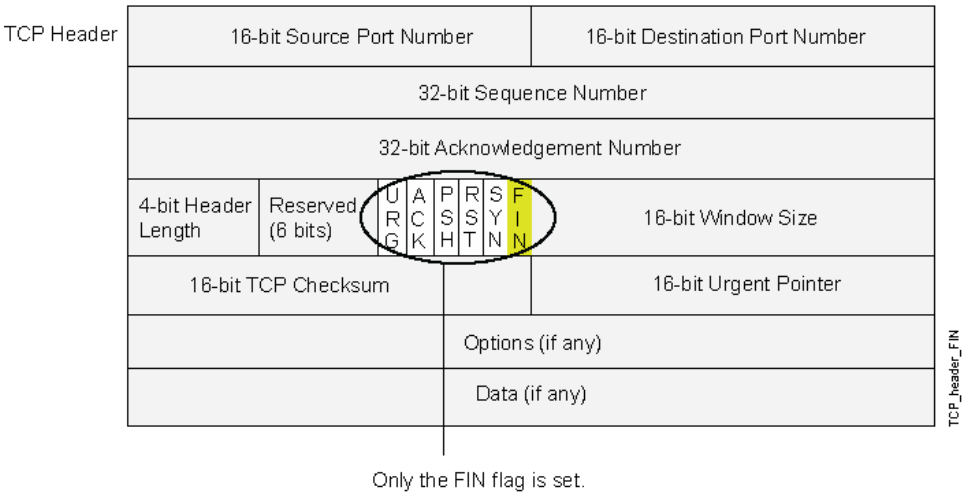

When you enable this screen option, Junos OS checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

NOTE: Junos OS supports TCP header with SYN and FIN flags set protection for both IPv4 and Ipv6 traffic.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set

This example shows how to create a screen to block packets with the FIN flag set but the ACK flag not set.

- Requirements on page 941
- Overview on page 942
- Configuration on page 942
- Verification on page 942

Requirements

Before you begin, understand how TCP headers work. See "Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set" on page 940.

Copyright © 2016, Juniper Networks, Inc.

Overview

The TCP segments with the FIN flag set also have the ACK flag set to acknowledge the previous packet received. Because a TCP header with the FIN flag set but the ACK flag not set is anomalous TCP behavior, there is no uniform response to this. When you enable the fin-no-ack screen option, Junos OS checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

In this example, you create a screen called screen-1 to block packets with the FIN flag set but the ACK flag not set.

Configuration

Step-by-Step Procedure

To block packets with the FIN flag set but the ACK flag not set:

1. Configure the screen.

   ```
   [edit ]
   user@host# set security screen ids-option screen-1 tcp fin-no-ack
   ```

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

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 942
- Verifying the Security Screen Configuration on page 942

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security screen ids-option** *screen-name* command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                          Value
  TCP FIN no ACK                enabled
```

Understanding TCP Header with No Flags Set

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running. See Figure 52.

Figure 52: TCP Header with No Flags Set

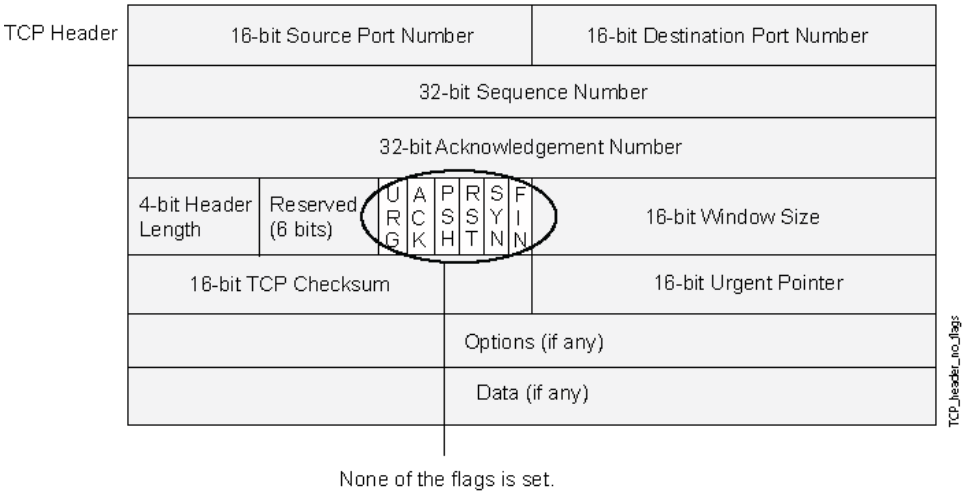

When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

NOTE: Junos OS supports TCP header with no flags set protection for both IPv4 and IPv6 traffic.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking Packets with No Flags Set

This example shows how to create a screen to block packets with no flags set.

- Requirements on page 944
- Overview on page 944
- Configuration on page 944
- Verification on page 944

Requirements

Before you begin, understand how a TCP header with no flags set works. See "Understanding TCP Header with No Flags Set" on page 943.

Overview

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running.

When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

In this example, you create a screen called screen-1 to block packets with no flags set.

Configuration

Step-by-Step Procedure

To block packets with no flags set:

1. Configure the screen.

   ```
   [edit ]
   user@host# set security screen ids-option screen-1 tcp tcp-no-flag
   ```

2. Enable the screen in the security zone.

   ```
   [edit ]
   user@host# set security zones security-zone zone-1 screen screen-1
   ```

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

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 945
- Verifying the Security Screen Configuration on page 945

Copyright © 2016, Juniper Networks, Inc.

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** ***screen-name*** command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                          Value
  TCP no flag                   enabled
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 41

Protecting Against Attacker Evasion Techniques

- Understanding Attacker Evasion Techniques on page 947
- Understanding FIN Scans on page 948
- Thwarting a FIN Scan on page 948
- Understanding TCP SYN Checking on page 948
- Setting TCP SYN Checking on page 950
- Setting Strict SYN Checking on page 951
- Understanding IP Spoofing on page 951
- Example: Blocking IP Spoofing on page 951
- Understanding IP Spoofing in Layer 2 Transparent Mode on page 953
- Configuring IP Spoofing in Layer 2 Transparent Mode on page 954
- Understanding IP Source Route Options on page 955
- Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set on page 957
- Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set on page 958

Understanding Attacker Evasion Techniques

Whether gathering information or launching an attack, it is generally expected that the attacker avoids detection. Although some IP address and port scans are blatant and easily detectable, more wily attackers use a variety of means to conceal their activity. Techniques such as using FIN scans instead of SYN scans—which attackers know most firewalls and intrusion detection programs detect—indicate an evolution of reconnaissance and exploit techniques for evading detection and successfully accomplishing their tasks.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Copyright © 2016, Juniper Networks, Inc.

Understanding FIN Scans

A FIN scan sends TCP segments with the FIN flag set in an attempt to provoke a response (a TCP segment with the RST flag set) and thereby discover an active host or an active port on a host. Attackers might use this approach rather than perform an address sweep with ICMP echo requests or an address scan with SYN segments, because they know that many firewalls typically guard against the latter two approaches but not necessarily against FIN segments. The use of TCP segments with the FIN flag set might evade detection and thereby help the attackers succeed in their reconnaissance efforts.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Thwarting a FIN Scan

To thwart FIN scans, take either or both of the following actions:

- Enable the screen option that specifically blocks TCP segments with the FIN flag set but not the ACK flag, which is anomalous for a TCP segment:

  ```
  user@host#set security screen fin-no-ack tcp fin-no-ack
  user@host#set security zones security-zone name screen fin-no-ack
  ```

  where ***name*** is the name of the zone to which you want to apply this screen option .

- Change the packet processing behavior to reject all non-SYN packets that do not belong to an existing session. The SYN check flag is set as the default.

NOTE: **Changing the packet flow to check that the SYN flag is set for packets that do not belong to existing sessions also thwarts other types of non-SYN scans, such as a null scan (when no TCP flags are set).**

Understanding TCP SYN Checking

By default, Junos OS checks for SYN flags in the first packet of a session and rejects any TCP segments with non-SYN flags attempting to initiate a session. You can leave this packet flow as is or change it so that Junos OS does not enforce SYN flag checking before creating a session. Figure 53 illustrates packet flow sequences both when SYN flag checking is enabled and when it is disabled.

Copyright © 2016, Juniper Networks, Inc.

Figure 53: SYN Flag Checking

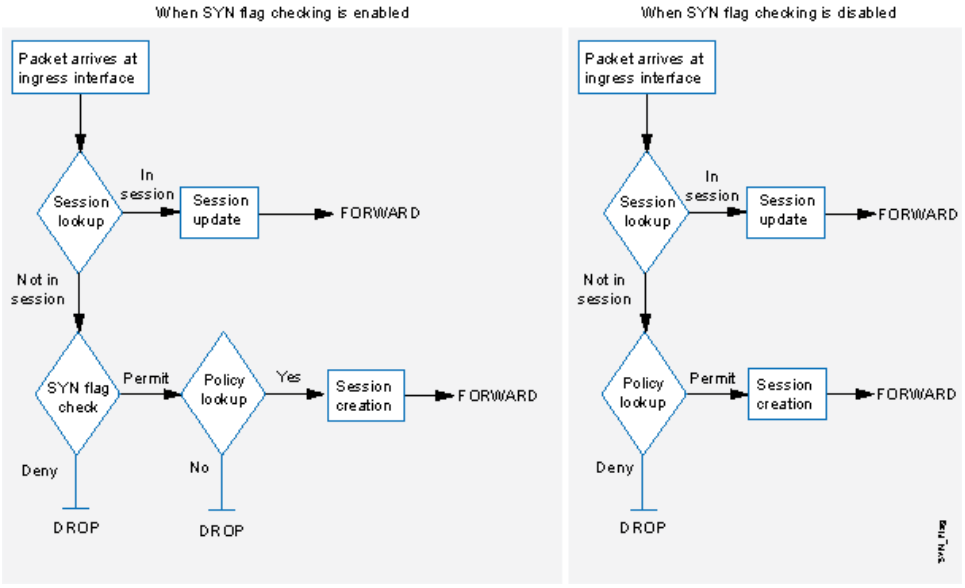

When Junos OS with SYN flag checking enabled receives a non-SYN TCP segment that does not belong to an existing session, it drops the packet. By default, Junos OS does not send a TCP RST to the source host on receiving the non-SYN segment. You can configure the device to send TCP RST to the source host by using the **set security zones security-zone trust tcp-rst** command. If the code bit of the initial non-SYN TCP packet is RST, the device does not send a TCP-RST.

Not checking for the SYN flag in the first packets offers the following advantages:

- NSRP with Asymmetric Routing—In an active/active NSRP configuration in a dynamic routing environment, a host might send the initial TCP segment with the SYN flag set to one device (Device-A), but the SYN/ACK might be routed to the other device in the cluster (Device-B). If this asymmetric routing occurs after Device-A has synchronized its session with Device-B, all is well. On the other hand, if the SYN/ACK response reaches Device-B before Device-A synchronizes the session and SYN checking is enabled, Device-B rejects the SYN/ACK, and the session cannot be established. With SYN checking disabled, Device-B accepts the SYN/ACK response—even though there is no existing session to which it belongs—and creates a new session table entry for it.
- Uninterrupted Sessions—If you reset the device or even change a component in the core section of a policy and SYN checking is enabled, all existing sessions or those sessions to which the policy change applies are disrupted and must be restarted. Disabling SYN checking avoids such disruptions to network traffic flows.

NOTE: A solution to this scenario is to install the device with SYN checking disabled initially. Then, after a few hours—when established sessions are running through the device—enable SYN checking. The core section in a policy contains the following main components: source and destination zones, source and destination addresses, one or more services, and an action.

Copyright © 2016, Juniper Networks, Inc.

However, the previous advantages exact the following security sacrifices:

- Reconnaissance Holes—When an initial TCP segment with a non-SYN flag—such as ACK, URG, RST, FIN—arrives at a closed port, many operating systems (Windows, for example) respond with a TCP segment that has the RST flag set. If the port is open, then the recipient does not generate any response.

  By analyzing these responses or lack thereof, an intelligence gatherer can perform reconnaissance on the protected network and also on the Junos OS policy set. If a TCP segment is sent with a non-SYN flag set and the policy permits it through, the destination host receiving such a segment might drop it and respond with a TCP segment that has the RST flag set. Such a response informs the perpetrator of the presence of an active host at a specific address and that the targeted port number is closed. The intelligence gatherer also learns that the firewall policy permits access to that port number on that host.

  By enabling SYN flag checking, Junos OS drops TCP segments without a SYN flag if they do not belong to an existing session. It does not return a TCP RST segment. Consequently, the scanner gets no replies regardless of the policy set or whether the port is open or closed on the targeted host.

- Session Table Floods—If SYN checking is disabled, an attacker can bypass the Junos OS SYN flood protection feature by flooding a protected network with a barrage of TCP segments that have non-SYN flags set. Although the targeted hosts drop the packets—and possibly send TCP RST segments in reply—such a flood can fill up the session table of the Juniper Networks device. When the session table is full, the device cannot process new sessions for legitimate traffic.

  By enabling SYN checking and SYN flood protection, you can thwart this kind of attack. Checking that the SYN flag is set on the initial packet in a session forces all new sessions to begin with a TCP segment that has the SYN flag set. SYN flood protection then limits the number of TCP SYN segments per second so that the session table does not become overwhelmed.

If you do not need SYN checking disabled, Juniper Networks strongly recommends that it be enabled (its default state for an initial installation of Junos OS). You can enable it with the **set flow tcp-syn-check** command. With SYN checking enabled, the device rejects TCP segments with non-SYN flags set unless they belong to an established session.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Setting TCP SYN Checking

With SYN checking enabled, the device rejects TCP segments with non-SYN flags set unless they belong to an established session. Enabling SYN checking can help prevent attacker reconnaissance and session table floods. TCP SYN checking is enabled by default.

To disable SYN checking:

```
user@host#set security flow tcp-session no-syn-check
```


Copyright © 2016, Juniper Networks, Inc.

Setting Strict SYN Checking

With strict SYN checking enabled, the device enables the strict three-way handshake check for the TCP session. It enhances security by dropping data packets before the three-way handshake is done. TCP strict SYN checking is disabled by default.

NOTE: **The strict-syn-check option cannot be enabled if no-syn-check or no-syn-check-in-tunnel is enabled.**

To enable strict SYN checking:

```
user@host#set security flow tcp-session strict-syn-check
```

Understanding IP Spoofing

One method of attempting to gain access to a restricted area of the network is to insert a false source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing. The mechanism to detect IP spoofing relies on route table entries. For example, if a packet with source IP address 10.1.1.6 arrives at ge-0/0/1, but Junos OS has a route to 10.1.1.0/24 through ge-0/0/0, a check for IP spoofing discovers that this address arrived at an invalid interface as defined in the route table. A valid packet from 10.1.1.6 can only arrive via ge-0/0/0, not ge-0/0/1. Therefore, Junos OS concludes that the packet has a spoofed source IP address and discards it.

NOTE: **Junos OS detects and drops both IPv4 and IPv6 spoofed packets.**

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking IP Spoofing

This example shows how to configure a screen to block IP spoof attacks.

- Requirements on page 951
- Overview on page 952
- Configuration on page 952
- Verification on page 952

Requirements

Before you begin, understand how IP Spoofing works. See "Understanding IP Spoofing" on page 951.

Copyright © 2016, Juniper Networks, Inc.

Overview

One method of attempting to gain access to a restricted area of a network is to insert a bogus source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing.

In this example, you configure a screen called screen-1 to block IP spoof attacks and enable the screen in the zone-1 security zone.

Configuration

Step-by-Step Procedure To block IP spoofing:

1. Configure the screen.

   ```
   [edit ]
   user@host# set security screen ids-option screen-1 ip spoofing
   ```

2. Enable the screen in the security zone.

   ```
   [edit]
   user@host# set security zone security-zone zone-1 screen screen-1
   ```

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

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 952
- Verifying the Security Screen Configuration on page 952

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security screen ids-option** *screen-name* command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                        Value
  IP spoofing                 enabled
```

Understanding IP Spoofing in Layer 2 Transparent Mode

In an IP spoofing attack, the attacker gains access to a restricted area of the network and inserts a false source address in the packet header to make the packet appear to come from a trusted source. IP spoofing is most frequently used in denial-of-service (DoS) attacks. When SRX Series devices are operating in transparent mode, the IP spoof-checking mechanism makes use of address book entries. Address books only exist on the Routing Engine. IP spoofing in Layer 2 transparent mode is performed on the Packet Forwarding Engine. Address book information cannot be obtained from the Routing Engine each time a packet is received by the Packet Forwarding Engine. Therefore, address books attached to the Layer 2 zones must be pushed to the Packet Forwarding Engine.

NOTE: **IP spoofing in Layer 2 transparent mode does not support DNS and wildcard addresses.**

When a packet is received by the Packet Forwarding Engine, the packet's source IP address is checked to determine if it is in the incoming zone's address-book. If the packet's source IP address is in the incoming zone's address book, then this IP address is allowed on the interface, and traffic is passed.

If the source IP address is not present in the incoming zone's address-book, but exists in other zones', then the IP address is considered a spoofed IP. Accordingly, actions such as drop and logging can be taken depending on the screen configuration (alarm-without-drop).

NOTE: **If the alarm-without-drop option is configured, the Layer 2 spoofing packet only triggers an alarm message, but the packet is not dropped.**

If a packet's source IP address is not present in the incoming zone's address book or other zones', then you cannot determine if the IP is spoofed or not. In such instances, the packet is passed.

Junos OS takes into account the following match conditions while it searches for source IP addresses in the address book:

Copyright © 2016, Juniper Networks, Inc.

- **Host-match**—The IP address match found in the address-book is an address without a prefix.
- **Prefix-match**—The IP address match found in the address-book is an address with a prefix.
- **Any-match**—The IP address match found in the address-book is “any”, “any-IPv4”, or “any-IPv6”.
- **No-match**—No IP address match is found.

Related Documentation

- Configuring IP Spoofing in Layer 2 Transparent Mode on page 954

Configuring IP Spoofing in Layer 2 Transparent Mode

You can configure the IP spoof-checking mechanism to determine whether or not an IP is being spoofed.

To configure IP spoofing in Layer 2 transparent mode:

1. Set the interface in Layer 2 transparent mode.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching
   ```

2. (Optional) Set the zone in Layer 2 transparent mode.

   ```
   [edit]
   user@host# set security zones security-zone untrust interfaces ge-0/0/1.0
   ```

3. Configure the address book.

   ```
   [edit]
   user@host# set security address-book my-book address myadd1 10.1.1.0/24
   user@host# set security address-book my-book address myadd2 10.1.2.0/24
   ```

4. Apply the address book to the zone.

   ```
   [edit]
   user@host# set security address-book my-book attach zone untrust
   ```

5. Configure screen IP spoofing.

   ```
   [edit]
   user@host# set security screen ids-option my-screen ip spoofing
   ```

6. Apply the screen to the zone.

   ```
   [edit]
   user@host# set security zones security-zone untrust screen my-screen
   ```

7. (Optional) Configure the **alarm-without-drop** option.

   ```
   [edit]
   user@host# set security screen ids-option my-screen alarm-without-drop
   ```


Copyright © 2016, Juniper Networks, Inc.

NOTE: If the **alarm-without-drop** option is configured, the Layer 2 spoofing packet only triggers an alarm message, but the packet is not dropped.

Related Documentation

- Understanding IP Spoofing in Layer 2 Transparent Mode on page 953

Understanding IP Source Route Options

Source routing was designed to allow users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as “hops”) along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis. If, for example, the transmission of a packet to a particular destination meets with irregular success, you might first use either the record route or the timestamp IP option to discover the addresses of devices along the path or paths that the packet takes. You can then use either the loose or the strict source route option to direct traffic along a specific path, using the addresses you learned from the results that the record route or timestamp options produced. By changing device addresses to alter the path and sending several packets along different paths, you can note changes that either improve or lessen the success rate. Through analysis and the process of elimination, you might be able to deduce where the trouble lies. See Figure 54.

Figure 54: IP Source Routing

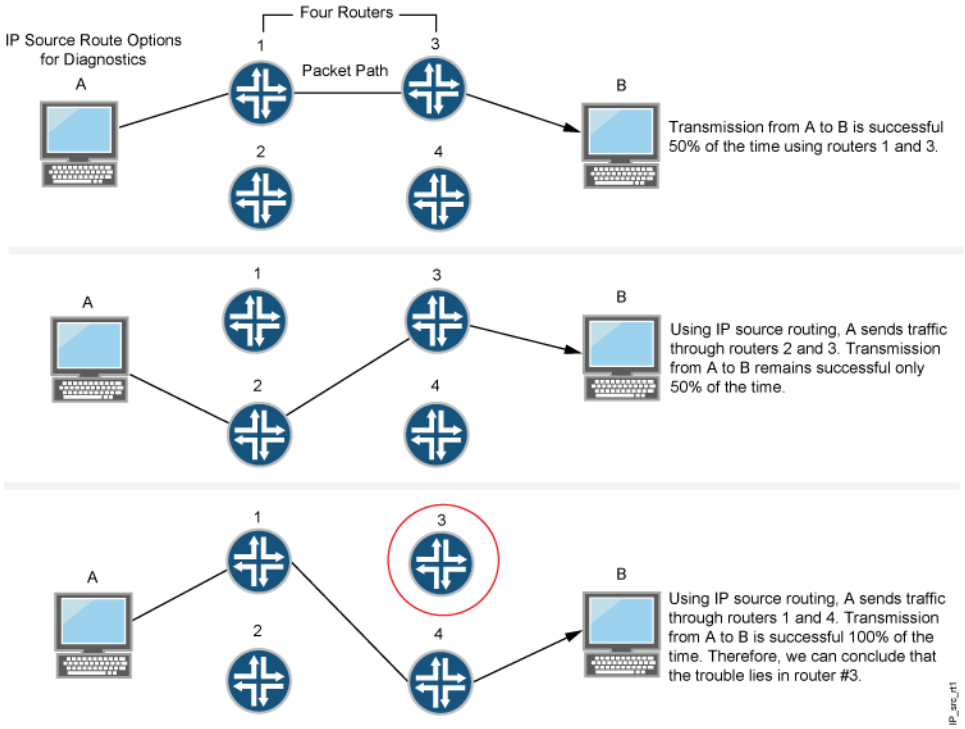

Although the uses of IP source route options were originally benign, attackers have learned to put them to more devious uses. They can use IP source route options to hide their true

Copyright © 2016, Juniper Networks, Inc.

address and access restricted areas of a network by specifying a different path. For an example showing how an attacker can put both deceptions to use, consider the following scenario as illustrated in Figure 55.

Figure 55: Loose IP Source Route Option for Deception

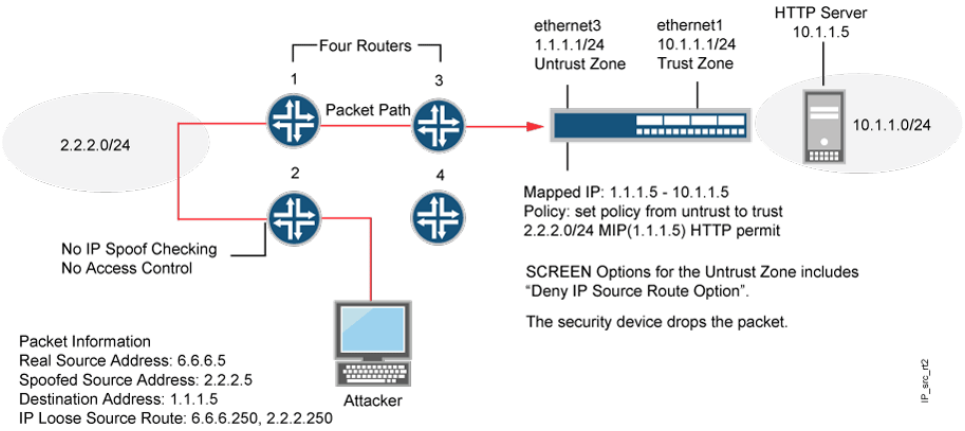

Junos OS only allows traffic 2.2.2.0/24 if it comes through ethernet1, an interface bound to zone_external. Devices 3 and 4 enforce access controls but devices 1 and 2 do not. Furthermore, device 2 does not check for IP spoofing. The attacker spoofs the source address and, by using the loose source route option, directs the packet through device 2 to the 2.2.2.0/24 network and from there out device 1. Device 1 forwards it to device 3, which forwards it to the Juniper Networks device. Because the packet came from the 2.2.2.0/24 subnet and has a source address from that subnet, it seems to be valid. However, one remnant of the earlier chicanery remains: the loose source route option. In this example, you have enabled the deny IP source route screen option for zone_external. When the packet arrives at ethernet3, the device rejects it.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface. The screen options are as follows:

- Deny IP Source Route Option—Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.
- Detect IP Loose Source Route Option—The device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.
- Detect IP Strict Source Route Option—The device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable to IPv4 only.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Reconnaissance Deterrence Overview on page 923

Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set

This example shows how to block packets with either a loose or a strict source route option set.

- Requirements on page 957
- Overview on page 957
- Configuration on page 957
- Verification on page 958

Requirements

Before you begin, understand how IP source route options work. See "Understanding IP Source Route Options" on page 955.

Overview

Source routing allows users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as "hops") along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface.

In this example you create the screen called screen-1 to block packets with either a loose or a strict source route option set and enable the screen in the zone-1 security zone.

Configuration

Step-by-Step Procedure

To block packets with either the loose or the strict source route option set:

1. Configure the screen.

   ```
   [edit ]
   user@host# set security screen ids-option screen-1 ip source-route-option
   ```

2. Enable the screen in the security zone.

   ```
   [edit ]
   user@host# set security zones security-zone zone-1 screen screen-1
   ```

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

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 958
- Verifying the Security Screen Configuration on page 958

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option *screen-name*** command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

Name                                    Value
IP source route option                  enabled
```

Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set

This example shows how to detect packets with either a loose or a strict source route option set.

- Requirements on page 958
- Overview on page 959
- Configuration on page 959
- Verification on page 959

Requirements

Before you begin, understand how IP source route options work. See “Understanding IP Source Route Options” on page 955.

Copyright © 2016, Juniper Networks, Inc.

Overview

Source routing allows users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as “hops”) along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface.

In this example, you create two screens called screen-1 and screen-2 to detect and record, but not block, packets with a loose or strict source route option set and enable the screens in zones zone-1 and zone-2.

Configuration

Step-by-Step Procedure

To detect and record, but not block, packets with a loose or strict source route option set:

1. Configure the loose source screen.

   ```
   [edit]
   user@host# set security screen ids-option screen-1 ip loose-source-route-option
   ```

2. Configure the strict source route screen.

   ```
   [edit]
   user@host# set security screen ids-option screen-2 ip strict-source-route-option
   ```

   NOTE: Currently, this screen option supports IPv4 only.

3. Enable the screens in the security zones.

   ```
   [edit]
   user@host# set security zones security-zone zone-1 screen screen-1
   user@host# set security zones security-zone zone-2 screen screen-2
   ```

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

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

Verification

Confirm that the configuration is working properly.

- Verifying the Screens in the Security Zone on page 959
- Verifying the Security Screen Configuration on page 960

Verifying the Screens in the Security Zone

Purpose Verify that the screen is enabled in the security zone.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security zones** command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
Security zone: zone-2
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-2
  Interfaces bound: 1
  Interfaces:
    ge-2/0/0.0
```

Verifying the Security Screen Configuration

Purpose Display the configuration information about the security screen.

Action From operational mode, enter the **show security screen ids-option** ***screen-name*** command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

Screen object status:

Name                                    Value
IP loose source route option            enabled


[edit]
user@host> show security screen ids-option screen-2
Screen object status:

Screen object status:

Name                                    Value
IP strict source route option           enabled
```


Copyright © 2016, Juniper Networks, Inc.

PART 8

Configuring Suspicious Packet Attributes for Security Devices

- Protecting Against ICMP and SYN Fragment Attacks on page 963
- Protecting Against IP Attacks on page 969

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 42

Protecting Against ICMP and SYN Fragment Attacks

- Suspicious Packet Attributes Overview on page 963
- Understanding ICMP Fragment Protection on page 963
- Example: Blocking Fragmented ICMP Packets on page 964
- Understanding Large ICMP Packet Protection on page 965
- Example: Blocking Large ICMP Packets on page 966
- Understanding SYN Fragment Protection on page 966
- Example: Dropping IP Packets Containing SYN Fragments on page 967

Suspicious Packet Attributes Overview

Attackers can craft packets to perform reconnaissance or launch denial-of-service (DoS) attacks. Sometimes it is unclear what the intent of a crafted packet is, but the very fact that it is crafted suggests that it is being put to some kind of insidious use.

The following topics describe screen options that block suspicious packets that might contain hidden threats:

- Understanding ICMP Fragment Protection on page 963
- Understanding Large ICMP Packet Protection on page 965
- Understanding Bad IP Option Protection on page 969
- Understanding Unknown Protocol Protection on page 971
- Understanding IP Packet Fragment Protection on page 972
- Understanding SYN Fragment Protection on page 966

Understanding ICMP Fragment Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.

Copyright © 2016, Juniper Networks, Inc.

When you enable the ICMP fragment protection screen option, Junos OS blocks any ICMP packet that has the More Fragments flag set or that has an offset value indicated in the offset field. See Figure 56.

Figure 56: Blocking ICMP Fragments

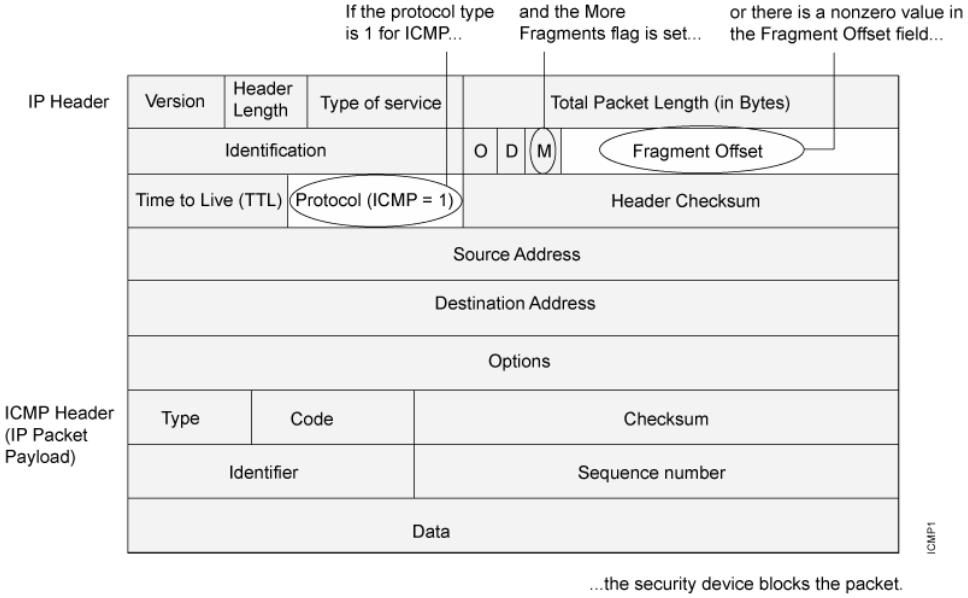

NOTE: Junos OS supports ICMP fragment protection for ICMPv6 packets.

Example: Blocking Fragmented ICMP Packets

This example shows how to block fragmented ICMP packets.

Requirements

Before you begin, Understand ICMP fragment protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When you enable the ICMP fragment protection screen option, Junos OS blocks any ICMP packet that has the more fragments flag set or that has an offset value indicated in the offset field.

In this example, you configure the ICMP fragment screen to block fragmented ICMP packets originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To block fragmented ICMP packets:

1. Configure the screen.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security screen ids-option icmp-fragment icmp fragment
```

2. Configure a security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen icmp-fragment
```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Understanding Large ICMP Packet Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for large ICMP packets. If an ICMP packet is unusually large, something is amiss.

See Figure 57.

Figure 57: Blocking Large ICMP Packets

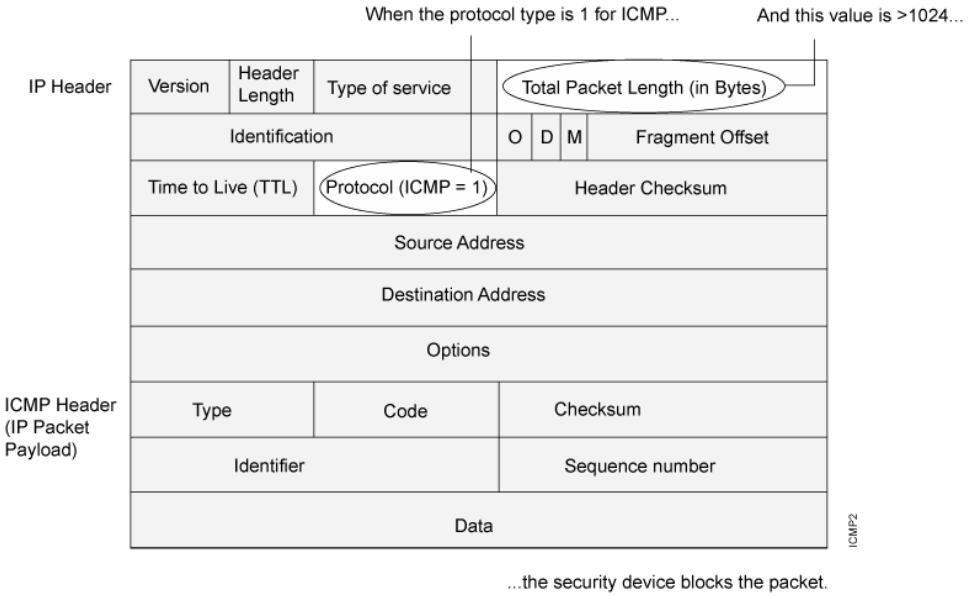

When you enable the large size ICMP packet protection screen option, Junos OS drops ICMP packets with a length greater than 1024 bytes.

NOTE: Junos OS supports large ICMP packet protection for both ICMP and ICMPv6 packets.

Copyright © 2016, Juniper Networks, Inc.

Example: Blocking Large ICMP Packets

This example shows how to block large ICMP packets.

Requirements

Before you begin, Understand large ICMP packet protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When you enable the large ICMP packet protection screen option, Junos OS drops ICMP packets that are larger than 1024 bytes.

In this example, you configure the ICMP large screen to block large ICMP packets originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To block large ICMP packets:

1. Configure the screen.

   ```
   [edit]
   user@host# set security screen ids-option icmp-large icmp large
   ```

2. Configure a security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone1 screen icmp-large
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Understanding SYN Fragment Protection

The IP encapsulates a TCP SYN segment in the IP packet that initiates a TCP connection. Because the purpose of this packet is to initiate a connection and invoke a SYN/ACK segment in response, the SYN segment typically does not contain any data. Because the IP packet is small, there is no legitimate reason for it to be fragmented.

A fragmented SYN packet is anomalous, and, as such, it is suspect. To be cautious, block such unknown elements from entering your protected network. See Figure 58.

Copyright © 2016, Juniper Networks, Inc.

Figure 58: SYN Fragments

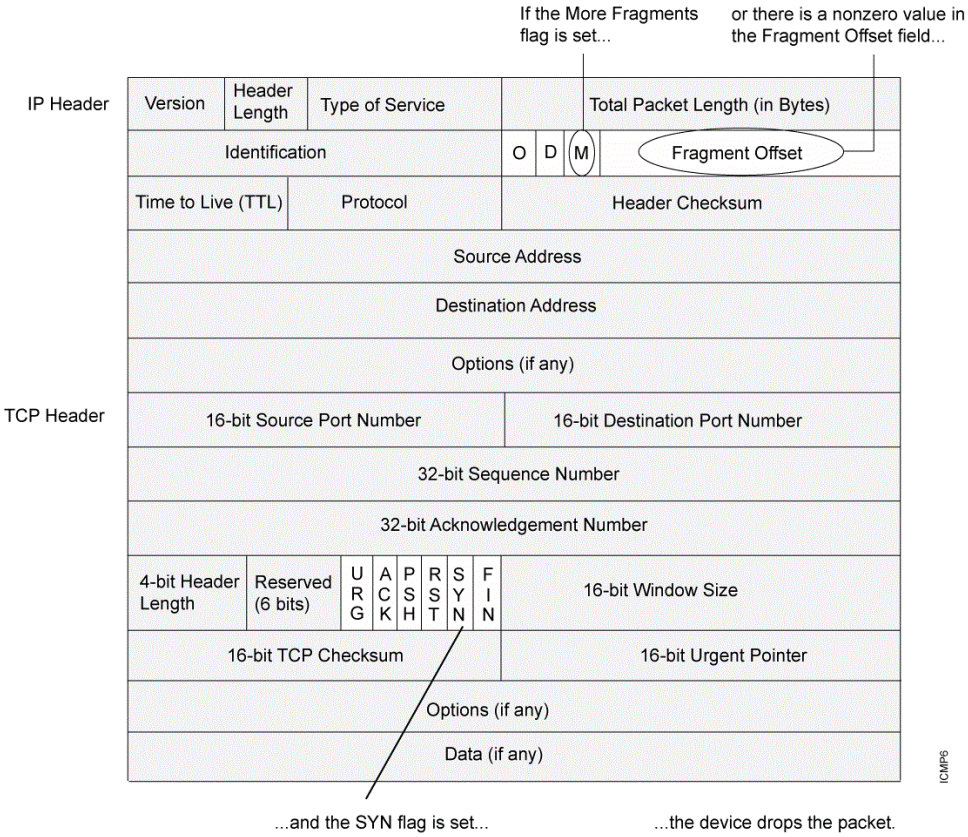

When you enable the SYN fragment detection screen option, Junos OS detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. Junos OS records the event in the screen counters list for the ingress interface.

NOTE: Junos OS supports SYN fragment protection for both IPv4 and IPv6 packets.

Example: Dropping IP Packets Containing SYN Fragments

This example shows how to drop IP packets containing SYN fragments.

Requirements

Before you begin, Understand IP packet fragment protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When you enable the SYN fragment detection screen option, Junos OS detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is

Copyright © 2016, Juniper Networks, Inc.

set in the TCP header. Also, Junos OS records the event in the screen counters list for the ingress interface.

In this example, you configure the SYN fragment screen to drop fragmented SYN packets originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To drop IP packets containing SYN fragments:

1. Configure the screen.

   ```
   [edit]
   user@host# set security screen ids-option syn-frag tcp syn-frag
   ```

2. Configure the security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone1 screen syn-frag
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 43

Protecting Against IP Attacks

- Understanding Bad IP Option Protection on page 969
- Example: Blocking IP Packets with Incorrectly Formatted Options on page 970
- Understanding Unknown Protocol Protection on page 971
- Example: Dropping Packets Using an Unknown Protocol on page 972
- Understanding IP Packet Fragment Protection on page 972
- Example: Dropping Fragmented IP Packets on page 973

Understanding Bad IP Option Protection

The IP standard RFC 791, *Internet Protocol*, specifies a set of eight options that provide special routing controls, diagnostic tools, and security. Although the original, intended uses for these options served worthy ends, people have figured out ways to twist these options to accomplish less commendable objectives.

Either intentionally or accidentally, attackers sometimes configure IP options incorrectly, producing either incomplete or malformed fields. Regardless of the intentions of the person who crafted the packet, the incorrect formatting is anomalous and potentially harmful to the intended recipient. See Figure 59.

Copyright © 2016, Juniper Networks, Inc.

Figure 59: Incorrectly Formatted IP Options

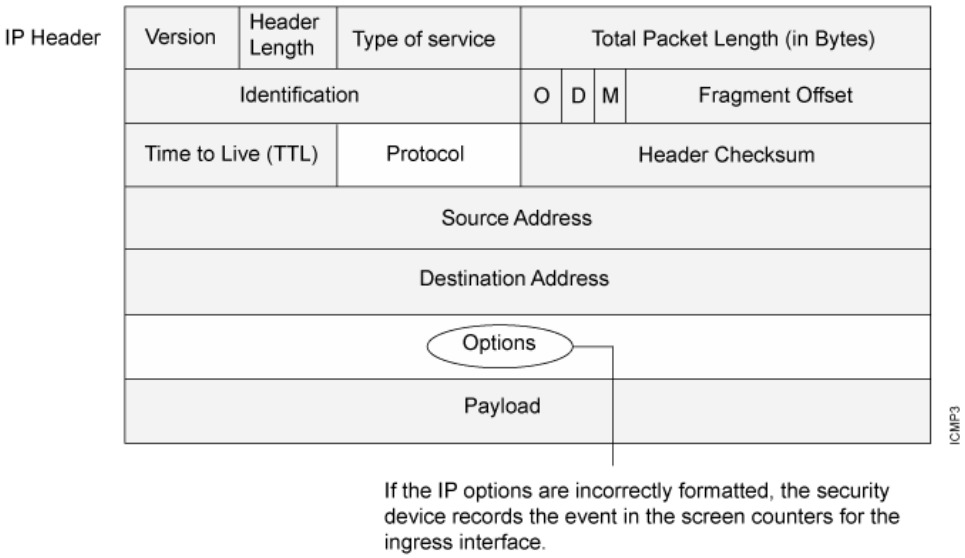

When you enable the bad IP option protection screen option, Junos OS blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, Junos OS records the event in the event log.

NOTE: Junos OS supports bad IP option protection for both IPv4 and IPv6 packets.

Example: Blocking IP Packets with Incorrectly Formatted Options

This example shows how to block large ICMP packets with incorrectly formatted options.

Requirements

Before you begin, Understand bad IP option protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When you enable the bad IP option protection screen option, Junos OS blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, Junos OS records the event in the event log.

In this example, you configure the IP bad option screen to block large ICMP packets originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To detect and block IP packets with incorrectly formatted IP options:

1. Configure the screen.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security screen ids-option ip-bad-option ip bad-option
```

NOTE: Currently this screen option is applicable only to IPv4.

2. Configure a security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen ip-bad-option
```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Understanding Unknown Protocol Protection

Based on RFC 1700, the protocol types with ID numbers of 137 or greater are reserved and undefined at this time. Precisely because these protocols are undefined, there is no way to know in advance if a particular unknown protocol is benign or malicious.

Unless your network makes use of a nonstandard protocol with an ID number of 137 or greater, a cautious stance is to block such unknown elements from entering your protected network. See Figure 60.

Figure 60: Unknown Protocols

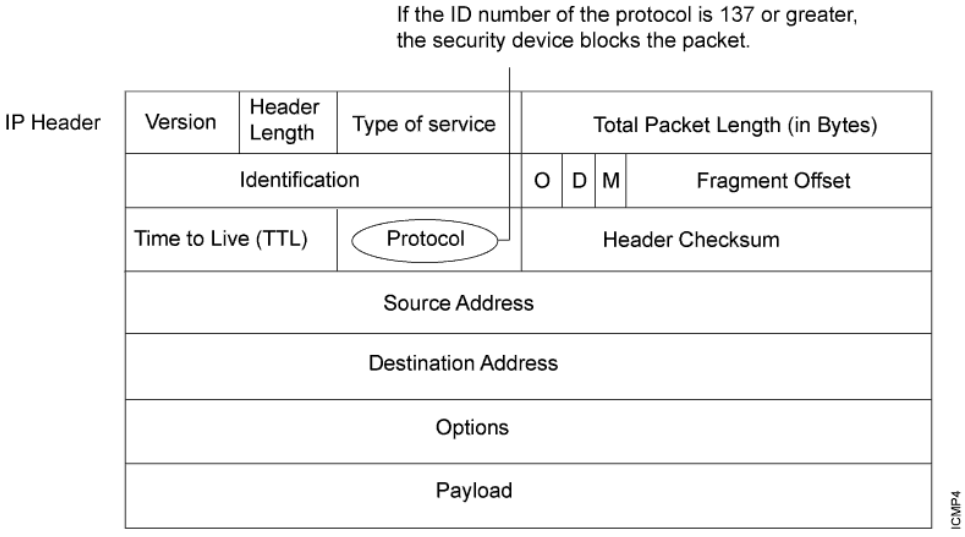

When you enable the unknown protocol protection screen option, Junos OS drops packets when the protocol field contains a protocol ID number of 137 or greater by default.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When you enable the unknown protocol protection screen option for IPv6 protocol , Junos OS drops packets when the protocol field contains a protocol ID number of 139 or greater by default.

Example: Dropping Packets Using an Unknown Protocol

This example shows how to drop packets using an unknown protocol.

Requirements

Before you begin, Understand unknown protocol protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When you enable the unknown protocol protection screen option, Junos OS drops packets when the protocol field contains a protocol ID number of 137 or greater by default.

In this example, you configure the unknown protocol screen to block packets with an unknown protocol originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To drop packets that use an unknown protocol:

1. Configure the unknown protocol screen.

   ```
   [edit]
   user@host# set security screen ids-option unknown-protocol ip unknown-protocol
   ```

2. Configure a security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone1 screen unknown-protocol
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Understanding IP Packet Fragment Protection

As packets traverse different networks, it is sometimes necessary to break a packet into smaller pieces (fragments) based upon the maximum transmission unit (MTU) of each network. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the victim

Copyright © 2016, Juniper Networks, Inc.

receives these packets, the results can range from processing the packets incorrectly to crashing the entire system. See Figure 61.

Figure 61: IP Packet Fragments

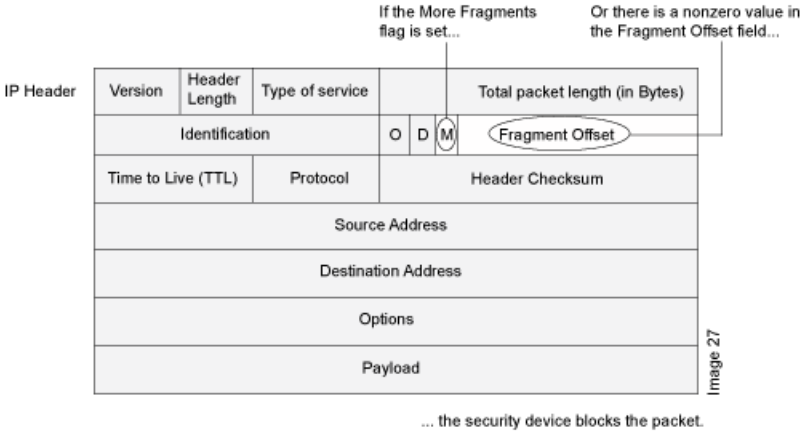

When you enable Junos OS to deny IP fragments on a security zone, it blocks all IP packet fragments that it receives at interfaces bound to that zone.

NOTE: **Junos OS supports IP fragment protection for both IPv4 and IPv6 packets.**

In IPv6 packets, fragment information is not present in the IPv6 header. The fragment information is present in the fragment extension header, which is responsible for IPv6 fragmentation and reassembly. The source node inserts the fragment extension header between the IPv6 header and the payload header if fragmentation is required. See Figure 62.

Figure 62: IPv6 Packet

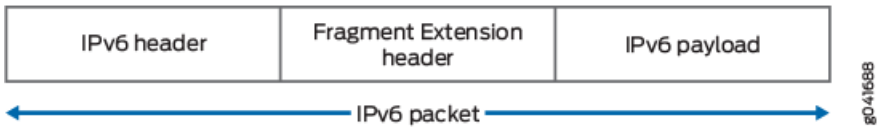

The general format of the fragment extension header is shown in Figure 63.

Figure 63: Fragment Extension Header

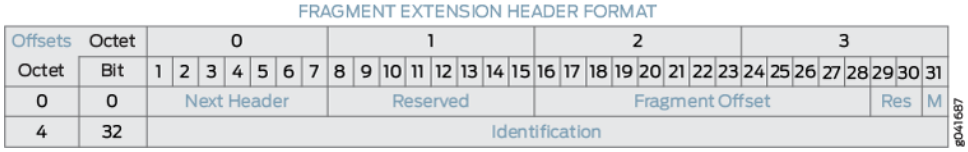

Example: Dropping Fragmented IP Packets

This example shows how to drop fragmented IP packets.

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, Understand IP packet fragment protection. See "Suspicious Packet Attributes Overview" on page 963.

Overview

When this feature is enabled, Junos OS denies IP fragments on a security zone and blocks all IP packet fragments that are received at interfaces bound to that zone.

In this example, you configure the block fragment screen to drop fragmented IP packets originating from the zone1 security zone.

Configuration

Step-by-Step Procedure

To drop fragmented IP packets:

1. Configure the screen.

   ```
   [edit]
   user@host# set security screen ids-option block-frag ip block-frag
   ```

2. Configure the security zone.

   ```
   [edit]
   user@host# set security zones security-zone zone1 screen block-frag
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security screen statistics zone *zone-name*** command.

Copyright © 2016, Juniper Networks, Inc.

PART 9

Configuration Statements and Operational Commands

- Configuration Statements on page 977
- Operational Commands on page 1021

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 44

Configuration Statements

- Security Configuration Statement Hierarchy on page 978
- [edit security screen] Hierarchy Level on page 979
- attack-threshold on page 982
- bad-inner-header on page 983
- description (Security Screen) on page 983
- destination-ip-based on page 984
- destination-threshold on page 985
- fin-no-ack on page 985
- flood (Security ICMP) on page 986
- flood (Security UDP) on page 987
- gre on page 988
- icmp (Security Screen) on page 989
- ids-option on page 990
- ipip on page 993
- ip (Security Screen) on page 994
- ip-sweep on page 997
- ip-in-udp on page 998
- land on page 998
- large on page 999
- limit-session on page 999
- no-syn-check on page 1000
- no-syn-check-in-tunnel on page 1000
- ping-death on page 1001
- port-scan on page 1002
- screen (Security Zones) on page 1003
- source-ip-based on page 1003
- source-threshold on page 1004
- strict-syn-check on page 1004

Copyright © 2016, Juniper Networks, Inc.

- syn-ack-ack-proxy on page 1005
- syn-check-required on page 1005
- syn-fin on page 1006
- syn-flood on page 1007
- syn-flood-protection-mode on page 1008
- syn-frag on page 1008
- tcp (Security Screen) on page 1009
- tcp-no-flag on page 1010
- tcp-sweep on page 1011
- timeout (Security Screen) on page 1012
- traceoptions (Security Screen) on page 1013
- trap on page 1014
- tunnel (Security Screen) on page 1015
- udp (Security Screen) on page 1016
- udp-sweep on page 1017
- white-list on page 1018
- winnuke on page 1019

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*

Copyright © 2016, Juniper Networks, Inc.

- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

[edit security screen] Hierarchy Level

```
security {
  screen {
    ids-option screen-name {
        alarm-without-drop;
        description text;
        icmp {
          flood {
            threshold number;
          }
          fragment;
          icmpv6-malformed;
          ip-sweep {
```

```
        threshold number;
    }
    large;
    ping-death;
}
ip {
    bad-option;
    block-frag;
    ipv6-extension-header {
        AH-header;
        ESP-header;
        HIP-header;
        destination-header {
            ILNP-nonce-option;
            home-address-option;
            line-identification-option;
            tunnel-encapsulation-limit-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        fragment-header;
        hop-by-hop-header {
            CALIPSO-option;
            RPL-option;
            SFM-DPD-option;
            jumbo-payload-option;
            quick-start-option;
            router-alert-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        mobility-header;
        no-next-header;
        routing-header;
        shim6-header
        user-defined-option-type <type-low> to <type-high>;
    }
    ipv6-extension-header-limit limit;
    ipv6-malformed-header;
    loose-source-route-option;
    record-route-option;
    security-option;
    source-route-option;
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
        gre {
            gre-4in4;
            gre-4in6;
            gre-6in4;
            gre-6in6;
        }
        ip-in-udp {
            teredo;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            ipip {
              ipip-4in4;
              ipip-4in6;
              ipip-6in4;
              ipip-6in6;
              ipip-6over4;
              ipip-6to4relay;
              isatap;
              dslite;
            }
            bad-inner-header;
          }
        }
        limit-session {
          destination-ip-based number;
          source-ip-based number;
        }
        tcp {
          fin-no-ack;
          land;
          port-scan {
            threshold number;
          }
          syn-ack-ack-proxy {
            threshold number;
          }
          syn-fin;
          syn-flood {
            alarm-threshold number;
            attack-threshold number;
            destination-threshold number;
            source-threshold number;
            timeout seconds;
            white-list name {
              destination-address destination-address;
              source-address source-address;
            }
          }
          syn-frag;
          tcp-no-flag;
          tcp-sweep {
            threshold threshold number;
          }
          winnuke;
        }
        udp {
          flood {
            threshold number;
          }
          udp-sweep {
            threshold threshold number;
          }
        }
      }
    }
```

```
        traceoptions {
          file filename {
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
        trap {
          interval trap interval;
        }
      }
    }
```

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

attack-threshold

Syntax attack-threshold *number*;

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement modified in Junos OS Release 9.2.

Description Define the number of SYN packets per second required to trigger the SYN proxy response.

Options ***number*** —Number of SYN packets per second required to trigger the SYN proxy response.

Range: 1 through 500,000 per second

Default: 200 per second

NOTE: For SRX Series devices, the applicable range is 1 through 1.000,000 per second.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603
- **destination-threshold on page 985**

Copyright © 2016, Juniper Networks, Inc.

bad-inner-header

Syntax bad-inner-header;

Hierarchy Level [edit security screen ids-option *ids-option-name* ip tunnel]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Enable IP tunnel bad inner header IDS option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Security Configuration Statement Hierarchy on page 603

description (Security Screen)

Syntax description *text*;

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for a screen.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options *text*—Descriptive text about a screen.
Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-ip-based

Syntax destination-ip-based *number*;

Hierarchy Level [edit security screen ids-option *screen-name* limit-session]

Release Information Statement modified in Junos OS Release 9.2.

Description Limit the number of concurrent sessions the device can direct to a single destination IP address.

Options ***number*** —Maximum number of concurrent sessions that can be directed to a destination IP address.

Range: 1 through 1,000,000

Default: 128

NOTE: For SRX Series devices, the applicable range is 1 through 8,000,000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-threshold

Syntax destination-threshold *number* ;

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement modified in Junos OS Release 9.2.

Description Specify the number of SYN segments received per second for a single destination IP address before the device begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based only on the destination IP address, regardless of the destination port number.

Options *number*—Number of SYN segments received per second before the device begins dropping connection requests.

Range: 4 through 1,000,000 per second

Default: 4000 per second

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603
- **attack-threshold on page 982**

fin-no-ack

Syntax fin-no-ack;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable detection of an illegal combination of flags, and reject packets that have this combination.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

flood (Security ICMP)

Syntax

```
flood {
  threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement modified in Junos OS Release 9.2.

Description Configure the device to detect and prevent Internet Control Message Protocol (ICMP) floods. An ICMP flood occurs when ICMP echo requests are broadcast with the purpose of flooding a system with so much data that it first slows down, and then times out and is disconnected. The threshold defines the number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.

Options

threshold ***number*** —Number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.

Range: 1 through 1,000,000 per second

Default: 1,000 per second

NOTE: For SRX Series devices the applicable range is 1 through 4,000,000 per second.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- **flood (Security UDP) on page 987**
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

flood (Security UDP)

Syntax

```
flood {
  threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* udp]

Release Information Statement modified in Junos OS Release 9.2.

Description Configure the device to detect and prevent UDP floods. UDP flooding occurs when an attacker sends UDP packets to slow down the system to the point that it can no longer process valid connection requests.

The threshold defines the number of UDP packets per second allowed to ping the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.

Options

threshold ***number*** —Number of UDP packets per second allowed to ping the same destination address before the device rejects further UDP packets.

Range: 1 through 1,000,000 per second

Default: 1,000 per second

NOTE: For SRX series devices the applicable range is 1 through 4,000,000 per second.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603
- **flood (Security ICMP) on page 986**

Copyright © 2016, Juniper Networks, Inc.

gre

Syntax

```
gre {
   gre-4in4;
   gre-4in6;
   gre-6in4;
   gre-6in6;
}
```

Hierarchy Level [edit security screen ids-option *ids-option-name* ip tunnel]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure IP tunnel GRE IDS option.

Options

gre-4in4— Enable IP tunnel GRE 4in4 IDS option.

gre-4in6— Enable IP tunnel GRE 4in6 IDS option.

gre-6in4— Enable IP tunnel GRE 6in4 IDS option.

gre-6in6— Enable IP tunnel GRE 6in6 IDS option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

icmp (Security Screen)

Syntax

```
icmp {
    flood {
        threshold number;
    }
    fragment;
    icmpv6-malformed;
    ip-sweep {
        threshold number;
    }
    large;
    ping-death;
}
```

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure ICMP intrusion detection service (IDS) options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

ids-option

Syntax

```
ids-option screen-name {
    alarm-without-drop;
    description text;
    icmp {
      flood {
        threshold number;
      }
      fragment;
      icmpv6-malformed;
      ip-sweep {
        threshold number;
      }
      large;
      ping-death;
    }
    ip {
      bad-option;
      block-frag;
      ipv6-extension-header {
        AH-header;
        ESP-header;
        HIP-header;
        }
        destination-header {
          ILNP-nonce-option;
          home-address-option;
          line-identification-option;
          tunnel-encapsulation-limit-option;
          user-defined-option-type <type-low> to <type-high>;
        }
        fragment-header;
        hop-by-hop-header {
          CALIPSO-option;
          RPL-option;
          SFM-DPD-option;
          jumbo-payload-option;
          quick-start-option;
          router-alert-option;
          user-defined-option-type <type-low> to <type-high>;
        }
        mobility-header;
        no-next-header;
        routing-header;
        shim6-header
        user-defined-option-type <type-low> to <type-high>;
      }
      ipv6-extension-header-limit limit;
      ipv6-malformed-header;
      loose-source-route-option;
      record-route-option;
      security-option;
      source-route-option;
```


Copyright © 2016, Juniper Networks, Inc.

```
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
      gre {
        gre-4in4;
        gre-4in6;
        gre-6in4;
        gre-6in6;
      }
      ip-in-udp {
        teredo;
      }
      ipip {
        ipip-4in4;
        ipip-4in6;
        ipip-6in4;
        ipip-6in6;
        ipip-6over4;
        ipip-6to4relay;
        isatap;
        dslite;
      }
      bad-inner-header;
    }
  }
  limit-session {
    destination-ip-based number;
    source-ip-based number;
  }
  tcp {
    fin-no-ack;
    land;
    port-scan {
      threshold number;
    }
    syn-ack-ack-proxy {
      threshold number;
    }
    syn-fin;
    syn-flood {
      alarm-threshold number;
      attack-threshold number;
      destination-threshold number;
      source-threshold number;
      timeout seconds;
      white-list name {
        destination-address destination-address;
        source-address source-address;
      }
    }
    syn-frag;
    tcp-no-flag;
```

```
        tcp-sweep {
          threshold threshold number;
        }
        winnuke;
      }
      udp {
        flood {
          threshold number;
        }
        port-scan {
          threshold number;
        }
        udp-sweep {
          threshold threshold number;
        }
      }
    }
}
```

Hierarchy Level [edit security screen]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **description** option added in Junos OS Release 12.1. UDP supports **port-scan** option starting from Junos OS Release 12.1X47-D10.

Description Define screens for the intrusion detection service (IDS).

Options **description *text***—Descriptive text about a screen.

loose-source-route-option—The device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.

source-route-option—Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.

strict-source-route-option—The device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable to IPv4 only.

NOTE: Loose source route option and strict source route option will only alarm and will not be dropped when there is overflow of traffic. When only IP source option is configured, the attacked packets are dropped.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

ipip

Syntax

```
ipip {
    ipip-4in4;
    ipip-4in6;
    ipip-6in4;
    ipip-6in6;
    ipip-6over4;
    ipip-6to4relay;
    isatap;
    dslite;
}
```

Hierarchy Level [edit security screen ids-option *ids-option-name* ip tunnel]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure IP tunnel IP-IP IDS options.

Options

dslite— Enable IP tunnel IP-IP DS-Lite IDS option.

ipip-4in4— Enable IP tunnel IP-IP 4in4 IDS option.

ipip-4in6— Enable IP tunnel IP-IP 4in6 IDS option.

ipip-6in4— Enable IP tunnel IP-IP 6in4 IDS option.

ipip-6in6— Enable IP tunnel IP-IP 6in6 IDS option.

ipip-6over4— Enable IP tunnel IP-IP 6over4 IDS option.

ipip-6to4relay— Enable IP tunnel IP-IP 6to4 Relay IDS option.

isatap— Enable IP tunnel IP-IP ISATAP IDS option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

ip (Security Screen)

Syntax

```
ip {
    bad-option;
    block-frag;
    ipv6-extension-header {
        AH-header;
        ESP-header;
        HIP-header;
        destination-header {
            ILNP-nonce-option;
            home-address-option;
            line-identification-option;
            tunnel-encapsulation-limit-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        fragment-header;
        hop-by-hop-header {
            CALIPSO-option;
            RPL-option;
            SFM-DPD-option;
            jumbo-payload-option;
            quick-start-option;
            router-alert-option;
            user-defined-option-type <type-low> to <type-high>;
        }
        mobility-header;
        no-next-header;
        routing-header;
        shim6-header
        user-defined-option-type <type-low> to <type-high>;
    }
    ipv6-extension-header-limit limit;
    ipv6-malformed-header;
    loose-source-route-option;
    record-route-option;
    security-option;
    source-route-option;
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
        gre {
            gre-4in4;
            gre-4in6;
            gre-6in4;
            gre-6in6;
        }
        ip-in-udp {
            teredo;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
    ipip {
      ipip-4in4;
      ipip-4in6;
      ipip-6in4;
      ipip-6in6;
      ipip-6over4;
      ipip-6to4relay;
      isatap;
      dslite;
    }
    bad-inner-header;
  }
}
```

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 bad-option extension header screens added in Junos OS Release 12.1X46-D10.

Description Configure IP layer IDS options.

Options

- **bad-option**—Detect and drop any packet with an incorrectly formatted IP option in the IP packet header. The device records the event in the screen counters list for the ingress interface. This screen option is applicable to IPv4 and IPv6.
- **block-frag**—Enable IP packet fragmentation blocking.
- **loose-source-route-option**—Detect packets where the IP option is 3 (loose source routing), and record the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified. The type 0 routing header of the loose source route option is the only related header defined in IPv6 .
- **record-route-option**—Detect packets where the IP option is 7 (record route), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **security-option**—Detect packets where the IP option is 2 (security), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **source-route-option**—Detect packets, and record the event in the screen counters list for the ingress interface.
- **spoofing**—Prevent spoofing attacks. Spoofing attacks occur when unauthorized agents attempt to bypass firewall security by imitating valid client IP addresses. Using the spoofing option invalidates such false source IP address connections.

  The default behavior is to base spoofing decisions on individual interfaces.
- **stream-option**—Detect packets where the IP option is 8 (stream ID), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **strict-source-route-option**—Detect packets where the IP option is 9 (strict source routing), and record the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable only to IPv4.
- **tear-drop**—Block the teardrop attack. Teardrop attacks occur when fragmented IP packets overlap and cause the host attempting to reassemble the packets to crash. The teardrop option directs the device to drop any packets that have such a discrepancy.
- **timestamp-option**—Detect packets where the IP option list includes option 4 (Internet timestamp), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **unknown-protocol**—Discard all received IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6. Such protocol numbers are undefined or reserved.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

ip-sweep

Syntax

```
ip-sweep {
    threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the device to detect and prevent an IP Sweep attack. An IP Sweep attack occurs when an attacker sends ICMP echo requests (pings) to multiple destination addresses. If a target host replies, the reply reveals the target's IP address to the attacker. If the device receives 10 ICMP echo requests within the number of microseconds specified in this statement, it flags this as an IP Sweep attack, and rejects the 11th and all further ICMP packets from that host for the remainder of the second.

Options

threshold *number*—Maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device. More than 10 requests from a host during this period triggers an IP Sweep attack response on the device during the remainder of the second.

Range: 1000 through 1,000,000 microseconds

Default: 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

ip-in-udp

Syntax

```
ip-in-udp {
    teredo;
}
```

Hierarchy Level [edit security screen ids-option *ids-option-name* ip tunnel]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure IP tunnel IPinUDP IDS option.

Options **teredo**— Enable IP tunnel IPinUDP Teredo IDS option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Security Configuration Statement Hierarchy on page 603

land

Syntax land;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable prevention of Land attacks by combining the SYN flood defense with IP spoofing protection. Land attacks occur when an attacker sends spoofed IP packets with headers containing the target's IP address for the source and destination IP addresses. The attacker sends these packets with the SYN flag set to any available port. The packets induce the target to create empty sessions with itself, filling its session table and overwhelming its resources.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

large

Syntax large;

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the device to detect and drop any ICMP frame with an IP length greater than 1024 bytes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

limit-session

Syntax
```
limit-session {
  destination-ip-based number;
  source-ip-based number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Limit the number of concurrent sessions the device can initiate from a single source IP address or the number of sessions it can direct to a single destination IP address.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

no-syn-check

Syntax no-syn-check;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Disable checking of the TCP SYN bit before creating a session. By default, the device checks that the SYN bit is set in the first packet of a session. If the bit is not set, the device drops the packet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

no-syn-check-in-tunnel

Syntax no-syn-check-in-tunnel;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Disable checking of the TCP SYN bit before creating a session for tunneled packets. By default, the device checks that the SYN bit is set in the first packet of a VPN session. If the bit is not set, the device drops the packet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Copyright © 2016, Juniper Networks, Inc.

ping-death

Syntax ping-death;

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the device to detect and reject oversized and irregular ICMP packets. Although the TCP/IP specification requires a specific packet size, many ping implementations allow larger packet sizes. Larger packets can trigger a range of adverse system reactions, including crashing, freezing, and restarting.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

port-scan

Syntax

```
port-scan {
  threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Prevent port scan attacks. A port scan attack occurs when an attacker sends packets with different port numbers to scan available services. The attack succeeds if a port responds. To prevent this attack, the device internally logs the number of different ports scanned from a single remote source. For example, if a remote host scans 10 ports in 0.005 seconds (equivalent to 5000 microseconds, the default threshold setting), the device flags this behavior as a port scan attack, and rejects further packets from the remote source.

Options

threshold ***number*** —Number of microseconds during which the device accepts packets from the same remote source with up to 10 different port numbers. If the number of ports during the threshold period reaches 10 or more, the device rejects additional packets from the source.

Range: 1000 through 1,000,000 microseconds

Default: 5000 microseconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

screen (Security Zones)

Syntax screen *screen-name*;

Hierarchy Level [edit security zones functional-zone management],
[edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a security screen for a security zone.

Options ***screen-name*** —Name of the screen.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

source-ip-based

Syntax source-ip-based *number*;

Hierarchy Level [edit security screen ids-option *screen-name* limit-session]

Release Information Statement modified in Junos OS Release 9.2.

Description Limit the number of concurrent sessions the device can initiate from a single source IP address.

Options ***number*** —Maximum number of concurrent sessions that can be initiated from a source IP address.

Range: 1 through 1,000,000

Default: 128

NOTE: For SRX Series devices the applicable range is 1 through 8,000,000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-threshold

Syntax source-threshold *number*;

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement modified in Junos OS Release 9.2.

Description Specify the number of SYN segments that the device can receive per second from a single source IP address (regardless of the destination IP address and port number) before the device begins dropping connection requests from that source.

Options ***number*** —Number of SYN segments to be received per second before the device starts dropping connection requests.

Range: 4 through 500,000 per second

Default: 4000 per second

NOTE: **For SRX Series devices the applicable range is 4 through 1,000,000 per second.**

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

strict-syn-check

Syntax strict-syn-check;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 9.4.

Description Enable the strict three-way handshake check for the TCP session. It enhances security by dropping data packets before the three-way handshake is done. By default, strict-syn-check is disabled.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Copyright © 2016, Juniper Networks, Inc.

syn-ack-ack-proxy

Syntax

```
syn-ack-ack-proxy; {
    threshold number,
}
```

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5; support for IPv6 addresses added in Junos OS Release 10.4.

Description Prevent the SYN-ACK-ACK attack, which occurs when the attacker establishes multiple telnet sessions without allowing each session to terminate. This behavior consumes all open slots, generating a denial-of-service (DoS) condition.

Options threshold ***number*** — Number of connections from any single IP address.

Range: 1 through 250,000

Default: 512

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

syn-check-required

Syntax syn-check-required;

Hierarchy Level [edit security policies from-zone ***zone-name*** to-zone ***zone-name*** policy ***policy-name*** then permit tcp-options]

Release Information Statement introduced in Junos OS Release 10.4.

Description Enable sync check per policy. The syn-check-required value overrides the global value no-syn-check.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Copyright © 2016, Juniper Networks, Inc.

syn-fin

Syntax syn-fin;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable detection of an illegal combination of flags that attackers can use to consume sessions on the target device, thus resulting in a denial-of-service (DoS) condition.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

syn-flood

Syntax

```
syn-flood {
  alarm-threshold number;
  attack-threshold number;
  destination-threshold number;
  source-threshold number;
  timeout seconds;
  white-list name {
    destination-address destination-address;
    source-address source-address;
  }
}
```

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure detection and prevention of SYN flood attacks. Such attacks occur when the connecting host continuously sends TCP SYN requests without replying to the corresponding ACK responses.

NOTE: On all SRX Series devices, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.

The synchronization cookie or proxy never drops packets; therefore the **alarm-without-drop** (not **drop**) action is shown in the system log.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

syn-flood-protection-mode

Syntax syn-flood-protection-mode (syn-cookie | syn-proxy);

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5; support for IPv6 addresses added in Junos OS Release 10.4.

Description Enable SYN cookie or SYN proxy defenses against SYN attacks. SYN flood protection mode is enabled globally on the device and is activated when the configured **syn-flood attack-threshold value** is exceeded.

Options
- **syn-cookie**—Uses a cryptographic hash to generate a unique Initial Sequence Number (ISN). This is enabled by default.
- **syn-proxy**—Uses a proxy to handle the SYN attack.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

syn-frag

Syntax syn-frag;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable detection of a SYN fragment attack and drops any packet fragments used for the attack. A SYN fragment attack floods the target host with SYN packet fragments. The host caches these fragments, waiting for the remaining fragments to arrive so it can reassemble them. The flood of connections that cannot be completed eventually fills the host's memory buffer. No further connections are possible, and damage to the host's operating system can occur.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

tcp (Security Screen)

Syntax

```
tcp {
  fin-no-ack;
  land;
  port-scan {
    threshold number;
  }
  syn-ack-ack-proxy {
    threshold number;
  }
  syn-fin;
  syn-flood {
    alarm-threshold number;
    attack-threshold number;
    destination-threshold number;
    source-threshold number;
    timeout seconds;
    white-list name {
      destination-address destination-address;
      source-address source-address;
    }
  }
  syn-frag;
  tcp-no-flag;
  tcp-sweep {
    threshold threshold number;
  }
  winnuke;
}
```

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure TCP-layer intrusion detection service (IDS) options.

NOTE: On all SRX Series devices, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.

The synchronization cookie or proxy never drops packets; therefore the **alarm-without-drop** (not **drop**) action is shown in the system log.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

tcp-no-flag

Syntax tcp-no-flag;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable the device to drop illegal TCP packets with a missing or malformed flag field.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

tcp-sweep

Syntax

```
tcp-sweep {
  threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the device to detect and prevent TCP sweep attack. In a TCP sweep attack, an attacker sends TCP SYN packets to the target device as part of the TCP handshake. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends TCP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as a TCP sweep attack.

If the **alarm-without-drop** option is not set, the device rejects the eleventh and all further TCP packets from that host for the remainder of the specified threshold period.

Options

threshold *number*—Maximum number of microseconds during which up to 10 TCP SYN packets from the same host are allowed into the device. More than 10 requests from a host during this period triggers TCP Sweep attack response on the router during the remainder of the second.

Range: 1000 through 1,000,000 microseconds

Default: 5000 microseconds

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

timeout (Security Screen)

Syntax timeout *seconds*;

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum length of time before a half-completed connection is dropped from the queue. You can decrease the timeout value until you see any connections dropped during normal traffic conditions.

Options ***seconds*** —Time interval before a half-completed connection is dropped from the queue.

Range: 1 through 50 seconds

Default: 20 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Screen)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security screen]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure screen tracing options.

To specify more than one tracing option, include multiple **flag** statements.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x*** **K** to specify KB, ***x*** **m** to specify MB, or ***x*** **g** to specify GB

Copyright © 2016, Juniper Networks, Inc.

Range: 10 KB through 1 GB

Default: 128 KB

- **world-readable** | **no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace all screen events
  - **configuration**—Trace screen configuration events
  - **flow**—Trace flow events
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

trap

Syntax
```
trap {
  interval trap interval;
}
```

Hierarchy Level
[edit security screen trap]

Release Information
Statement introduced in Junos OS Release 12.3X48-D20.

Description
Configure trap interval.

Options
interval—The trap interval is 1 through 3600 seconds, and the default interval is 2 seconds.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

tunnel (Security Screen)

Syntax

```
tunnel {
  gre {
    gre-4in4;
    gre-4in6;
    gre-6in4;
    gre-6in6;
  }
  ip-in-udp {
    teredo;
  }
  ipip {
    ipip-4in4;
    ipip-4in6;
    ipip-6in4;
    ipip-6in6;
    ipip-6over4;
    ipip-6to4relay;
    isatap;
    dslite;
  }
  bad-inner-header;
}
```

Hierarchy Level [edit security screen ids-option *ids-option-name* ip]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Enable IP tunnel IDS options.

Options The remaining options are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

udp (Security Screen)

Syntax

```
udp {
  flood {
    threshold number;
  }
  udp-sweep {
    threshold threshold number;
  }
}
```

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the number of packets allowed per second to the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

udp-sweep

Syntax

```
udp-sweep {
  threshold number;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* udp]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the device to detect and prevent UDP sweep attack. In a UDP sweep attack, an attacker sends UDP packets to the target device. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends UDP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as an UDP sweep attack.

If the **alarm-without-drop** option is not set, the device rejects the eleventh and all further UDP packets from that host for the remainder of the specified threshold period.

Options

threshold *number*—Maximum number of microseconds during which up to 10 UDP packets from the same host are allowed into the device. More than 10 requests from a host during this period triggers an UDP Sweep attack response on the device during the remainder of the second.

Range: 1000 through 1,000,000 microseconds

Default: 5000 microseconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

white-list

Syntax

```
white-list name {
  destination-address [address];
  source-address [address];
}
```

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement introduced in Junos OS Release 12.1.

Description Configure a whitelist of IP addresses that are to be exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process.

Both IP version 4 (IPv4) and IP version 6 (IPv6) whitelists are supported. Addresses in a whitelist must be all IPv4 or all IPv6. Each whitelist can have up to 32 IP address prefixes.

Options

- **destination-address *address***—Destination IP address or an address prefix. You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.
- **source-address *address***—Source IP address or an address prefix. You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

winnuke

Syntax winnuke;

Hierarchy Level [edit security screen ids-option *screen-name* tcp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable detection of attacks on Windows NetBios communications. Packets are modified as necessary and passed on. Each WinNuke attack triggers an attack log entry in the event alarm log.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 45

Operational Commands

- clear security screen statistics
- clear security screen statistics interface
- clear security screen statistics zone
- show security screen ids-option
- show security screen statistics
- show security screen status

Copyright © 2016, Juniper Networks, Inc.

clear security screen statistics

Syntax clear security screen statistics
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 9.0.

Description Clear intrusion detection service (IDS) security screen statistics on the device.

Options **node**—(Optional) For chassis cluster configurations, clear security screen statistics on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all** —Clear all nodes.
- **local** —Clear the local node.
- **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation
- show security screen statistics on page 1032
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output clear security screen statistics node 0 on page 1022

Output Fields This command produces no output.

Sample Output

clear security screen statistics node 0

```
user@host> clear security screen statistics node 0
```


Copyright © 2016, Juniper Networks, Inc.

clear security screen statistics interface

Syntax clear security screen statistics interface *interface-name*

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Clear intrusion detection service (IDS) security screen statistics for an interface.

Options

- **interface** ***interface-name*** —Name of the interface on which to clear security screen statistics.
- **node**—(Optional) For chassis cluster configurations, clear security screen statistics on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- show security screen statistics on page 1032
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

clear security screen statistics interface fab0 on page 1023
clear security screen statistics interface fab0 node 0 on page 1023

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security screen statistics interface fab0

```
user@host> clear security screen statistics interface fab0
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
node1:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Sample Output

clear security screen statistics interface fab0 node 0

```
user@host> clear security screen statistics interface fab0 node 0
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

clear security screen statistics zone

Syntax clear security screen statistics zone *zone-name*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Clear IDS security screen statistics for a security zone.

Options

- **zone** *zone-name* —Name of the security zone for which to clear security screen statistics.
- **node**—(Optional) For chassis cluster configurations, clear security screen statistics for a security zone on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- show security screen statistics on page 1032
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output
clear security screen statistics zone abc node all on page 1025
clear security screen statistics node 0 zone my-zone on page 1025

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security screen statistics zone abc node all

```
user@host> clear security screen statistics zone abc node all
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
node1:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Sample Output

clear security screen statistics node 0 zone my-zone

```
user@host> clear security screen statistics node 0 zone my-zone
node0:
--------------------------------------------------------------------------
  IDS statistics has been cleared.
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show security screen ids-option

Syntax

show security screen ids-option
screen-name
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5. Support for **node** options added in Junos OS Release 9.0. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10. Support for UDP **port scan** added in Junos OS Release 12.1X47-D10.

Description

Display configuration information about the specified security screen.

Options

- ***screen-name*** —Name of the screen.
- **node**—(Optional) For chassis cluster configurations, display the configuration status of the security screen on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- ids-option on page 990
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

show security screen ids-option jscreen on page 1029
show security screen ids-option jscreen (IPv6) on page 1030
show security screen ids-option jscreen1 node all on page 1030

Output Fields

Table 68 lists the output fields for the **show security screen ids-option** command. Output fields are listed in the approximate order in which they appear.

Table 68: show security screen ids-option Output Fields

Field Name	Field Description
TCP address sweep threshold	Number of microseconds for which the device accepts 10 TCP packets from the same remote source to different destination addresses.
TCP port scan threshold	Number of microseconds during which the device accepts packets from the same remote source with up to 10 different port numbers.
ICMP address sweep threshold	Maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device.

Copyright © 2016, Juniper Networks, Inc.

Table 68: show security screen ids-option Output Fields *(continued)*

Field Name	Field Description
UDP flood threshold	Number of UDP packets per second allowed to ping the same destination address before the device rejects further UDP packets.
UDP port scan threshold	Number of microseconds during which the device accepts packets from the same remote source IP with up to 10 different destination port numbers.
TCP winnuke	Enable or disable the detection of TCP WinNuke attacks.
TCP SYN flood attack threshold	Number of SYN packets per second required to trigger the SYN proxy response.
TCP SYN flood alarm threshold	Number of half-complete proxy connections per second at which the device makes entries in the event alarm log.
TCP SYN flood source threshold	Number of SYN segments to be received per second before the device begins dropping connection requests.
TCP SYN flood destination threshold	Number of SYN segments received per second before the device begins dropping connection requests.
TCP SYN flood timeout	Maximum length of time before a half-completed connection is dropped from the queue.
TCP SYN flood queue size	Number of proxy connection requests that can be held in the proxy connection queue before the device begins rejecting new connection requests.
ICMP large packet	Enable or disable the detection of any ICMP frame with an IP length greater than 1024 bytes.
UDP address sweep threshold	Number of microseconds for which the device accepts 10 UDP packets from the same remote source to different destination addresses.
IPv6 extension routing	Enable or disable the IPv6 extension routing screen option.
IPv6 extension shim6	Enable or disable the IPv6 extension shim6 screen option.
IPv6 extension fragment	Enable or disable the IPv6 extension fragment screen option.
IPv6 extension AH	Enable or disable the IPv6 extension Authentication Header Protocol screen option.
IPv6 extension ESP	Enable or disable the IPv6 extension Encapsulating Security Payload screen option.
IPv6 extension mobility	Enable or disable the IPv6 extension mobility screen option.
IPv6 extension HIP	Enable or disable the IPv6 extension Host Identify Protocol screen option.
IPv6 extension no next	Enable or disable the IPv6 extension no-next screen option.
IPv6 extension user-defined	Enable or disable the IPv6 extension user-defined screen option.

Copyright © 2016, Juniper Networks, Inc.

Table 68: show security screen ids-option Output Fields *(continued)*

Field Name	Field Description
IPv6 extension HbyH jumbo	Enable or disable the IPv6 extension HbyH jumbo screen option.
IPv6 extension HbyH RPL	Enable or disable the IPv6 extension HbyH RPL screen option.
IPv6 extension HbyH router alert	Enable or disable the IPv6 extension HbyH router screen option.
IPv6 extension HbyH quick start	Enable or disable the IPv6 extension HbyH quick-start screen option.
IPv6 extension HbyH CALIPSO	Enable or disable the IPv6 extension HbyH Common Architecture Label IPv6 Security Screen option.
IPv6 extension HbyH SMF DPD	Enable or disable the IPv6 extension HbyH Simplified Multicast Forwarding IPv6 Duplicate Packet Detection screen option.
IPv6 extension HbyH user-defined	Enable or disable the IPv6 extension HbyH user-defined screen option.
IPv6 extension Dst tunnel encap limit	Enable or disable the IPv6 extension distributed (network) storage tunnel encapsulation limit screen option.
IPv6 extension Dst home address	Enable or disable the IPv6 extension DST home address screen option.
IPv6 extension Dst ILNP nonce	Enable or disable the IPv6 extension DST Identifier-Locator Network Protocol nonce screen option.
IPv6 extension Dst line-id	Enable or disable the IPv6 extension DST line-ID screen option.
IPv6 extension Dst user-defined	Enable or disable the IPv6 extension DST user-defined screen option.
IPv6 extension header limit	Threshold for the number of IPv6 extension headers that can pass through the screen.
IPv6 malformed header	Enable or disable the IPv6 malformed header screen option.
ICMPv6 malformed header	Enable or disable the ICMPv6 malformed packet screen option.

Sample Output

show security screen ids-option jscreen

```
user@host> show security screen ids-option jscreen
Screen object status:
Name                                        Value
  TCP port scan threshold                   5000
  UDP port scan threshold                   10000
  ICMP address sweep threshold              5000
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security screen ids-option jscreen (IPv6)

```
user@host> show security screen ids-option jscreen

Screen object status:

   Name                                        Value
  ICMP ping of death                          enabled
  ……

  IPv6 extension routing                      enabled
  IPv6 extension shim6                        enabled
  IPv6 extension fragment                     enabled
  IPv6 extension AH                           enabled
  IPv6 extension ESP                          enabled
  IPv6 extension mobility                     enabled
  IPv6 extension HIP                          enabled
  IPv6 extension no next                      enabled
  IPv6 extension user-defined                 enabled
  IPv6 extension HbyH jumbo                   enabled
  IPv6 extension HbyH RPL                     enabled
  IPv6 extension HbyH router alert            enabled
  IPv6 extension HbyH quick start             enabled
  IPv6 extension HbyH CALIPSO                 enabled
  IPv6 extension HbyH SMF DPD                 enabled
  IPv6 extension HbyH user-defined            enabled
  IPv6 extension Dst tunnel encap limit       enabled
  IPv6 extension Dst home address             enabled
  IPv6 extension Dst ILNP nonce               enabled
  IPv6 extension Dst line-id                  enabled
  IPv6 extension Dst user-defined             enabled
  IPv6 extension header limit                 20
  IPv6 Malformed header                       enabled
  ICMPv6 malformed packet                     enabled
```

Sample Output

show security screen ids-option jscreen1 node all

```
user@host> show security screen ids-option jscreen1 node all
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  UDP flood threshold                       1000
  TCP winnuke                               enabled
  TCP SYN flood attack threshold            200
  TCP SYN flood alarm threshold             512
  TCP SYN flood source threshold            4000
  TCP SYN flood destination threshold       4000
  TCP SYN flood timeout                     20
  TCP SYN flood queue size                  1024
  ICMP large packet                         enabled
node1:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  UDP flood threshold                       1000
```


Copyright © 2016, Juniper Networks, Inc.

```
TCP winnuke                                enabled
TCP SYN flood attack threshold             200
TCP SYN flood alarm threshold              512
TCP SYN flood source threshold             4000
TCP SYN flood destination threshold        4000
TCP SYN flood timeout                      20
TCP SYN flood queue size                   1024
ICMP large packet                          enabled
```

Copyright © 2016, Juniper Networks, Inc.

show security screen statistics

Syntax

show security screen statistics (zone *zone-name* | interface *interface-name*)
<logical-system (*logical-system-name* | all)>
<node (*node-id* | all | local | primary)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 8.5. **node** options added in Junos OS Release 9.0. **logical-system all** option added in Junos OS Release 11.2R6. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

Description

Display intrusion detection service (IDS) security screen statistics.

Options

- **zone** ***zone-name*** —Display screen statistics for this security zone.
- **interface** *interface-name*—Display screen statistics for this interface.
- **logical-system**—(Optional) Display screen statistics for configured logical systems.
  - ***logical-system-name***—Display screen statistics for the named logical system.
  - **all**—Display screen statistics for all logical systems, including the master (root) logical system.
- **node**—(Optional) For chassis cluster configurations, display screen statistics on a specific node.
  - ***node-id***—Identification number of a node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.
- **root-logical-system**—(Optional) Display screen statistics for the master logical system only.

Required Privilege Level

view

Related Documentation

- clear security screen statistics on page 1022
- clear security screen statistics interface on page 1023
- clear security screen statistics zone on page 1025
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

show security screen statistics zone scrzone on page 1035
show security screen statistics zone untrust (IPv6) on page 1035
show security screen statistics interface ge-0/0/3 on page 1036
show security screen statistics interface ge-0/0/1 (IPv6) on page 1036
show security screen statistics interface ge-0/0/1 node primary on page 1037
show security screen statistics zone trust logical-system all on page 1037

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 69 lists the output fields for the **show security screen statistics** command. Output fields are listed in the approximate order in which they appear.

Table 69: show security screen statistics Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.

Copyright © 2016, Juniper Networks, Inc.

Table 69: show security screen statistics Output Fields *(continued)*

Field Name	Field Description
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.
UDP address sweep	Number of UDP address sweeps.
IPv6 extension header	Number of packets filtered for the defined IPv6 extension headers.
IPv6 extension hop by hop option	Number of packets filtered for the defined IPv6 hop-by-hop option types.
IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).

Copyright © 2016, Juniper Networks, Inc.

Table 69: show security screen statistics Output Fields *(continued)*

ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                            Statistics
  ICMP flood                                    0
  UDP  flood                                    0
  TCP  winnuke                                  0
  ……
  IPv6   extension header                        0
  IPv6   extension hop by hop option             0
  IPv6   extension destination option            0
  IPv6   extension header limit                  0
  IPv6   malformed header                        0
```

Copyright © 2016, Juniper Networks, Inc.

```
  ICMPv6 malformed packet                    0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                                 Statistics
  ICMP flood                                        0
  UDP flood                                         0
  ……
  IPv6 extension header                             0
  IPv6 extension hop by hop option                  0
  IPv6 extension destination option                 0
  IPv6 extension header limit                       0
  IPv6 malformed header                             0
  ICMPv6 malformed packet                           0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 1
  UDP flood                                  1
  TCP winnuke                                1
  TCP port scan                              1
  ICMP address sweep                         1
  TCP sweep                                  1
  UDP sweep                                  1
  IP tear drop                               1
  TCP SYN flood                              1
  IP spoofing                                1
  ICMP ping of death                         1
  IP source route option                     1
  TCP land attack                            1
  TCP SYN fragment                           1
  TCP no flag                                1
  IP unknown protocol                        1
  IP bad options                             1
  IP record route option                     1
  IP timestamp option                        1
  IP security option                         1
  IP loose source route option               1
  IP strict source route option              1
  IP stream option                           1
  ICMP fragment                              1
  ICMP large packet                          1
  TCP SYN FIN                                1
  TCP FIN no ACK                             1
  Source session limit                       1
  TCP SYN-ACK-ACK proxy                      1
  IP block fragment                          1
  Destination session limit                  1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
```

Copyright © 2016, Juniper Networks, Inc.

```
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls1
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls2
Screen statistics:

IDS attack type                             Statistics
```


Copyright © 2016, Juniper Networks, Inc.

```
ICMP flood                          0
UDP flood                           0
TCP winnuke                         0
TCP port scan                       0
ICMP address sweep                  0
TCP sweep                           0
UDP sweep                           0
IP tear drop                        0
TCP SYN flood                       0
IP spoofing                         0
ICMP ping of death                  0
IP source route option              0
TCP land attack                     0
TCP SYN fragment                    0
TCP no flag                         0
IP unknown protocol                 0
IP bad options                      0
IP record route option              0
IP timestamp option                 0
IP security option                  0
IP loose source route option        0
IP strict source route option       0
IP stream option                    0
ICMP fragment                       0
ICMP large packet                   0
TCP SYN FIN                         0
TCP FIN no ACK                      0
Source session limit                0
TCP SYN-ACK-ACK proxy               0
IP block fragment                   0
Destination session limit           0
```

Copyright © 2016, Juniper Networks, Inc.

show security screen status

Syntax **show security screen status**

Release Information Command introduced in Junos OS Release 12.3X48-D20.

Description Show screen status data.

Required Privilege Level view

List of Sample Output show security screen status on page 1040

Sample Output

show security screen status

```
user@host> show security screen status
Screen status:
      Screen trap interval : 2 second(s)
```


Copyright © 2016, Juniper Networks, Inc.

Building Blocks Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 10

Overview

- Introduction to Security Building Blocks on page 1045

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 46

Introduction to Security Building Blocks

- Understanding Security Building Blocks for Security Devices on page 1045

Understanding Security Building Blocks for Security Devices

This guide provides information about the building blocks used to configure features for security devices.

- A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic through policies. Security zones are logical entities to which one or more interfaces are bound. With many types of Juniper Networks devices, you can define multiple security zones, the exact number of which you determine based on your network needs.
- An address book is a collection of addresses and address sets. Junos OS allows you to configure multiple address books. Address books are like components, or building blocks, that are referenced in other configurations such as security policies or NAT. You can add addresses to address books or use the predefined addresses available to each address book by default.
- A security policy is a stateful firewall policy that provides a set of tools to network administrators, enabling them to implement network security for their organizations. Security policies enforce rules for transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on traffic as it passes through the firewall.
- An application set is a group of applications. Junos OS simplifies the process by allowing you to manage a small number of application sets, rather than a large number of individual application entries. The application (or application set) is referred to by security policies as match criteria for packets initiating sessions.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding Address Books on page 1069
- Security Policies Overview on page 1085
- Security Policy Applications Overview on page 1171

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 11

Configuring Security Zones and Interfaces

- Configuring Security Zones on page 1049
- Managing Inbound Traffic for Security Zones on page 1055
- Identifying Duplicate Sessions by Configuring TCP-Reset Parameters on page 1065

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 47

Configuring Security Zones

- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050
- Understanding Security Zones on page 1050
- Example: Creating Security Zones on page 1051

Security Zones and Interfaces Overview

Interfaces act as a doorway through which traffic enters and exits a Juniper Networks device. Many interfaces can share exactly the same security requirements; however, different interfaces can also have different security requirements for inbound and outbound data packets. Interfaces with identical security requirements can be grouped together into a single security zone.

A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic through policies.

Security zones are logical entities to which one or more interfaces are bound. With many types of Juniper Networks devices, you can define multiple security zones, the exact number of which you determine based on your network needs.

On a single device, you can configure multiple security zones, dividing the network into segments to which you can apply various security options to satisfy the needs of each segment. At a minimum, you must define two security zones, basically to protect one area of the network from the other. On some security platforms, you can define many security zones, bringing finer granularity to your network security design—and without deploying multiple security appliances to do so.

From the perspective of security policies, traffic enters into one security zone and goes out on another security zone. This combination of a **from-zone** and a **to-zone** is defined as a *context*. Each context contains an ordered list of policies. For more information on policies, see "Security Policies Overview" on page 1085.

This topic includes the following sections:

- Understanding Security Zone Interfaces on page 1050

Copyright © 2016, Juniper Networks, Inc.

Understanding Security Zone Interfaces

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone.

Through the policies you define, you can permit traffic between zones to flow in one direction or in both. With the routes that you define, you specify the interfaces that traffic from one zone to another must use. Because you can bind multiple interfaces to a zone, the routes you chart are important for directing traffic to the interfaces of your choice.

An interface can be configured with an IPv4 address, IPv6 address, or both.

Related Documentation

- Understanding Functional Zones on page 1050
- Example: Creating Security Zones on page 1051
- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055

Understanding Functional Zones

A functional zone is used for special purposes, like management interfaces. Currently, only the management (MGT) zone is supported. Management zones have the following properties:

- Management zones host management interfaces.
- Traffic entering management zones does not match policies; therefore, traffic cannot transit out of any other interface if it was received in the management interface.
- Management zones can only be used for dedicated management interfaces.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Example: Creating Security Zones on page 1051

Understanding Security Zones

Security zones are the building blocks for policies; they are logical entities to which one or more interfaces are bound. Security zones provide a means of distinguishing groups of hosts (user systems and other hosts, such as servers) and their resources from one another in order to apply different security measures to them.

Security zones have the following properties:

- Policies—Active security policies that enforce rules for the transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on the traffic as it passes through the firewall. For more information, see "Security Policies Overview" on page 1085.
- Screens—A Juniper Networks stateful firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security

Copyright © 2016, Juniper Networks, Inc.

zone to another. For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful. For more information, see "Reconnaissance Deterrence Overview" on page 923.

- Address books—IP addresses and address sets that make up an address book to identify its members so that you can apply policies to them. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, and Domain Name System (DNS) names. For more information, see "Example: Configuring Address Books and Address Sets" on page 1076.
- TCP-RST—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYNchronize flag set.
- Interfaces—List of interfaces in the zone.

Security zones have the following preconfigured zone:

- Trust zone—Available only in the factory configuration and is used for initial connection to the device. After you commit a configuration, the trust zone can be overridden.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050
- Example: Creating Security Zones on page 1051

Example: Creating Security Zones

This example shows how to configure zones and assign interfaces to them. When you configure a security zone, you can specify many of its parameters at the same time.

- Requirements on page 1051
- Overview on page 1051
- Configuration on page 1052
- Verification on page 1053

Requirements

Before you begin, configure network interfaces. See the *Interfaces Feature Guide for Security Devices*.

Overview

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone.

NOTE: By default, interfaces are in the null zone. The interfaces will not pass traffic until they have been assigned to a zone.

Copyright © 2016, Juniper Networks, Inc.

NOTE: You can configure 2000 interfaces within a security zone on SRX5400, SRX5600, and SRX5800 devices.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/1 unit 0 family inet address 10.12.12.1/24
set interfaces ge-0/0/1 unit 0 family inet6 address 2001:db8:1::1/64
set security security-zone ABC interfaces ge-0/0/1.0
```

Step-by-Step Procedure

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

To create zones and assign interfaces to them:

1. Configure an Ethernet interface and assign an IPv4 address to it.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.12.12.1/24
   ```

2. Configure an Ethernet interface and assign an IPv6 address to it.

   ```
   user@host# set interfaces ge-0/0/1 unit 0 family inet6 address 2001:db8:1::1/64
   ```

3. Configure a security zone and assign it to an Ethernet interface.

   ```
   user@host# set security security-zone ABC interfaces ge-0/0/1.0
   ```

Results

From configuration mode, confirm your configuration by entering the **show security zones security-zone ABC** and **show interfaces ge-0/0/1** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
[edit]
user@host# show security zones security-zone ABC
...
    interfaces {
        ge-0/0/1.0 {
            ...
        }
    }

[edit]
user@host# show interfaces ge-0/0/1
...
        unit 0 {
        family inet {
            address 10.12.12.1/24;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        family inet6 {
            address 2001:db8:1::1/64;
        }
    }
```

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

Verification

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050
- Understanding Security Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 48

Managing Inbound Traffic for Security Zones

- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Example: Controlling Inbound Traffic Based on Traffic Types on page 1056
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058
- Example: Controlling Inbound Traffic Based on Protocols on page 1059
- Supported System Services for Host Inbound Traffic on page 1061

Understanding How to Control Inbound Traffic Based on Traffic Types

This topic describes how to configure zones to specify the kinds of traffic that can reach the device from systems that are directly connected to its interfaces.

Note the following:

- You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)
- You must enable all expected host-inbound traffic. Inbound traffic destined to this device is dropped by default.
- You can also configure a zone's interfaces to allow for use by dynamic routing protocols.

This feature allows you to protect the device against attacks launched from systems that are directly or indirectly connected to any of its interfaces. It also enables you to selectively configure the device so that administrators can manage it using certain applications on certain interfaces. You can prohibit use of other applications on the same or different interfaces of a zone. For example, most likely you would want to ensure that outsiders not use the Telnet application from the Internet to log into the device because you would not want them connecting to your system.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter on page 1065

Copyright © 2016, Juniper Networks, Inc.

- Example: Controlling Inbound Traffic Based on Traffic Types on page 1056

Example: Controlling Inbound Traffic Based on Traffic Types

This example shows how to configure inbound traffic based on traffic types.

- Requirements on page 1056
- Overview on page 1056
- Configuration on page 1056
- Verification on page 1058

Requirements

Before you begin:

- Configure network interfaces. See *Interfaces Feature Guide for Security Devices*.
- Understand Inbound traffic types. See "Understanding How to Control Inbound Traffic Based on Traffic Types" on page 1055.

Overview

By allowing system services to run, you can configure zones to specify different types of traffic that can reach the device from systems that are directly connected to its interfaces. You can configure the different system services at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)

You must enable all expected host-inbound traffic. Inbound traffic from devices directly connected to the device's interfaces is dropped by default.

Configuration

CLI Quick Configuration

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

```
set security zones security-zone ABC host-inbound-traffic system-services all
set security zones security-zone ABC interfaces ge-0/0/1.3 host-inbound-traffic
    system-services telnet
set security zones security-zone ABC interfaces ge-0/0/1.3 host-inbound-traffic
    system-services ftp
set security zones security-zone ABC interfaces ge-0/0/1.3 host-inbound-traffic
    system-services snmp
set security zones security-zone ABC interfaces ge-0/0/1.0 host-inbound-traffic
    system-services all
set security zones security-zone ABC interfaces ge-0/0/1.0 host-inbound-traffic
    system-services ftp except
set security zones security-zone ABC interfaces ge-0/0/1.0 host-inbound-traffic
    system-services http except
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure inbound traffic based on traffic types:

1. Configure a security zone.

   ```
   [edit]
   user@host# edit security zones security-zone ABC
   ```

2. Configure the security zone to support inbound traffic for all system services.

   ```
   [edit security zones security-zone ABC]
   user@host# set host-inbound-traffic system-services all
   ```

3. Configure the Telnet, FTP, and SNMP system services at the interface level (not the zone level) for the first interface.

   ```
   [edit security zones security-zone ABC]
   user@host# set interfaces ge-0/0/1.3 host-inbound-traffic system-services telnet
   user@host# set interfaces ge-0/0/1.3 host-inbound-traffic system-services ftp
   user@host# set interfaces ge-0/0/1.3 host-inbound-traffic system-services snmp
   ```

4. Configure the security zone to support inbound traffic for all system services for a second interface.

   ```
   [edit security zones security-zone ABC]
   user@host# set interfaces ge-0/0/1.0 host-inbound-traffic system-services all
   ```

5. Exclude the FTP and HTTP system services from the second interface.

   ```
   [edit security zones security-zone ABC]
   user@host# set interfaces ge-0/0/1.0 host-inbound-traffic system-services ftp
     except
   user@host# set interfaces ge-0/0/1.0 host-inbound-traffic system-services http
     except
   ```

Results

From configuration mode, confirm your configuration by entering the **show security zones security-zone ABC**. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host#  show security zones security-zone ABC
host-inbound-traffic {
    system-services {
        all;
    }
}
    interfaces {
    ge-0/0/1.3 {
        host-inbound-traffic {
            system-services {
                ftp;
                telnet;
                snmp;
            }
        }
    }
    ge-0/0/1.0 {
```

Copyright © 2016, Juniper Networks, Inc.

```
            host-inbound-traffic {
                system-services {
                    all;
                    ftp {
                        except;
                    }
                    http {
                        except;
                    }
                }
            }
        }
    }
```

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

Verification

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation
- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055

Understanding How to Control Inbound Traffic Based on Protocols

This topic describes the inbound system protocols on the specified zone or interface.

Any host-inbound traffic that corresponds to a protocol listed under the host-inbound traffic option is allowed. For example, if anywhere in the configuration, you map a protocol to a port number other than the default, you can specify the protocol in the host-inbound traffic option, and the new port number will be used. Table 70 lists the supported protocols. A value of **all** indicates that traffic from all of the following protocols is allowed inbound on the specified interfaces (of the zone, or a single specified interface).

Table 70: Supported Inbound System Protocols

Supported System Services			
all	igmp	pim	sap
bfd	ldp	rip	vrrp
bgp	msdp	ripng	nhrp
router-discovery	dvmrp	ospf	rsvp
pgm	ospf3		

Copyright © 2016, Juniper Networks, Inc.

NOTE: If DVMRP or PIM is enabled for an interface, IGMP and MLD host-inbound traffic is enabled automatically. Because IS-IS uses OSI addressing and should not generate any IP traffic, there is no host-inbound traffic option for the IS-IS protocol.

NOTE: You do not need to configure Neighbor Discovery Protocol (NDP) on host-inbound traffic, because the NDP is enabled by default.

Configuration option for IPv6 Neighbor Discovery Protocol (NDP) is available. The configuration option is **set protocol neighbor-discovery onlink-subnet-only** command. This option will prevent the device from responding to a Neighbor Solicitation (NS) from a prefix which was not included as one of the device interface prefixes.

NOTE: The Routing Engine needs to be rebooted after setting this option to remove any possibility of a previous IPv6 entry from remaining in the forwarding-table.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter on page 1065
- Example: Controlling Inbound Traffic Based on Protocols on page 1059

Example: Controlling Inbound Traffic Based on Protocols

This example shows how to enable inbound traffic for an interface.

- Requirements on page 1059
- Overview on page 1060
- Configuration on page 1060
- Verification on page 1061

Requirements

Before you begin:

- Configure security zones. See "Example: Creating Security Zones" on page 1051.
- Configure network interfaces. See the*Interfaces Feature Guide for Security Devices*.

Overview

Any host-inbound traffic that corresponds to a protocol listed under the host-inbound traffic option is allowed. For example, if anywhere in the configuration you map a protocol to a port number other than the default, you can specify the protocol in the host-inbound traffic option, and the new port number will be used.

A value of **all** indicates that traffic from all of the protocols is allowed inbound on the specified interfaces (of the zone, or a single specified interface).

Configuration

CLI Quick Configuration

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

```
set security zones security-zone ABC interfaces ge-0/0/1.0 host-inbound-traffic protocols
  ospf
set security zones security-zone ABC interfaces ge-0/0/1.0 host-inbound-traffic protocols
  ospf3
```

Step-by-Step Procedure

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

To configure inbound traffic based on protocols:

1. Configure a security zone.

   ```
   [edit]
   user@host# edit security zones security-zone ABC
   ```

2. Configure the security zone to support inbound traffic based on the ospf protocol for an interface.

   ```
   [edit security zones security-zone ABC]
   user@host# set interfaces ge-0/0/1.0 host-inbound-traffic protocols ospf
   ```

3. Configure the security zone to support inbound traffic based on the ospf3 protocol for an interface.

   ```
   [edit security zones security-zone ABC]
   user@host# set interfaces ge-0/0/1.0 host-inbound-traffic protocols ospf3
   ```

Results

From configuration mode, confirm your configuration by entering the **show security zones security-zone ABC**. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security zones security-zone ABC
interfaces {
    ge-0/0/1.0 {
        host-inbound-traffic {
            protocols {
                ospf;
```


Copyright © 2016, Juniper Networks, Inc.

```
                ospf3;
            }
        }
    }
}
```

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

Verification

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058

Supported System Services for Host Inbound Traffic

This topic describes the supported system services for host inbound traffic on the specified zone or interface.

For example, suppose a user whose system was connected to interface **203.0.113.4** in zone **ABC** wanted to telnet into interface **198.51.100.4** in zone **ABC**. For this action to be allowed, the Telnet application must be configured as an allowed inbound service on both interfaces and a policy must permit the traffic transmission.

Table 71 shows the system services that can be used for host inbound traffic.

Table 71: System Services for Host Inbound Traffic

Host Inbound System Services	
all	any-service
dns	finger
ftp	http
https	indent-reset
ike	netconf
ntp	ping
reverse-ssh	reverse-telnet
rlogin	rpm

Copyright © 2016, Juniper Networks, Inc.

Table 71: System Services for Host Inbound Traffic *(continued)*

Host Inbound System Services	
rsh	sip
snmp	snmp-trap
ssh	telnet
tftp	traceroute
xnm-clear-text	xnm-ssl

NOTE: On SRX Series Services Gateways, the **xnm-clear-text** field is enabled in the factory-default configuration. This setting enables incoming Junos XML protocol traffic in the trust zone for the device when the device is operating with factory-default settings. We recommend that you replace the factory-default settings with a user-defined configuration that provides additional security once the box is configured. You must delete the **xnm-clear-text** field manually by using the CLI command **delete system services xnm-clear-text**.

Table 72 shows the supported protocols that can be used for host inbound traffic.

Table 72: Protocols for Host Inbound Traffic

Protocols	
all	bfd
bgp	dvmrp
igmp	msdp
ospf	nhrp
pgm	ospf3
rip	pim
sap	ripng
	vrrp

Copyright © 2016, Juniper Networks, Inc.

NOTE: All services (except DHCP and BOOTP) can be configured either per zone or per interface. A DHCP server is configured only per interface because the incoming interface must be known by the server to be able to send out DHCP replies.

NOTE: You do not need to configure Neighbor Discovery Protocol (NDP) on host-inbound traffic, because the NDP is enabled by default.

Configuration option for IPv6 Neighbor Discovery Protocol (NDP) is available. The configuration option is **set protocol neighbor-discovery onlink-subnet-only** command. This option will prevent the device from responding to a Neighbor Solicitation (NS) from a prefix which was not included as one of the device interface prefixes.

NOTE: The Routing Engine needs to be rebooted after setting this option to remove any possibility of a previous IPv6 entry from remaining in the forwarding-table.

Related Documentation

- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Example: Controlling Inbound Traffic Based on Traffic Types on page 1056

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 49

Identifying Duplicate Sessions by Configuring TCP-Reset Parameters

- Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter on page 1065
- Example: Configuring the TCP-Reset Parameter on page 1065

Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter

When the TCP-RST feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYNchronize flag set.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058
- Example: Configuring the TCP-Reset Parameter on page 1065

Example: Configuring the TCP-Reset Parameter

This example shows how to configure the TCP-Reset parameter for a zone.

- Requirements on page 1065
- Overview on page 1066
- Configuration on page 1066
- Verification on page 1066

Requirements

Before you begin, configure security zones. See “Example: Creating Security Zones” on page 1051.

Copyright © 2016, Juniper Networks, Inc.

Overview

When the TCP-Reset parameter feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the SYN flag set.

Configuration

Step-by-Step Procedure

To configure the TCP-Reset parameter for a zone:

1. Configure a security zone.

   ```
   [edit]
   user@host# edit security zones security-zone ABC
   ```

2. Configure the TCP-Reset parameter for the zone.

   ```
   [edit security zones security-zone ABC]
   user@host# set tcp-rst
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security zones** command.

Related Documentation

- Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter on page 1065

Copyright © 2016, Juniper Networks, Inc.

PART 12

Configuring Address Books and Address Sets

- Configuring Address, Address Books, and Address Sets on page 1069

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 50

Configuring Address, Address Books, and Address Sets

- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071
- Understanding Global Address Books on page 1071
- Configuring Addresses and Address Sets on page 1072
- Example: Configuring Address Books and Address Sets on page 1076
- Limitations of Addresses and Address Sets in a Security Policy on page 1081

Understanding Address Books

An address book is a collection of addresses and address sets. Junos OS allows you to configure multiple address books. Address books are like components, or building blocks, that are referenced in other configurations such as security policies or NAT. You can add addresses to address books or use the predefined addresses available to each address book by default.

Address book entries include addresses of hosts and subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated. These addresses can be any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, or Domain Name System (DNS) names.

- Predefined Addresses on page 1069
- Network Prefixes in Address Books on page 1070
- Wildcard Addresses in Address Books on page 1070
- DNS Names in Address Books on page 1070

Predefined Addresses

You can either create addresses or use any of the following predefined addresses that are available by default:

- **Any**—This address matches any IP address. When this address is used as a source or destination address in a policy configuration, it matches the source and destination address of any packet.

Copyright © 2016, Juniper Networks, Inc.

- **Any-ipv4**—This address matches any IPv4 address.
- **Any-ipv6**—This address matches any IPv6 address.

Network Prefixes in Address Books

You can specify addresses as network prefixes in the prefix/length format. For example, 203.0.113.0/24 is an acceptable address book address because it translates to a network prefix. However, 203.0.113.4/24 is not acceptable for an address book because it exceeds the subnet length of 24 bits. Everything beyond the subnet length must be entered as 0 (zero). In special scenarios, you can enter a hostname because it can use the full 32-bit address length.

An IPv6 address prefix is a combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form ipv6-prefix/prefix-length and represents a block of address space (or a network). The ipv6-prefix variable follows general IPv6 addressing rules. The /prefix-length variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 2001:db8::/32 is a possible IPv6 prefix. For more information on text representation of IPv6 addresses and address prefixes, see RFC 4291, *IP Version 6 Addressing Architecture.*

Wildcard Addresses in Address Books

Besides IP addresses and domain names, you can specify a wildcard address in an address book. A wildcard address is represented as A.B.C.D/wildcard-mask. The wildcard mask determines which of the bits in the IP address A.B.C.D should be ignored. For example, the source IP address 192.168.0.11/255.255.0.255 in a security policy implies that the security policy match criteria can discard the third octet in the IP address (symbolically represented as 192.168.*.11). Therefore, packets with source IP addresses such as 192.168.1.11 and 192.168.22.11 conform to the match criteria. However, packets with source IP addresses such as 192.168.0.1 and 192.168.1.21 do not satisfy the match criteria.

The wildcard address usage is not restricted to full octets only. You can configure any wildcard address. For example, the wildcard address 192.168. 7.1/255.255.7.255 implies that you need to ignore only the first 5 bits of the third octet of the wildcard address while making the policy match. If the wildcard address usage is restricted to full octets only, then wildcard masks with either 0 or 255 in each of the four octets only will be permitted.

DNS Names in Address Books

By default, you can resolve IPv4 and IPv6 addresses for a DNS. If IPv4 or IPv6 addresses are designated, you can resolve only those addresses by using the keywords ipv4-only and ipv6-only, respectively.

Before you can use domain names for address entries, you must configure the security device for DNS services. For information about DNS, see *DNS Overview*.

Related Documentation

- Understanding Global Address Books on page 1071
- Understanding Address Sets on page 1071
- Configuring Addresses and Address Sets on page 1072

Copyright © 2016, Juniper Networks, Inc.

Understanding Address Sets

An address book can grow to contain large numbers of addresses and become difficult to manage. You can create groups of addresses called address sets to manage large address books. Using address sets, you can organize addresses in logical groups and use them to easily configure other features, such as policies and NAT rules.

The predefined address set, **any**, which contains both **any-ipv4** and **any-ipv6** addresses, is automatically created for each security zone.

You can create address sets with existing users, or create empty address sets and later fill them with users. When creating address sets, you can combine IPv4 and IPv6 addresses, but the addresses must be in the same security zone.

You can also create an address set within an address set. This allows you to apply policies more effectively. For example, if you want to apply a policy to two address sets, **set1** and **set2**, instead of using two statements, you can use just one statement to apply the policy to a new address set, **set3**, that includes address sets **set1** and **set2**.

When you add addresses to policies, sometimes the same subset of addresses can be present in multiple policies, making it difficult to manage how policies affect each address entry. Reference an address set entry in a policy like an individual address book entry to allow you to manage a small number of address sets, rather than manage a large number of individual address entries.

Related Documentation

- Understanding Address Books on page 1069
- Configuring Addresses and Address Sets on page 1072
- Limitations of Addresses and Address Sets in a Security Policy on page 1081

Understanding Global Address Books

An address book called “global” is always present on your system. Similar to other address books, the global address book can include any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, or Domain Name System (DNS) names.

You can create addresses in the global address book or use the predefined addresses (any, any-ipv4, and any-ipv6). However, to use the addresses in the global address book, you do not need to attach the security zones to it. The global address book is available to all security zones that have no address books attached to them.

Global address books are used in the following cases:

- NAT configurations–NAT rules can use address objects only from the global address book. They cannot use addresses from zone-based address books.
- Global policies–Addresses used in a global policy must be defined in global address book. Global address book objects do not belong to any particular zone.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071
- Configuring Addresses and Address Sets on page 1072

Configuring Addresses and Address Sets

You can define addresses and address sets in an address book and then use them when configuring different features. You can also use predefined addresses **any**, **any-ipv4**, and **any-ipv6** that are available by default. However, you cannot add the predefined address **any** to an address book.

After address books and sets are configured, they are used in configuring different features, such as security policies, security zones, and NAT.

- Addresses and Address Sets on page 1072
- Address Books and Security Zones on page 1073
- Address Books and Security Policies on page 1073
- Address Books and NAT on page 1075

Addresses and Address Sets

You can define IPv4 addresses, IPv6 addresses, wildcard addresses, or Domain Name System (DNS) names as address entries in an address book.

The following sample address book called **book1** contains different types of addresses and address sets. Once defined, you can leverage these addresses and address sets when you configure security zones, policies, or NAT rules.

```
[edit security address-book book1]
user@host# set address a1 203.0.113.1
user@host# set address a2 203.0.113.4/30
user@host# set address a4 2001:db8::/32
user@host# set address a5 2001:db8:1::1/127
user@host# set address example dns-name www.example.com
user@host# set address-set set1 address a1
user@host# set address-set set1 address a2
user@host# set address-set set1 address a2
user@host# set address-set set2 address bbc
```

When defining addresses and address sets, follow these guidelines:

- Address sets can only contain address names that belong to the same security zone.
- Address names **any**, **any-ipv4** and **any-ipv6** are reserved; you cannot use them to create any addresses.
- Addresses and address sets in the same zone must have distinct names.

Copyright © 2016, Juniper Networks, Inc.

- Address names cannot be the same as address set names. For example, if you configure an address with the name **add1**, do not create the address set with the name **add1**.
- When deleting an individual address book entry from the address book, you must remove the address (wherever it is referred) from all the address sets; otherwise, the system will cause a commit failure.

Address Books and Security Zones

A security zone is a logical group of interfaces with identical security requirements. You attach security zones to address books that contain entries for the addressable networks and end hosts (and, thus, users) belonging to the zone.

A zone can use two address books at a time—the global address book and the address book that the zone is attached to. When a security zone is not attached to any address book, it automatically uses the global address book. Thus, when a security zone is attached to an address book, the system looks up addresses from this attached address book; otherwise, the system looks up addresses from the default global address book. The global address book is available to all security zones by default; you do not need to attach zones to the global address book.

The following guidelines apply when attaching security zones to address books:

- Addresses attached to a security zone conform to the security requirements of the zone.
- The address book that you attach to a security zone must contain all IP addresses that are reachable within that zone.
- When you configure policies between two zones, you must define the addresses for each of the zone's address books.
- Addresses in a user-defined address book have a higher lookup priority than addresses in the global address book. Thus, for a security zone that is attached to a user-defined address book, the system searches the user-defined address book first; if no address is found, then it searches the global address book.

Address Books and Security Policies

Addresses and address sets are used when specifying the match criteria for a policy. Before you can configure policies to permit, deny, or tunnel traffic to and from individual hosts and subnets, you must make entries for them in address books. You can define different types of addresses, such as IPv4 addresses, IPv6 addresses, wildcard addresses, and DNS names, as match criteria for security policies.

Policies contain both source and destination addresses. You can refer to an address or address set in a policy by the name you give to it in the address book attached to the zone specified in the policy.

- When traffic is sent to a zone, the zone and address to which the traffic is sent are used as the destination zone and address-matching criteria in policies.

Copyright © 2016, Juniper Networks, Inc.

- When traffic is sent from a zone, the zone and address from which the traffic is sent are used as the source zone and address-matching criteria in policies.

Addresses Available for Security Policies

When configuring the source and destination addresses for a policy rule, you can type a question mark in the CLI to list all the available addresses that you can choose from.

You can use the same address name for different addresses that are in different address books. However, the CLI lists only one of these addresses—the address that has the highest lookup priority.

For example, suppose you configure addresses in two address books—**global** and **book1**. Then, display the addresses that you can configure as source or destination addresses in a policy (see Table 73).

Table 73: Available Addresses Displayed in the CLI

Addresses Configured	Addresses Displayed in the CLI
[edit security address-book] **set global address a1 203.0.113.0/24;** **set global address a2 198.51.100.0/24;** **set global address a3 192.0.2.0/24;** **set book1 address a1 203.113..128/25;**	[edit security policies from-zone trust to-zone untrust] user@host# **set policy p1 match set match source-address ?** Possible completions: [Open a set of values a1 The address in address book book1 a2 The address in address book global a3 The address in address book global any Any IPv4 or IPv6 address any-ipv4 Any IPv4 address any-ipv6 Any IPv6 address

The addresses displayed in this example illustrate:

- Addresses in a user-defined address book have a higher lookup priority than addresses in the global address book.
- Addresses in a global address book have a higher priority than the predefined addresses **any**, **any-ipv4**, and **any-ipv6**.
- When the same address name is configured for two or more different addresses, only the highest priority address, based on the address lookup, is available. In this example, the CLI displays address **a1** from **book1** (203.0.113.128/25) because that address has a higher lookup priority than the global address **a1** (203.0.113.0/24).

Applying Policies to Address Sets

When you specify an address set in policies, Junos OS applies the policies automatically to each address set member, so you do not have to create them one by one for each address. Also, if an address set is referenced in a policy, the address set cannot be removed without removing its reference in the policy. It can, however, be edited.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Consider that for each address set, the system creates individual rules for its members. It creates an internal rule for each member in the group as well as for each service configured for each user. If you configure address books without taking this into account, you can exceed the number of available policy resources, especially if both the source and destination addresses are address groups and the specified service is a service group.

Figure 64 shows how policies are applied to address sets.

Figure 64: Applying Policies to Address Sets

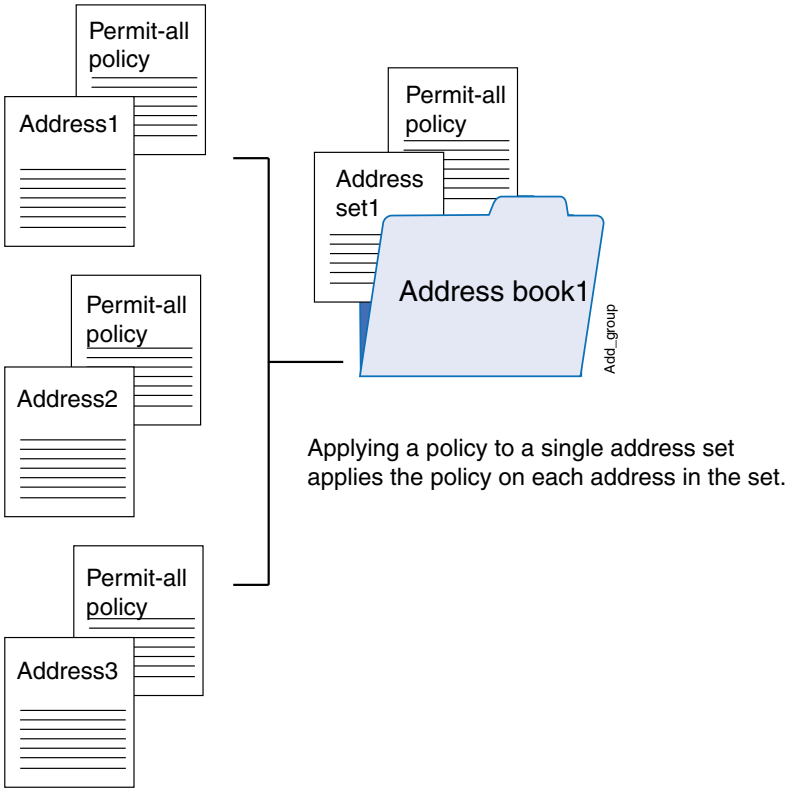

Address Books and NAT

Once you define addresses in address books, you can specify them in the source, destination, or static NAT rules. It is simpler to specify meaningful address names instead of IP prefixes as source and destination addresses in the NAT rule configuration. For example, instead of specifying 10.208.16.0/22 as source address, you can specify an address called **local** that includes address 10.208.16.0/22.

You can also specify address sets in NAT rules, allowing you to add multiple addresses within an address set and therefore manage a small number of address sets, rather than manage a large number of individual address entries. When you specify an address set in a NAT rule, Junos OS applies the rule automatically to each address set member, so you do not have to specify each address one by one.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The following address and address set types are not supported in NAT rules—wildcard addresses, DNS names, and a combination of IPv4 and IPv6 addresses.

When configuring address books with NAT, follow these guidelines:

- In a NAT rule, you can specify addresses from a global address book only. User-defined address books are not supported with NAT.
- You can configure an address set as a source address name in a source NAT rule. However, you cannot configure an address set as a destination address name in a destination NAT rule.

  The following sample NAT statements show the address and address set types that are supported with source and destination NAT rules:

  ```
  [edit security nat source rule-set src-nat rule src-rule1]
  set match source-address 2001:db8.1::/64
  set match source-address-name add1
  set match source-address-name add-set1
  set match destination-address 2001:db8.2::/64
  set match destination-address-name add2
  set match destination-address-name add-set2

  [edit security nat destination rule-set dst-nat rule dst-rule1]
  set match source-address 2001:db8.1::/64
  set match source-address-name add2
  set match source-address-name add-set2
  set match destination-address-name add1
  ```

- In a static NAT rule, you cannot configure an address set as a source or destination address name. The following sample NAT statements show the types of address that are supported with static NAT rules:

  ```
  [edit security nat static rule-set stat]
  set rule stat-rule1 match destination-address 203.0.113.0/24
  set rule stat-rule2 match destination-address-name add1
  ```

Related Documentation

- Understanding Address Books on page 1069
- Understanding Global Address Books on page 1071
- Understanding Address Sets on page 1071

Example: Configuring Address Books and Address Sets

This example shows how to configure addresses and address sets in address books. It also shows how to attach address books to security zones.

- Requirements on page 1077
- Overview on page 1077

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1078
- Verification on page 1080

Requirements

Before you begin:

- Configure the Juniper Networks security devices for network communication.
- Configure network interfaces on server and member devices. See the *Interfaces Feature Guide for Security Devices*.
- Configure Domain Name System (DNS) services. For information about DNS, see *DNS Overview*.

Overview

In this example, you configure an address book with addresses and address sets (see Figure 65) to simplify configuring your company's network. You create an address book called **Eng-dept** and add addresses of members from the Engineering department. You create another address book called **Web** and add a DNS name to it. Then you attach a security zone trust to the **Eng-dept** address book and security zone untrust to the **Web** address book. You also create address sets to group software and hardware addresses in the Engineering department. You plan to use these addresses as source address and destination addresses in your future policy configurations.

In addition, you add an address to the global address book, to be available to any security zone that has no address book attached to it.

Copyright © 2016, Juniper Networks, Inc.

Figure 65: Configuring Addresses and Address Sets

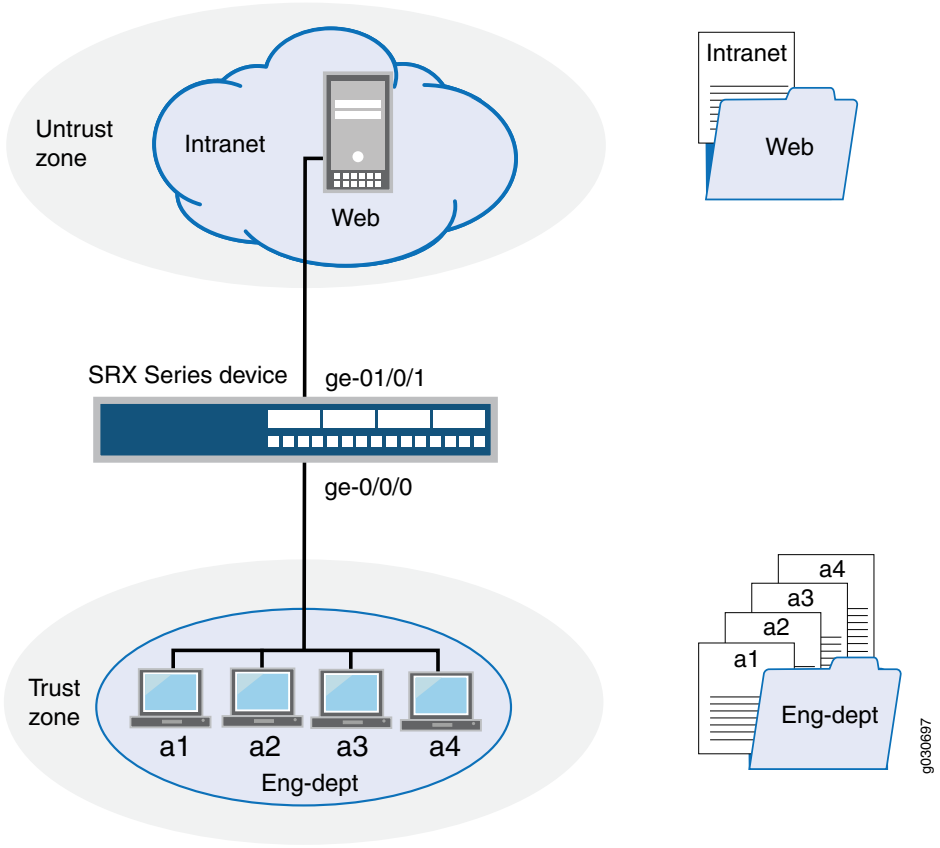

Configuration

CLI Quick Configuration

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

```
set security zones security-zone trust interfaces ge-0/0/0
set security zones security-zone untrust interfaces ge-0/0/1
set security address-book Eng-dept address a1 203.0.113.1
set security address-book Eng-dept address a2 203.0.113.2
set security address-book Eng-dept address a3 203.0.113.3
set security address-book Eng-dept address a4 203.0.113.4
set security address-book Eng-dept address-set sw-eng address a1
set security address-book Eng-dept address-set sw-eng address a2
set security address-book Eng-dept address-set hw-eng address a3
set security address-book Eng-dept address-set hw-eng address a4
set security address-book Eng-dept attach zone trust
set security address-book Web address Intranet dns-name www-int.example.com
set security address-book Web attach zone untrust
set security address-book global address g1 198.51.100.2/24
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure addresses and address sets:

1. Create security zones and assign interfaces to them.

```
[edit]
user@host# set security zones security-zone trust interfaces ge-0/0/0
user@host# set security zones security-zone untrust interfaces ge-0/0/1
```

2. Create an address book and define addresses in it.

```
[edit security address-book Eng-dept ]
user@host# set address a1 203.0.113.1
user@host# set address a2 203.0.113.2
user@host# set address a3 203.0.113.3
user@host# set address a4 203.0.113.4
```

3. Create address sets.

```
[edit security address-book Eng-dept]
user@host# set address-set sw-eng address a1
user@host# set address-set sw-eng address a2
user@host# set address-set hw-eng address a3
user@host# set address-set hw-eng address a4
```

4. Attach the address book to a security zone.

```
[edit security address-book Eng-dept]
user@host# set attach zone trust
```

5. Create another address book and attach it to a security zone.

```
[edit security address-book Web ]
user@host# set address Intranet dns-name www-int.example.com
user@host# set attach zone untrust
```

6. Define an address in the global address book.

```
[edit]
user@host# set security address-book global address g1 198.51.100.2/24
```

Results

From configuration mode, confirm your configuration by entering the **show security zones** and **show security address-book** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security zones
  security-zone untrust {
    interfaces {
      ge-0/0/1.0;
    }
  }
    security-zone trust {
      interfaces {
        ge-0/0/0.0;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
[edit]
user@host# show security address-book
    Eng-dept {
      address a1 203.0.113.1/24;
      address a2 203.0.113.2/24;
      address a3 203.0.113.3/24;
      address a4 203.0.113.4/24;
      address-set sw-eng {
        address a1;
        address a2;
      }
      address-set hw-eng {
        address a3;
        address a4;
      }
      attach {
        zone trust;
      }
    }
      Web {
        address Intranet {
          dns-name www-int.example.com ;
        }
        attach {
          zone untrust;
        }
      }
      global {
        address g1 198.51.100.2/24;
      }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying Address Book Configuration on page 1080
- Verifying Global Address Book Configuration on page 1080

Verifying Address Book Configuration

Purpose Display information about configured address books and addresses.

Action From configuration mode, enter the **show security address-book** command.

Verifying Global Address Book Configuration

Purpose Display information about configured addresses in the global address book.

Action From configuration mode, enter the **show security address-book global** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

Limitations of Addresses and Address Sets in a Security Policy

On SRX Series devices, one policy can reference multiple address sets, multiple address entries, or both. One address set can reference a maximum of 1024 address entries and a maximum of 256 address sets. There is a limit to the number of address objects that a policy can reference; the maximum number of address objects per policy is 1024.

For example, if policy p1 references 10 address entries and 10 address sets (which in turn reference 3 address entries), then the number of address objects under this policy is 40 (10 address entries + [10 address sets x 3 address entries] = 40).

Note that every IPv6 address entry is equal to 4 IPv4 address entries. For example, a policy configured for 1000 IPv4 address entries and 5 IPv6 address entries has 1020 address objects (1000 + [5 x4] = 1020), which is within the 1024 value, and can be committed. However, a policy configured for 1000 IPv4 address entries and 7 IPv6 address entries has 1028 address objects (1000 +[7 x 4] = 1028), which exceeds the 1024 value, cannot be committed, and consequently generates an error message.

Related Documentation
- Understanding Address Sets on page 1071
- Example: Configuring Address Books and Address Sets on page 1076

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 13

Configuring Security Policies

- Enforcing Transit Traffic Rules by Configuring Security Policies on page 1085
- Configuring Negated Addresses on page 1109
- Configuring Global Security Policy on page 1115
- Managing Security Policy Activation By Configuring Schedulers on page 1123
- Configuring User Role Firewall Security Policies on page 1127
- Setting Security Policy Reorder on page 1149
- Monitoring and Troubleshooting Security Policies on page 1153
- Handling Security Policy Violations on page 1165

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 51

Enforcing Transit Traffic Rules by Configuring Security Policies

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Understanding Security Policies for Self Traffic on page 1092
- Security Policies Configuration Overview on page 1093
- Configuring Policies Using the Firewall Wizard on page 1094
- Example: Configuring a Security Policy to Permit or Deny All Traffic on page 1094
- Example: Configuring a Security Policy to Permit or Deny Selected Traffic on page 1098
- Example: Configuring a Security Policy to Permit or Deny Wildcard Address Traffic on page 1102
- Example: Configuring a Security Policy to Redirect Traffic Logs to an External System Log Server on page 1105

Security Policies Overview

To secure their business, organizations must control access to their LAN and their resources. Security policies are commonly used for this purpose. Secure access is required both within the company across the LAN and in its interactions with external networks such as the Internet. Junos OS provides powerful network security features through its stateful firewall, application firewall, and user identity firewall. All three types of firewall enforcement are implemented through security policies. The stateful firewall policy syntax is widened to include additional tuples for the application firewall and the user identity firewall.

In a Junos OS stateful firewall, the security policies enforce rules for transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on traffic as it passes through the firewall. From the perspective of security policies, the traffic enters one security zone and exits another security zone. This combination of a *from-zone* and *to-zone* is called a *context*. Each context contains an *ordered list* of policies. Each policy is processed in the order that it is defined within a context.

Copyright © 2016, Juniper Networks, Inc.

A security policy, which can be configured from the user interface, controls the traffic flow from one zone to another zone by defining the kind(s) of traffic permitted from specified IP sources to specified IP destinations at scheduled times.

Policies allow you to deny, permit, reject (deny and send a TCP RST or ICMP port unreachable message to the source host), encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another. You decide which users and what data can enter and exit, and when and where they can go.

NOTE: For an SRX Series device that supports virtual systems, policies set in the root system do not affect policies set in virtual systems.

An SRX Series device secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another.

Logging capability can also be enabled with security policies during session initialization (**session-init**) or session close (**session-close**) stage.

- To view logs from denied connections, enable log on **session-init**.
- To log sessions after their conclusion/tear-down, enable log on **session-close**.

NOTE: Session log is enabled at real time in the flow code which impacts the user performance. If both session-close and session-init are enabled, performance is further degraded as compared to enabling session-init only.

For SRX Series branch devices, a factory-default security policy is provided that:

- Allows all traffic from the trust zone to the untrust zone.
- Allows all traffic between trusted zones, that is from the trust zone to intrazone trusted zones.
- Denies all traffic from the untrust zone to the trust zone.

Through the creation of policies, you can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations at scheduled times.

At the broadest level, you can allow all kinds of traffic from any source in one zone to any destination in all other zones without any scheduling restrictions. At the narrowest level, you can create a policy that allows only one kind of traffic between a specified host in one zone and another specified host in another zone during a scheduled interval of time. See Figure 66.

Copyright © 2016, Juniper Networks, Inc.

Figure 66: Security Policy

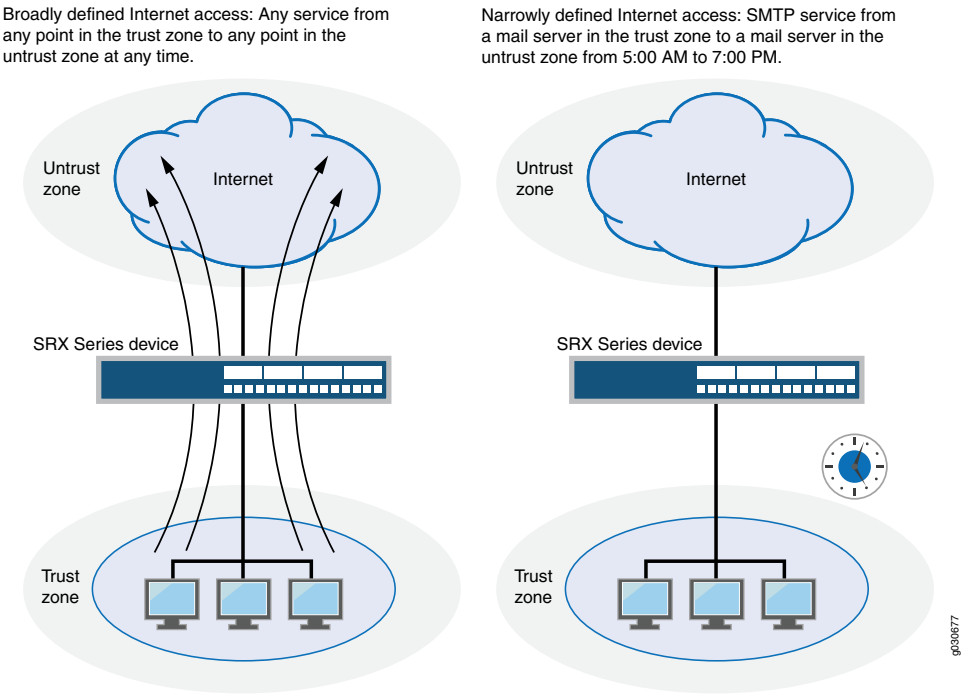

Every time a packet attempts to pass from one zone to another or between two interfaces bound to the same zone, the device checks for a policy that permits such traffic (see "Understanding Security Zones" on page 1050 and "Policy Application Sets Overview" on page 1172). To allow traffic to pass from one security zone to another—for example, from zone A to zone B—you must configure a policy that permits zone A to send traffic to zone B. To allow traffic to flow the other way, you must configure another policy permitting traffic from zone B to zone A.

To allow data traffic to pass between zones, you must configure firewall policies.

Related Documentation

- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Security Policies Configuration Overview on page 1093
- Understanding Security Policy Ordering on page 1149
- Security Zones and Interfaces Overview on page 1049

Understanding Security Policy Rules

The security policy applies the security rules to the transit traffic within a context (**from-zone** to **to-zone**). Each policy is uniquely identified by its name. The traffic is classified by matching its source and destination zones, the source and destination addresses, and the application that the traffic carries in its protocol headers with the policy database in the data plane.

Copyright © 2016, Juniper Networks, Inc.

Each policy is associated with the following characteristics:

- A source zone
- A destination zone
- One or many source address names or address set names
- One or many destination address names or address set names
- One or many application names or application set names

These characteristics are called the *match criteria*. Each policy also has actions associated with it: permit, deny, reject, count, log, and VPN tunnel. You have to specify the match condition arguments when you configure a policy, source address, destination address, and application name.

You can specify to configure a policy with IPv4 or IPv6 addresses using the wildcard entry **any**. When flow support is not enabled for IPv6 traffic, **any** matches IPv4 addresses. When flow support is enabled for IPv6 traffic, **any** matches both IPv4 and IPv6 addresses. To enable flow-based forwarding for IPv6 traffic, use the **set security forwarding-options family inet6 mode flow-based** command. You can also specify the wildcard **any-ipv4** or **any-ipv6** for the source and destination address match criteria to include only IPv4 or only IPv6 addresses, respectively.

When flow support for IPv6 traffic is enabled, the maximum number of IPv4 or IPv6 addresses that you can configure in a security policy is based on the following match criteria:

- Number_of_src_IPv4_addresses + number_of_src_IPv6_addresses * 4 <= 1024
- Number_of_dst_IPv4_addresses + number_of_dst_IPv6_addresses * 4 <= 1024

Thr reason for the match criteria is that an IPv6 address uses four times the memory space that an IPv4 address uses.

NOTE: You can configure a security policy with IPv6 addresses only if flow support for IPv6 traffic is enabled on the device.

If you do not want to specify a specific application, enter **any** as the default application. To look up the default applications, from configuration mode, enter **show groups junos-defaults | find applications (predefined applications)**. For example, if you do not supply an application name, the policy is installed with the application as a wildcard (default). Therefore, any data traffic that matches the rest of the parameters in a given policy would match the policy regardless of the application type of the data traffic.

NOTE: If a policy is configured with multiple applications, and more than one of the applications match the traffic, then the application that best meets the match criteria is selected.

Copyright © 2016, Juniper Networks, Inc.

The action of the first policy that the traffic matches is applied to the packet. If there is no matching policy, the packet is dropped. Policies are searched from top to bottom, so it is a good idea to place more specific policies near the top of the list. You should also place IPsec VPN tunnel policies near the top. Place the more general policies, such as one that would allow certain users access to all Internet applications, at the bottom of the list. For example, place deny-all or reject-all policies at the bottom after all of the specific policies have been parsed before and legitimate traffic has been allowed/count/logged.

NOTE: Support for IPv6 addresses added in Junos OS Release 10.2 and support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

Policies are looked up during flow processing after firewall filters and screens have been processed and route look up has been completed by the Services Processing Unit (SPU) (for high-end SRX Series devices). Policy look up determines the destination zone, destination address, and egress interface.

When you are creating a policy, the following policy rules apply:

- Security policies are configured in a **from-zone** to **to-zone** direction. Under a specific zone direction, each security policy contains a name, match criteria, an action, and miscellaneous options.
- The policy name, match criteria, and action are required.
- The policy name is a keyword.
- The source address in the match criteria is composed of one or more address names or address set names in the **from-zone**.
- The destination address of the match criteria is composed of one or more address names or address set names in the **to-zone**.
- The application name in the match criteria is composed of the name of one or more applications or application sets.
- One of the following actions is required: permit, deny, or reject.
- Accounting and auditing elements can be specified: count and log.
- You can enable logging at the end of a session with the **session-close** command, or at the beginning of the session with the **session-init** command.
- When the count alarm is turned on, specify alarm thresholds in bytes per second or kilobytes per minute.
- You cannot specify **global** as either the **from-zone** or the **to-zone** except under following condition:

  Any policy configured with the **to-zone** as a global zone must have a single destination address to indicate that either static NAT or incoming NAT has been configured in the policy.

Copyright © 2016, Juniper Networks, Inc.

- In SRX Series Services Gateways, the policy permit option with NAT is simplified. Each policy will optionally indicate whether it allows NAT translation, does not allow NAT translation, or does not care.
- Address names cannot begin with the following reserved prefixes. These are used only for address NAT configuration:
  - **static_nat_**
  - **incoming_nat_**
  - **junos_**
- Application names cannot begin with the **junos_** reserved prefix.

Understanding Wildcard Addresses

Source and destination addresses are two of the five match criteria that should be configured in a security policy. You can now configure wildcard addresses for the source and destination address match criteria in a security policy. A wildcard address is represented as A.B.C.D/wildcard-mask. The wildcard mask determines which of the bits in the IP address A.B.C.D should be ignored by the security policy match criteria. For example, the source IP address 192.168.0.11/255.255.0.255 in a security policy implies that the security policy match criteria can discard the third octet in the IP address (symbolically represented as 192.168.*.11). Therefore, packets with source IP addresses such as 192.168.1.11 and 192.168.22.11 conform to the match criteria. However, packets with source IP addresses such as 192.168.0.1 and 192.168.1.21 do not satisfy the match criteria.

The wildcard address usage is not restricted to full octets only. You can configure any wildcard address. For example, the wildcard address 192.168. 7.1/255.255.7.255 implies that you need to ignore only the first 5 bits of the third octet of the wildcard address while making the policy match. If the wildcard address usage is restricted to full octets only, then wildcard masks with either 0 or 255 in each of the four octets only will be permitted.

NOTE: **The first octet of the wildcard mask should be greater than 128. For example, a wildcard mask represented as 0.255.0.255 or 1.255.0.255 is invalid.**

A wildcard security policy is a simple firewall policy that allows you to permit, deny, and reject the traffic trying to cross from one security zone to another. You should not configure security policy rules using wildcard addresses for services such as Unified Threat Management (UTM).

Copyright © 2016, Juniper Networks, Inc.

NOTE: Only Intrusion and Prevention (IDP) for IPv6 sessions is supported for all high-end SRX Series devices. UTM for IPv6 sessions is not supported. If your current security policy uses rules with the IP address wildcard any, and UTM features are enabled, you will encounter configuration commit errors because UTM features do not yet support IPv6 addresses. To resolve the errors, modify the rule returning the error so that the any-ipv4 wildcard is used; and create separate rules for IPv6 traffic that do not include UTM features.

Configuring wildcard security policies on a device affects performance and memory usage based on the number of wildcard policies configured per from-zone and to-zone context. Therefore, you can only configure a maximum of 480 wildcard policies for a specific from-zone and to-zone context.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Elements on page 1091
- Security Policies Configuration Overview on page 1093
- Understanding Security Policy Ordering on page 1149

Understanding Security Policy Elements

A security policy is a set of statements that controls traffic from a specified source to a specified destination using a specified service. A policy permits, denies, or tunnels specified types of traffic unidirectionally between two points.

Each policy consists of:

- A unique name for the policy.
- A **from-zone** and a **to-zone**, for example: user@host# **set security policies from-zone untrust to-zone untrust**
- A set of match criteria defining the conditions that must be satisfied to apply the policy rule. The match criteria are based on a source IP address, destination IP address, and applications. The user identify firewall provides greater granularity by including an additional tuple, source-identity, as part of the policy statement.
- A set of actions to be performed in case of a match—permit, deny, or reject.
- Accounting and auditing elements—counting, logging, or structured system logging.

If the SRX Series receives a packet that matches those specifications, it performs the action specified in the policy.

Security policies enforce a set of rules for transit traffic, identifying which traffic can pass through the firewall and the actions taken on the traffic as it passes through the firewall. Actions for traffic matching the specified criteria include permit, deny, reject, log, or count.

For SRX Series branch devices, a factory default security policy is provided that:

Copyright © 2016, Juniper Networks, Inc.

- permits all traffic from the trust zone to the untrust zone.
- denies all traffic from the untrust zone to the trust zone.
- allows all traffic from the trust zone to the untrust zone.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Ordering on page 1149
- Security Policies Configuration Overview on page 1093

Understanding Security Policies for Self Traffic

Security policies are configured on the devices to apply services to the traffic flowing through the device. For example UAC and UTM policies are configured to apply services to the transient traffic.

Self-traffic or host traffic, is the host-inbound traffic; that is, the traffic terminating on the device or the host-outbound traffic that is the traffic originating from the device. You can now configure policies to apply services on self traffic. Services like the SSL stack service that must terminate the SSL connection from a remote device and perform some processing on that traffic, IDP services on host-inbound traffic, or IPsec encryption on host-outbound traffic must be applied through the security policies configured on self-traffic.

When you configure a security policy for self-traffic, the traffic flowing through the device is first checked against the policy, then against the **host-inbound-traffic** option configured for the interfaces bound to the zone.

You can configure the security policy for self-traffic to apply services to self-traffic. The host-outbound policies will work only in cases where the packet that originated in the host device goes through the flow and the incoming interface of this packet is set to local.

The advantages of using the self-traffic are:

- You can leverage most of the existing policy or flow infrastructure used for the transit traffic.
- You do not need a separate IP address to enable any service.
- You can apply services or policies to any host-inbound traffic with the destination IP address of any interface on the device.

NOTE: You can configure the security policy for self-traffic with relevant services only. For example, it is not relevant to configure the fwauth service on host-outbound traffic, and gprs-gtp services are not relevant to the security policies for self-traffic.

Copyright © 2016, Juniper Networks, Inc.

The security policies for the self traffic are configured under the new default security zone called the *junos-host* zone. The junos-host zone will be part of the junos-defaults configuration, so users can delete it. The existing zone configurations such as interfaces, screen, tcp-rst, and host-inbound-traffic options are not meaningful to the junos-host zone. Therefore there is no dedicated configuration for the junos-host zone.

NOTE: You can use host-inbound-traffic to control incoming connections to a device; however it does not restrict traffic going out of the device. Whereas, junos-host-zone allows you to select the application of your choice and also restrict outgoing traffic. For example, services like NAT, IDP, UTM, and so forth can now be enabled for traffic going in or out of the SRX Series device using junos-host-zone.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Ordering on page 1149

Security Policies Configuration Overview

You must complete the following tasks to create a security policy:

1. Create zones. See "Example: Creating Security Zones" on page 1051.
2. Configure an address book with addresses for the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
3. Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.
4. Create the policy. See "Example: Configuring a Security Policy to Permit or Deny All Traffic" on page 1094, "Example: Configuring a Security Policy to Permit or Deny Selected Traffic" on page 1098, and "Example: Configuring a Security Policy to Permit or Deny Wildcard Address Traffic" on page 1102.
5. Create schedulers if you plan to use them for your policies. See "Example: Configuring Schedulers for a Daily Schedule Excluding One Day" on page 1124.

The Firewall Policy Wizard enables you to perform basic security policy configuration. For more advanced configuration, use the J-Web interface or the CLI.

Related Documentation

- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Troubleshooting Security Policies on page 1162

Copyright © 2016, Juniper Networks, Inc.

Configuring Policies Using the Firewall Wizard

The Firewall Policy Wizard enables you to perform basic security policy configuration. For more advanced configuration, use the J-Web interface or the CLI.

To configure policies using the Firewall Policy Wizard:

1. Select **Configure>Tasks>Configure FW Policy** in the J-Web interface.
2. Click the Launch Firewall Policy Wizard button to launch the wizard.
3. Follow the prompts in the wizard.

The upper-left area of the wizard page shows where you are in the configuration process. The lower-left area of the page shows field-sensitive help. When you click a link under the Resources heading, the document opens in your browser. If the document opens in a new tab, be sure to close only the tab (not the browser window) when you close the document.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Ordering on page 1149

Example: Configuring a Security Policy to Permit or Deny All Traffic

This example shows how to configure a security policy to permit or deny all traffic.

- Requirements on page 1094
- Overview on page 1094
- Configuration on page 1095
- Verification on page 1097

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book and create addresses for use in the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.

Overview

In the Junos OS, security policies enforce rules for transit traffic, in terms of what traffic can pass through the device, and the actions that need to take place on traffic as it passes through the device. From the perspective of security policies, the traffic enters one security

Copyright © 2016, Juniper Networks, Inc.

zone and exits another security zone. In this example, you configure the trust and untrust interfaces, ge-0/0/2 and ge-0/0/1. See Figure 67.

Figure 67: Permitting All Traffic

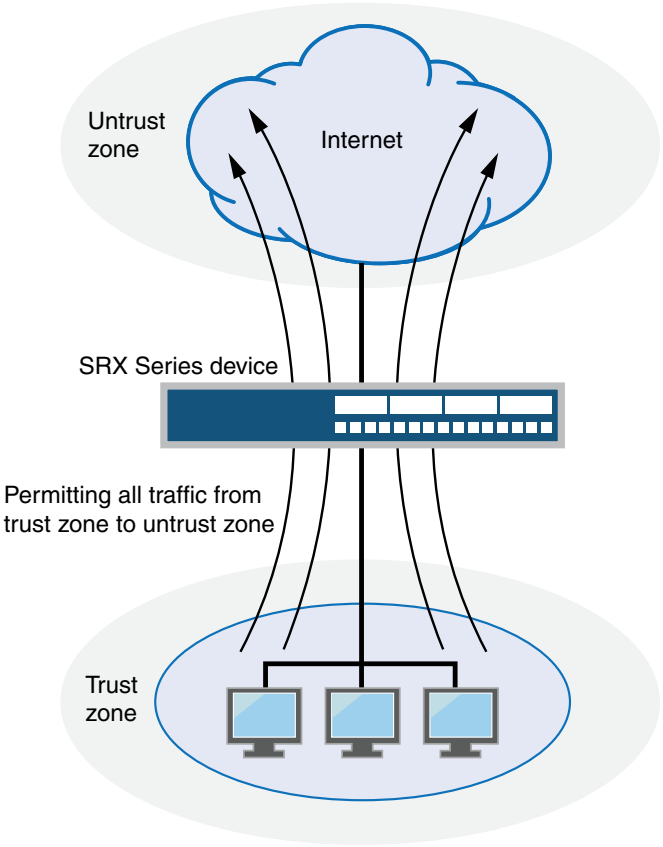

This configuration example shows how to:

- Permit or deny all traffic from the trust zone to the untrust zone but block everything from the untrust zone to the trust zone.
- Permit or deny selected traffic from a host in the trust zone to a server in the untrust zone at a particular time.

Configuration

CLI Quick Configuration

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

```
set security zones security-zone trust interfaces ge-0/0/2 host-inbound-traffic
    system-services all
set security zones security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
    system-services all
set security policies from-zone trust to-zone untrust policy permit-all match
    source-address any
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone untrust policy permit-all match
  destination-address any
set security policies from-zone trust to-zone untrust policy permit-all match application
  any
set security policies from-zone trust to-zone untrust policy permit-all then permit
set security policies from-zone untrust to-zone trust policy deny-all match source-address
  any
set security policies from-zone untrust to-zone trust policy deny-all match
  destination-address any
set security policies from-zone untrust to-zone trust policy deny-all match application
  any
set security policies from-zone untrust to-zone trust policy deny-all then deny
```

Step-by-Step Procedure

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

To configure a security policy to permit or deny all traffic:

1. Configure the interfaces and security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/0/2 host-inbound-traffic
  system-services all
user@host# set security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
  system-services all
```

2. Create the security policy to permit traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-all match source-address any
user@host# set policy permit-all match destination-address any
user@host# set policy permit-all match application any
user@host# set policy permit-all then permit
```

3. Create the security policy to deny traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy deny-all match source-address any
user@host# set policy deny-all match destination-address any
user@host# set policy deny-all match application any
user@host# set policy deny-all then deny
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

NOTE: The configuration example is a default permit-all from the trust zone to the untrust zone.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
```


Copyright © 2016, Juniper Networks, Inc.

```
    policy permit-all {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone untrust to-zone trust {
    policy deny-all {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        deny;
      }
    }
  }
```

```
user@host# show security zones
security-zone trust {
  interfaces {
    ge-0/0/2.0 {
      host-inbound-traffic {
        system-services {
          all;
        }
      }
    }
  }
}
security-zone untrust {
  interfaces {
    ge-0/0/1.0 {
      host-inbound-traffic {
        system-services {
          all;
        }
      }
    }
  }
}
```

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

Verification

Verifying Policy Configuration

Purpose Verify information about security policies.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security policies detail** command to display a summary of all security policies configured on the device.

Meaning The output displays information about policies configured on the system. Verify the following information:

- From and to zones
- Source and destination addresses
- Match criteria

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Example: Configuring a Security Policy to Permit or Deny Selected Traffic

This example shows how to configure a security policy to permit or deny selected traffic.

- Requirements on page 1098
- Overview on page 1098
- Configuration on page 1099
- Verification on page 1101

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book and create addresses for use in the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.
- Permit traffic to and from trust and untrust zones. See "Example: Configuring a Security Policy to Permit or Deny All Traffic" on page 1094.

Overview

In Junos OS, security policies enforce rules for the transit traffic, in terms of what traffic can pass through the device, and the actions that need to take place on the traffic as it passes through the device. From the perspective of security policies, the traffic enters one security zone and exits another security zone. In this example, you configure a specific security policy to allow only e-mail traffic from a host in the trust zone to a server in the untrust zone. No other traffic is allowed. See Figure 68.

Copyright © 2016, Juniper Networks, Inc.

Figure 68: Permitting Selected Traffic

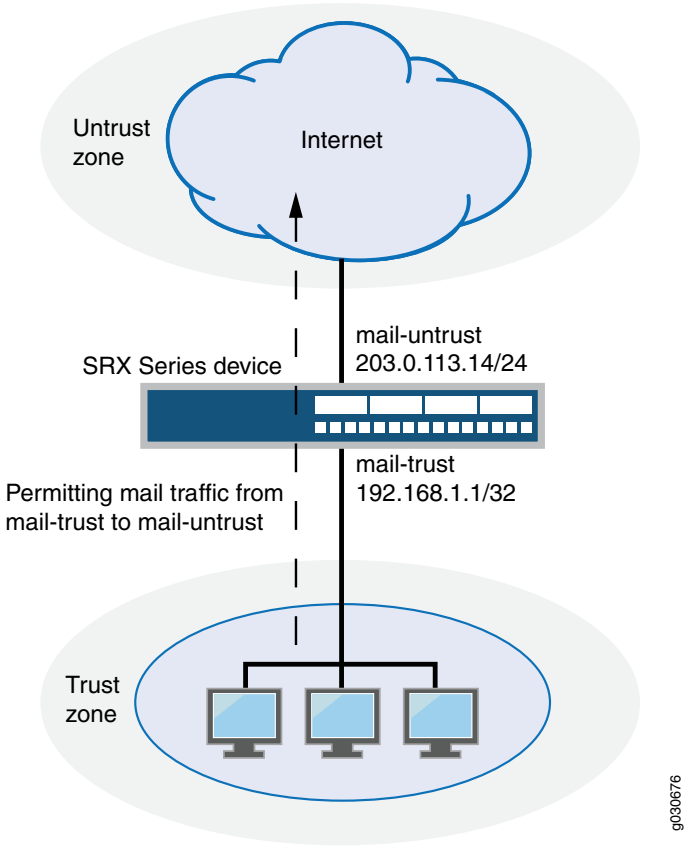

Configuration

CLI Quick Configuration

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

```
set security zones security-zone trust interfaces ge-0/0/2 host-inbound-traffic
    system-services all
set security zones security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
    system-services all
set security address-book book1 address mail-untrust 203.0.113.14/24
set security address-book book1 attach zone untrust
set security address-book book2 address mail-trust 192.168.1.1/32
set security address-book book2 attach zone trust
set security policies from-zone trust to-zone untrust policy permit-mail match
    source-address mail-trust
set security policies from-zone trust to-zone untrust policy permit-mail match
    destination-address mail-untrust
set security policies from-zone trust to-zone untrust policy permit-mail match application
    junos-mail
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone untrust policy permit-mail then permit
```

Step-by-Step Procedure

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

To configure a security policy to allow selected traffic:

1. Configure the interfaces and security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/0/2 host-inbound-traffic
  system-services all
user@host# set security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
  system-services all
```

2. Create address book entries for both the client and the server. Also, attach security zones to the address books.

```
[edit security address-book book1]
user@host# set address mail-untrust 203.0.113.14/24
user@host# set attach zone untrust

[edit security address-book book2]
user@host# set address mail-trust 192.168.1.1/32
user@host# set attach zone trust
```

3. Define the policy to permit mail traffic.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-mail match source-address mail-trust
user@host# set policy permit-mail match destination-address mail-untrust
user@host# set policy permit-mail match application junos-mail
user@host# set policy permit-mail then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy permit-mail {
    match {
      source-address mail-trust;
      destination-address mail-untrust;
      application junos-mail;
    }
    then {
      permit;
    }
  }
}

user@host# show security zones
security-zone trust {
  host-inbound-traffic {
```


Copyright © 2016, Juniper Networks, Inc.

```
    system-services {
      all;
    }
    interfaces {
      ge-0/0/2 {
        host-inbound-traffic {
          system-services {
            all;
          }
        }
      }
    }
  }
  security-zone untrust {
    interfaces {
      ge-0/0/1 {
        host-inbound-traffic {
          system-services {
            all;
          }
        }
      }
    }
  }
```

```
user@host# show security address-book
book1 {
  address mail-untrust 203.0.113.14/24;
  attach {
    zone untrust;
  }
}
  book2 {
    address mail-trust 192.168.1.1/32;
    attach {
      zone trust;
    }
  }
```

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

Verification

Verifying Policy Configuration

Purpose Verify information about security policies.

Action From operational mode, enter the **show security policies detail** command to display a summary of all security policies configured on the device.

Meaning The output displays information about policies configured on the system. Verify the following information:

- From and to zones

Copyright © 2016, Juniper Networks, Inc.

- Source and destination addresses
- Match criteria

Related Documentation

- Security Policies Overview on page 1085
- Example: Configuring a Security Policy to Permit or Deny All Traffic on page 1094

Example: Configuring a Security Policy to Permit or Deny Wildcard Address Traffic

This example shows how to configure a security policy to permit or deny wildcard address traffic.

- Requirements on page 1102
- Overview on page 1102
- Configuration on page 1102
- Verification on page 1104

Requirements

Before you begin:

- Understand wildcard addresses. See "Understanding Wildcard Addresses" on page 1090.
- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book and create addresses for use in the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.
- Permit traffic to and from trust and untrust zones. See "Example: Configuring a Security Policy to Permit or Deny All Traffic" on page 1094.
- Permit e-mail traffic to and from trust and untrust zones. See "Example: Configuring a Security Policy to Permit or Deny Selected Traffic" on page 1098

Overview

In the Junos operating system (Junos OS), security policies enforce rules for the transit traffic, in terms of what traffic can pass through the device, and the actions that need to take place on the traffic as it passes through the device. From the perspective of security policies, the traffic enters one security zone and exits another security zone. In this example, you configure a specific security to allow only wildcard address traffic from a host in the trust zone to the untrust zone. No other traffic is allowed.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** in configuration mode.

```
set security zones security-zone trust interfaces ge-0/0/2 host-inbound-traffic
  system-services all
set security zones security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
  system-services all
set security address-book book1 address wildcard-trust wildcard-address
  192.168.0.11/255.255.0.255
set security address-book book1 attach zone trust
set security policies from-zone trust to-zone untrust policy permit-wildcard match
  source-address wildcard-trust
set security policies from-zone trust to-zone untrust policy permit-wildcard match
  destination-address any
set security policies from-zone trust to-zone untrust policy permit-wildcard match
  application any
set security policies from-zone trust to-zone untrust policy permit-wildcard then permit
```

Step-by-Step Procedure

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

To configure a security policy to allow selected traffic:

1. Configure the interfaces and security zones.

   ```
   [edit security zones]
   user@host# set security-zone trust interfaces ge-0/0/2 host-inbound-traffic
     system-services all
   user@host# set security-zone untrust interfaces ge-0/0/1 host-inbound-traffic
     system-services all
   ```

2. Create an address book entry for the host and attach the address book to a zone.

   ```
   [edit security address-book book1]
   user@host# set address wildcard-trust wildcard-address 192.168.0.11/255.255.0.255
   user@host# set attach zone trust
   ```

3. Define the policy to permit wildcard address traffic.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy permit-wildcard match source-address wildcard-trust
   user@host# set policy permit-wildcard match destination-address any
   user@host# set policy permit-wildcard match application any
   user@host# set policy permit-wildcard then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show security policies** and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy permit-wildcard {
    match {
      source-address wildcard-trust;
```

Copyright © 2016, Juniper Networks, Inc.

```
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

```
user@host# show security zones
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
    interfaces {
      ge-0/0/2 {
        host-inbound-traffic {
          system-services {
            all;
          }
        }
      }
    }
  }
}
  security-zone untrust {
    interfaces {
      ge-0/0/1 {
        host-inbound-traffic {
          system-services {
            all;
          }
        }
      }
    }
  }
user@host# show security address-book
book1 {
  address wildcard-trust {
    wildcard-address 192.168.0.11/255.255.0.255;
  }
  attach {
    zone trust;
  }
}
```

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

Verification

Verifying Policy Configuration

Purpose Verify information about security policies.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security policies policy-name permit-wildcard detail** command to display details about the permit-wildcard security policy configured on the device.

Meaning The output displays information about the permit-wildcard policy configured on the system. Verify the following information:

- From and To zones
- Source and destination addresses
- Match criteria

Related Documentation

- Security Policies Configuration Overview on page 1093
- Understanding Security Policy Rules on page 1087
- Example: Configuring a Security Policy to Permit or Deny All Traffic on page 1094
- Example: Configuring a Security Policy to Permit or Deny Selected Traffic on page 1098

Example: Configuring a Security Policy to Redirect Traffic Logs to an External System Log Server

This example shows how to configure a security policy to send traffic logs generated on the device to an external system log server.

- Requirements on page 1105
- Overview on page 1105
- Configuration on page 1106
- Verification on page 1108

Requirements

This example uses the following hardware and software components:

- A client connected to an SRX5600 device at the interface ge-4/0/5
- A server connected to the SRX5600 device at the interface ge-4/0/1

  The logs generated on the SRX5600 device are stored in a Linux-based system log server.
- An SRX5600 device connected to the Linux-based server at the interface ge-4/0/4

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure a security policy on the SRX5600 device to send traffic logs, generated by the device during data transmission, to a Linux-based server. Traffic logs record details of every session. The logs are generated during session establishment

Copyright © 2016, Juniper Networks, Inc.

and termination between the source and the destination device that are connected to the SRX5600 device.

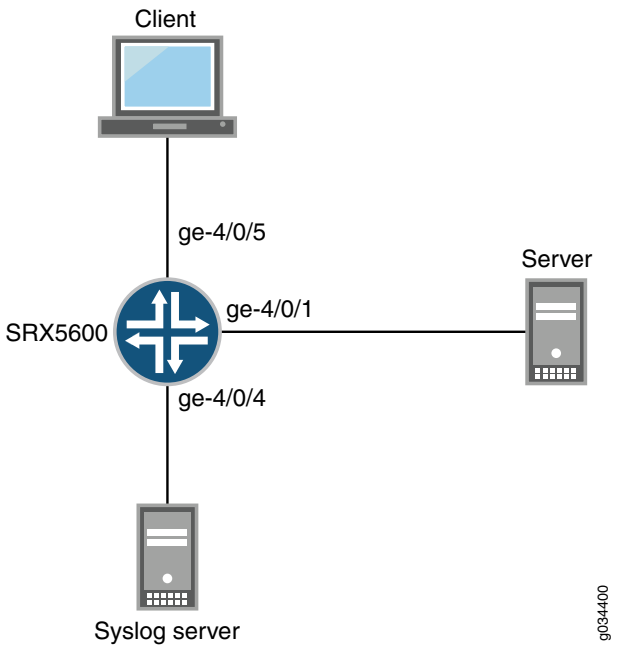

Configuration

CLI Quick Configuration

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

```
set security log source-address 127.0.0.1
set security log stream trafficlogs severity debug
set security log stream trafficlogs host 203.0.113.2
set security zones security-zone client host-inbound-traffic system-services all
set security zones security-zone client host-inbound-traffic protocols all
set security zones security-zone client interfaces ge-4/0/5.0
set security zones security-zone server host-inbound-traffic system-services all
set security zones security-zone server interfaces ge-4/0/4.0
set security zones security-zone server interfaces ge-4/0/1.0
set security policies from-zone client to-zone server policy match source-address any
set security policies from-zone client to-zone server policy match destination-address
    any
set security policies from-zone client to-zone server policy match application any
set security policies from-zone client to-zone server policy match then permit
set security policies from-zone client to-zone server policy match then log session-init
set security policies from-zone client to-zone server policy match then log session-close
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure a security policy to send traffic logs to an external system log server:

1. Configure security logs to transfer traffic logs generated at the SRX5600 device to an external system log server with the IP address 203.0.113.2. The IP address 127.0.0.1 is the loopback address of the SRX5600 device.

   ```
   [edit security log]
   user@host# set source-address 127.0.0.1
   user@host# set stream trafficlogs severity debug
   user@host# set stream trafficlogs host 203.0.113.2
   ```

2. Configure a security zone and specify the types of traffic and protocols that are allowed on interface ge-4/0/5.0 of the SRX5600 device.

   ```
   [edit security zones]
   user@host# set security-zone client host-inbound-traffic system-services all
   user@host# set security-zone client host-inbound-traffic protocols all
   user@host# set security-zone client interfaces ge-4/0/5.0
   ```

3. Configure another security zone and specify the types of traffic that are allowed on the interfaces ge-4/0/4.0 and ge-4/0/1.0 of the SRX5600 device.

   ```
   [edit security zones]
   user@host# set security-zone server host-inbound-traffic system-services all
   user@host# set security-zone server interfaces ge-4/0/4.0
   user@host# set security-zone server interfaces ge-4/0/1.0
   ```

4. Create a policy and specify the match criteria for that policy. The match criteria specifies that the device can allow traffic from any source, to any destination, and on any application.

   ```
   [edit security policies from-zone client to-zone server]
   user@host# set policy match source-address any
   user@host# set policy match destination-address any
   user@host# set policy match application any
   user@host# set policy match then permit any
   ```

5. Enable the policy to log traffic details at the beginning and at the end of the session.

   ```
   [edit security policies from-zone client to-zone server]
   user@host# set policy match then log session-init
   user@host# set policy match then log session-close
   ```

Results

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

```
[edit]
user@host# show security log
format syslog;
source-address 127.0.0.1;
stream trafficlogs {
   severity debug;
```

Copyright © 2016, Juniper Networks, Inc.

```
    host {
      203.0.113.2;
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying Zones on page 1108
- Verifying Policies on page 1108

Verifying Zones

Purpose Verify that the security zone is enabled or not.

Action From operational mode, enter the **show security zones** command.

Verifying Policies

Purpose Verify that the policy is working.

Action From operational mode, enter the **show security policies** command on all the devices.

Related Documentation

- Security Policies Overview on page 1085
- Security Policies Configuration Overview on page 1093
- Example: Configuring a Security Policy to Permit or Deny All Traffic on page 1094

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 52

Configuring Negated Addresses

- Understanding Negated Address Support on page 1109
- Example: Configuring Negated Addresses on page 1110

Understanding Negated Address Support

Junos OS allows users to add any number of source and destination addresses to a policy. If you need to exclude certain addresses from a policy, you can configure them as negated addresses. When an address is configured as a negated address, it is excluded from a policy. You cannot, however, exclude the following IP addresses from a policy:

- Wildcard
- IPv6
- any
- any-ipv4
- any-ipv6
- 0.0.0.0

When a range of addresses or a single address is negated, it can be divided into multiple addresses. These negated addresses are shown as a prefix or a length that requires more memory for storage on a Packet Forwarding Engine.

Each platform has a limited number of policies with negated addresses. A policy can contain 10 source or destination addresses. The capacity of the policy depends on the maximum number of policies that the platform supports.

Copyright © 2016, Juniper Networks, Inc.

Before you configure a negated source address, destination address, or both, perform the following tasks:

1. Create a source, destination, or both address book.
2. Create address names and assign source and destination addresses to the address names.
3. Create address sets to group source, destination, or both address names.
4. Attach source and destination address books to security zones. For example, attach the source address book to the from-zone **trust** and the destination address book to the to-zone **untrust**.
5. Specify the match source, destination, or both address names.
6. Execute source-address-excluded, destination-address excluded, or both commands. A source, destination, or both addresses added in the source, destination, or both address books will be excluded from the policy.

NOTE: **The global address book does not need to be attached to any security zone.**

Related Documentation

- Example: Configuring Address Books and Address Sets on page 1076
- Example: Configuring Negated Addresses on page 1110

Example: Configuring Negated Addresses

This example shows how to configure negated source and destination addresses. It also shows how to configure address books and address sets.

- Requirements on page 1110
- Overview on page 1111
- Configuration on page 1111
- Verification on page 1113

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- A PC
- Junos OS Release 12.1X45-D10

Before you begin, configure address books and address sets. See "Example: Configuring Address Books and Address Sets" on page 1076.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you create source and destination address books, SOUR-ADDR and DES-ADDR, and add source and destination addresses to it. You create source and destination address sets, as1 and as2, and group source and destination addresses to them. Then you attach source address book to the security zone trust and the destination address book to the security zone untrust.

You create security zones from-zone trust and to-zone untrust. You specify the policy name to p1 and then you set the name of the match source address to as1 and the match destination address to as2. You specify the commands **source -address-excluded** and **destination -address-excluded** to exclude source and destination addresses configured in the policy p1. Finally, you set the policy p1 to permit traffic from-zone trust to to-zone untrust.

Configuration

CLI Quick Configuration

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

```
set security address-book SOU-ADDR address ad1 255.255.255.255/32
set security address-book SOU-ADDR address ad2 203.0.113.130/25
set security address-book SOU-ADDR address ad3 range-address 192.0.2.6 to 192.0.2.116
set security address-book SOU-ADDR address ad4 192.0.2.128/25
set security address-book SOU-ADDR address-set as1 address ad1
set security address-book SOU-ADDR address-set as1 address ad2
set security address-book SOU-ADDR address-set as1 address ad3
set security address-book SOU-ADDR address-set as1 address ad4
set security address-book SOU-ADDR attach zone trust
set security address-book DES-ADDR address ad8 198.51.100.1/24
set security address-book DES-ADDR address ad9 range-address 192.0.2.117 to 192.0.2.199
set security address-book DES-ADDR address ad10 198.51.100.0/24
set security address-book DES-ADDR address ad11 range-address 192.0.2.199 to
  192.0.2.250
set security address-book DES-ADDR address-set as2 address ad8
set security address-book DES-ADDR address-set as2 address ad9
set security address-book DES-ADDR address-set as2 address ad10
set security address-book DES-ADDR address-set as2 address ad11
set security address-book DES-ADDR attach zone untrust
set security policies from-zone trust to-zone untrust policy p1 match source-address as1
set security policies from-zone trust to-zone untrust policy p1 match
  source-address-excluded
set security policies from-zone trust to-zone untrust policy p1 match destination-address
  as2
set security policies from-zone trust to-zone untrust policy p1 match
  destination-address-excluded
set security policies from-zone trust to-zone untrust policy p1 match application any
set security policies from-zone trust to-zone untrust policy p1 then permit
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure negated addresses:

1. Create a source address book and address names. Add the source addresses to the address book.

   ```
   [edit security address book ]
   user@host#set SOU-ADDR address ad1 255.255.255.255/32
   user@host#set SOU-ADDR address ad2 203.0.113.130/25
   user@host#set SOU-ADDR ad3 range-address 192.0.2.6 to 192.0.2.116
   user@host#set SOU-ADDR address ad4 192.0.2.128/25
   ```

2. Create an address set to group source address names.

   ```
   [edit security address book ]
   user@host# set SOU-ADDR address-set as1 address ad1
   user@host# set SOU-ADDR address-set as1 address ad2
   user@host# set SOU-ADDR address-set as1 address ad3
   user@host# set SOU-ADDR address-set as1 address ad4
   ```

3. Attach the source address book to the security from zone.

   ```
   [edit security address book ]
   user@host# set SOU-ADDR attach zone trust
   ```

4. Create a destination address book and address names. Add the destination addresses to the address book.

   ```
   [edit security address book ]
   user@host#set DES-ADDR address ad8 198.51.100.1/24
   user@host#set DES-ADDR address ad9 range-address 192.0.2.117 to 192.0.2.199
   user@host#set DES-ADDR address ad10 198.51.100.0/24
   user@host#set DES-ADDR address ad11 range-address 192.0.2.199 to 192.0.2.250
   ```

5. Create another address set to group destination address names.

   ```
   [edit security address book ]
   user@host# set DES-ADDR address-set as1 address ad8
   user@host# set DES-ADDR address-set as1 address ad9
   user@host# set DES-ADDR address-set as1 address ad10
   user@host# set DES-ADDR address-set as1 address ad11
   ```

6. Attach the destination address book to the security to zone.

   ```
   [edit security address book ]
   user@host# set DES-ADDR attach zone untrust
   ```

7. Specify the policy name and source address.

   ```
   [edit security policies]
   user@host# set from-zone trust to-zone untrust policy p1 match source-address
     as1
   ```

8. Exclude source addresses from the policy.

   ```
   [edit security policies]
   user@host# set from-zone trust to-zone untrust policy p1 match
     source-address-excluded
   ```


Copyright © 2016, Juniper Networks, Inc.

9. Specify the destination address.

```
[edit security policies]
user@host# set from-zone trust to-zone untrust policy p1 match destination-address
  as2
```

10. Exclude destination addresses from the policy.

```
[edit security policies]
user@host#  set from-zone trust to-zone untrust policy p1 match
  destination-address-excluded
```

11. Configure the security policy application.

```
[edit security policies]
user@host# set from-zone trust to-zone untrust policy p1 match application any
```

12. Permit the traffic from-zone trust to to-zone untrust.

```
[edit security policies]
user@host# set from-zone trust to-zone untrust policy p1 then permit
```

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

```
[edit]
user@host# show security policies
  from-zone trust to-zone untrust {
    policy p1 {
      match {
        source-address as1;
        destination-address as2;
        source-address-excluded;
        destination-address-excluded;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Policy Configuration on page 1113
- Verifying the Policy Configuration Detail on page 1114

Verifying the Policy Configuration

Purpose Verify that the policy configuration is correct.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security policies policy-name p1** command.

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
From zone: trust, To zone: untrust
Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
Source addresses(excluded): as1
Destination addresses(excluded): as2
Applications: any
Action: permit
```

This output summarizes the policy configuration.

Verifying the Policy Configuration Detail

Purpose Verify that the policy and the negated source and destination address configurations are correct.

Action From operational mode, enter the **show security policies policy-name p1 detail** command.

```
user@host>show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(SOU-ADDR): 255.255.255.255/32
    ad2(SOU-ADDR): 203.0.113.130/25
    ad3(SOU-ADDR): 192.0.2.6 ~ 192.0.2.116
    ad4(SOU-ADDR): 192.0.2.128/25
    Destination addresses(excluded):
    ad8(DES-ADDR): 198.51.100.1/24
    ad9(DES-ADDR): 192.0.2.117 ~ 192.0.2.199
    ad10(DES-ADDR): 198.51.100.0/24
    ad11(DES-ADDR): 192.0.2.199 ~ 192.0.2.250
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

This output summarizes the policy configuration and shows the names of negated source and destination addresses excluded from the policy.

Related Documentation

- Understanding Negated Address Support on page 1109
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 53

Configuring Global Security Policy

- Global Policy Overview on page 1115
- Example: Configuring a Global Policy with No Zone Restrictions on page 1117
- Example: Configuring a Global Policy with Multiple Zones on page 1119

Global Policy Overview

In a Junos OS stateful firewall, security policies enforce rules for transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on traffic as it passes through the firewall. Security policies require traffic to enter one security zone and exit another security zone. This combination of a from-zone and to-zone is called a *context*. Each context contains an ordered list of policies. Each policy is processed in the order that it is defined within a context. Traffic is classified by matching the policy's from-zone, to-zone, source address, destination address, and the application that the traffic carries in its protocol header. Each global policy, as with any other security policy, has the following actions: permit, deny, reject, log, count.

You can configure a security policy from the user interface. Security policies control traffic flow from one zone to another zone by defining the kind(s) of traffic permitted from specific IP sources to specific IP destinations at scheduled times. This works well in most cases, but it is not flexible enough. For example, if you want to perform actions on traffic you have to configure policies for each possible context. To avoid creating multiple policies across every possible context, you can create a global policy that encompasses all zones, or a multizone policy that encompasses several zones.

Using a global policy, you can regulate traffic with addresses and applications, regardless of their security zones, by referencing user-defined addresses or the predefined address "any." These addresses can span multiple security zones. For example, if you want to provide access to or from multiple zones, you can create a global policy with the address "any," which encompasses all addresses in all zones. Selecting the "any" address matches any IP address, and when "any" is used as a source/destination address in any global policy configuration, it matches the source/destination address of any packet.

Using a global policy you can also provide access to multiple source zones and multiple destination zones in one policy. However, we recommend that, for security reasons and to avoid spoofing traffic, when you create a multizone policy you use identical matching criteria (source address, destination address, application) and an identical action. In Figure 69, for example, if you create a multizone policy that includes DMZ and Untrust

Copyright © 2016, Juniper Networks, Inc.

from-zones, spoofing traffic from 203.0.113.0/24 from the DMZ zone could match the policy successfully and reach the protected host in the Trust to-zone.

Figure 69: Multizone Global Policy Security Consideration

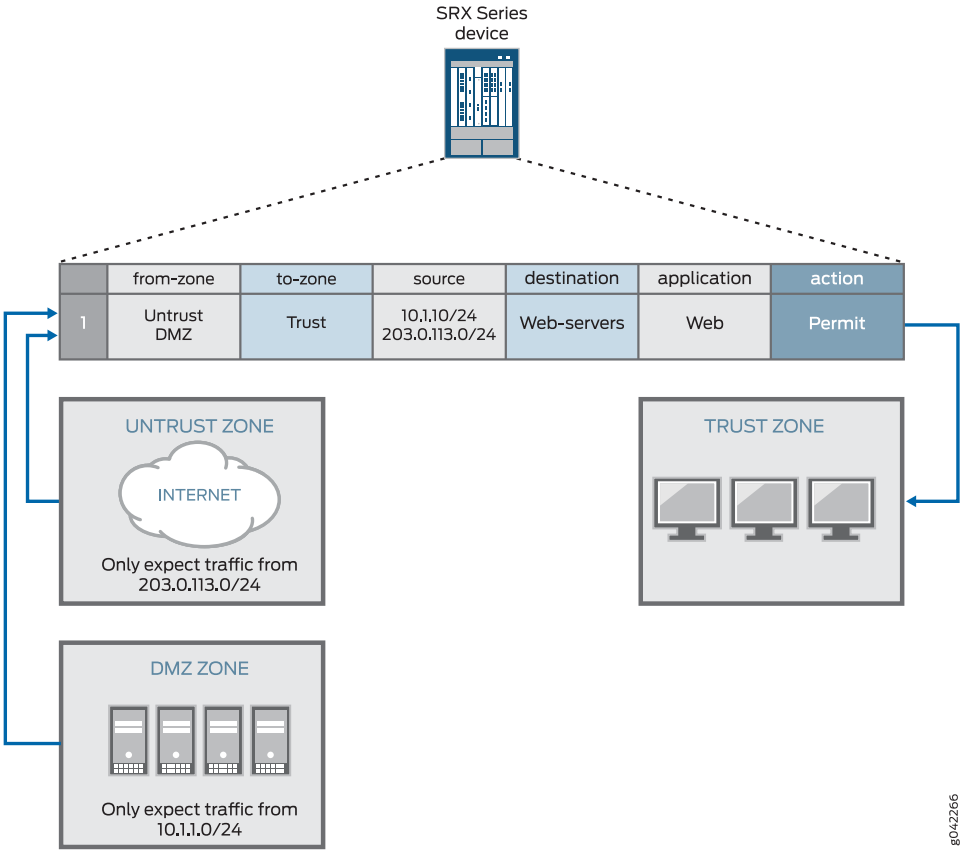

NOTE: Global policies without from-zone and to-zone information do not support VPN tunnels because VPN tunnels require specific zone information.

When policy lookup is performed, policies are checked in the following order: intra-zone (trust-to-trust), inter-zone (trust-to-untrust), then global. Similar to regular policies, global policies in a context are ordered, such that the first matched policy is applied to the traffic.

Copyright © 2016, Juniper Networks, Inc.

NOTE: If you have a global policy, make sure you have not defined a "catch-all" rule such as, match source any, match destination any, or match application any in the intra-zone or inter-zone policies because the global policies will not be checked. If you do not have a global policy, then it is recommended that you include a "deny all" action in your intra-zone or inter-zone policies. If you do have a global policy, then you should include a "deny all" action in the global policy.

In logical systems, you can define global policies for each logical system. Global policies in one logical system are in a separate context than other security policies, and have a lower priority than regular security policies in a policy lookup. For example, if a policy lookup is performed, regular security policies have priority over global policies. Therefore, in a policy lookup, regular security policies are searched first and if there is no match, global policy lookup is performed.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Example: Configuring a Global Policy with No Zone Restrictions on page 1117

Example: Configuring a Global Policy with No Zone Restrictions

Unlike other security policies in Junos OS, global policies do not reference specific source and destination zones. Global policies reference the predefined address "any" or user-defined addresses that can span multiple security zones. Global policies give you the flexibility of performing actions on traffic without any zone restrictions. For example, you can create a global policy so that every host in every zone can access the company website, for example, www.example.com. Using a global policy is a convenient shortcut when there are many security zones. Traffic is classified by matching its source address, destination address, and the application that the traffic carries in its protocol header.

This example shows how to configure a global policy to deny or permit traffic.

- Requirements on page 1117
- Overview on page 1118
- Configuration on page 1118
- Verification on page 1119

Requirements

Before you begin:

- Review the firewall security policies.

Copyright © 2016, Juniper Networks, Inc.

See "Security Policies Overview" on page 1085, "Global Policy Overview" on page 1115, "Understanding Security Policy Rules" on page 1087, and "Understanding Security Policy Elements" on page 1091.

- Configure an address book and create addresses for use in the policy.

  See "Example: Configuring Address Books and Address Sets" on page 1076.

- Create an application (or application set) that indicates that the policy applies to traffic of that type.

  See "Example: Configuring Applications and Application Sets" on page 1172.

Overview

This configuration example shows how to configure a global policy that accomplishes what multiple security policies (using zones) would have accomplished. Global policy gp1 permits all traffic while policy gp2 denies all traffic.

Configuration

CLI Quick Configuration

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

```
set security address-book global address server1 www.example.com
set security address-book global address server2 www.mail.example.com
set security policies global policy gp1 match source-address server1
set security policies global policy gp1 match destination-address server2
set security policies global policy gp1 match application any
set security policies global policy gp1 then permit
set security policies global policy gp2 match source-address server2
set security policies global policy gp2 match destination-address server1
set security policies global policy gp2 match application junos-ftp
set security policies global policy gp2 then deny
```

Step-by-Step Procedure

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

To configure a global policy to permit or deny all traffic:

1. Create addresses.

   ```
   [edit security]
   user@host# set security address-book global address server1 www.example.com
   user@host# set security address-book global address server2
     www.mail.example.com
   ```

2. Create the global policy to permit all traffic.

   ```
   [edit security]
   user@host# set policy global policy gp1 match source-address server1
   user@host# set policy global policy gp1 match destination-address server2
   user@host# set policy global policy gp1 match application any
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy global policy gp1 then permit
```

3. Create the global policy to deny all traffic.

```
[edit security]
user@host# set policy global policy gp2 match source-address server2
user@host# set policy global policy gp2 match destination-address server1
user@host# set policy global policy gp2 match application junos-ftp
user@host# set policy global policy gp2 then deny
```

Results From configuration mode, confirm your configuration by entering the **show security policies** and **show security policies** <*global*> commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host> show security policies
Default policy: permit-all
Global policies:
  Policy: gp1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: server1
    Destination addresses: server2
    Applications: any
    Action: permit
  Policy: gp2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 2
    Source addresses: server2
    Destination addresses: server1
    Applications: junos-ftp
    Action: deny
```

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

Verification

Verifying Global Policy Configuration

Purpose Verify that global policies gp1 and gp2 are configured as required.

Action From operational mode, enter the **show security policies** <*global*>command.

Related Documentation

- Security Policies Overview on page 1085
- Global Policy Overview on page 1115

Example: Configuring a Global Policy with Multiple Zones

Unlike other security policies in Junos OS, global policies allow you to create multizone policies. A global policy is a convenient shortcut when there are many security zones, because it enables you to configure multiple source zones and multiple destination zones in one global policy instead of having to create a separate policy for each

Copyright © 2016, Juniper Networks, Inc.

from-zone/to-zone pair, even when other attributes, such as source-address or destination-address, are identical.

- Requirements on page 1120
- Overview on page 1120
- Configuration on page 1120
- Verification on page 1121

Requirements

Before you begin:

- Review the firewall security policies.

  See "Security Policies Overview" on page 1085, "Global Policy Overview" on page 1115, "Understanding Security Policy Rules" on page 1087, and "Understanding Security Policy Elements" on page 1091.

- Create security zones.

  See "Example: Creating Security Zones" on page 1051

Overview

This configuration example shows how to configure a global policy that accomplishes what multiple security policies would have accomplished. Global policy Pa permits all traffic from zones 1 and 2 to zones 3 and 4.

Configuration

CLI Quick Configuration

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

```
set security policies global policy Pa match source-address any
set security policies global policy Pa match destination-address any
set security policies global policy Pa match application any
set security policies global policy Pa match from-zone zone1
set security policies global policy Pa match from-zone zone2
set security policies global policy Pa match to-zone zone3
set security policies global policy Pa match to-zone zone4
set security policies global policy Pa then permit
```

Step-by-Step Procedure

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

To configure a global policy with multiple zones:

1. Create a global policy to allow any traffic from zones 1 and 2 to zones 3 and 4.

   ```
   [edit security]
   set security policies global policy Pa match source-address any
   ```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies global policy Pa match destination-address any
set security policies global policy Pa match application any
set security policies global policy Pa match from-zone zone1
set security policies global policy Pa match from-zone zone2
set security policies global policy Pa match to-zone zone3
set security policies global policy Pa match to-zone zone4
set security policies global policy Pa then permit
```

Results From configuration mode, confirm your configuration by entering the **show security policies** *global* command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security policies global
policy Pa {
    match {
        source-address any;
        destination-address any;
        application any;
        from-zone [ zone1 zone2 ];
        to-zone [ zone3 zone4 ];
    }
    then {
        permit;
    }
}
```

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

Verification

Verifying Global Policy Configuration

Purpose Verify that the global policy is configured as required.

Action From operational mode, enter the **show security policies** *global* command.

Related Documentation

- Security Policies Overview on page 1085
- Global Policy Overview on page 1115

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 54

Managing Security Policy Activation By Configuring Schedulers

- Security Policy Schedulers Overview on page 1123
- Example: Configuring Schedulers for a Daily Schedule Excluding One Day on page 1124

Security Policy Schedulers Overview

Schedulers are powerful features that allow a policy to be activated for a specified duration. You can define schedulers for a single (nonrecurrent) or recurrent time slot within which a policy is active. You can create schedulers irrespective of a policy, meaning that a scheduler cannot be used by any policies. However, if you want a policy to be active within a scheduled time, then you must first create a scheduler.

When a scheduler times out, the associated policy is deactivated and all sessions associated with the policy are also timed out.

If a policy contains a reference to a scheduler, the schedule determines when the policy is active, that is, when it can be used as a possible match for traffic. Schedulers allow you to restrict access to a resource for a period of time or remove a restriction.

The following guidelines apply to schedulers:

- A scheduler can have multiple policies associated with it; however, a policy cannot be associated with multiple schedulers.
- A policy is active during the time when the scheduler it refers to is also active.
- When a scheduler is off, the policy is unavailable for policy lookup.
- A scheduler can be configured as one of the following:
  - Scheduler can be active for a single time slot, as specified by a start date and time and a stop date and time.
  - Scheduler can be active forever (recurrent), but as specified by the daily schedule. The schedule on a specific day (time slot) takes priority over the daily schedule.
  - Scheduler can be active within a time slot as specified by the weekday schedule.
  - Scheduler can have a combination of two time slots (daily and timeslot).

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Security Policies Overview on page 1085
- Example: Configuring Schedulers for a Daily Schedule Excluding One Day on page 1124
- Verifying Scheduled Policies on page 1159

Example: Configuring Schedulers for a Daily Schedule Excluding One Day

This example shows how to configure schedulers for packet match checks every day except Sunday.

- Requirements on page 1124
- Overview on page 1124
- Configuration on page 1124
- Verification on page 1126

Requirements

Before you begin:

- Understand security policies schedulers. See "Security Policies Overview" on page 1085.
- Configure security zones before applying this configuration. See "Example: Creating Security Zones" on page 1051.

Overview

Schedulers are powerful features that allow a policy to be activated for a specified duration. You can define schedulers for a single (nonrecurrent) or recurrent time slot within which a policy is active. If you want a policy to be active within a scheduled time, then you must first create a scheduler.

To configure a scheduler, you enter a meaningful name and a start and stop time for the scheduler. You can also attach comments.

In this example, you:

- Specify the scheduler, sch1, that allows a policy, which refers to it, to be used for packet match checks every day except Sunday.
- Create a policy, abc, and specify the match conditions and action to be taken on traffic that matches the specified conditions, and bind the schedulers to the policy to allow access during the specified days.

Configuration

CLI Quick Configuration

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

```
set schedulers scheduler sch1 daily all-day
```


Copyright © 2016, Juniper Networks, Inc.

```
set schedulers scheduler sch1 sunday exclude
set security policies from-zone green to-zone red policy abc match source-address any
set security policies from-zone green to-zone red policy abc match destination-address
  any
set security policies from-zone green to-zone red policy abc match application any
set security policies from-zone green to-zone red policy abc then permit
set security policies from-zone green to-zone red policy abc scheduler-name sch1
set security policies default-policy permit-all
```

Step-by-Step Procedure

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

To configure a scheduler:

1. Set a scheduler.

```
[edit schedulers ]
user@host# set scheduler sch1 daily all-day
user@host# set scheduler sch1 sunday exclude
```

2. Specify the match conditions for the policy.

```
[edit security policies from-zone green to-zone red policy abc]
user@host# set match source-address any destination-address any application
  any
```

3. Specify the action.

```
[edit security policies from-zone green to-zone red policy abc]
user@host# set then permit
```

4. Associate the scheduler to the policy.

```
[edit security policies from-zone green to-zone red policy abc ]
user@host# set scheduler-name sch1
```

Results

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

```
[edit]
[user@host]show schedulers
  scheduler sch1 {
    daily all-day;
    sunday exclude;
  }
[edit]
[user@host]show security policies
from-zone green to-zone red {
  policy abc {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
```

Copyright © 2016, Juniper Networks, Inc.

```
        permit;
      }
      scheduler-name sch1;
    }
  }
    default-policy {
      permit-all;
    }
```

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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Schedulers are Active on page 1126
- Verifying Policies on page 1126

Verifying Schedulers are Active

Purpose Verify if schedulers are enabled or not.

Action From operational mode, enter the **show schedulers** command.

Verifying Policies

Purpose Verify if the policies are working.

Action From operational mode, enter the **show security policies** command.

Related Documentation

- Verifying Scheduled Policies on page 1159

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 55

Configuring User Role Firewall Security Policies

- Understanding User Role Firewalls on page 1127
- User Role Retrieval and the Policy Lookup Process on page 1128
- Understanding the User Identification Table on page 1130
- Obtaining Username and Role Information Through Firewall Authentication on page 1136
- Configuring a User Role Firewall For Captive Portal Redirection on page 1137
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138
- Configuring Resource Policies Using UAC on page 1145

Understanding User Role Firewalls

Network security enforcement, monitoring, and reporting based solely on IP information soon will not be sufficient for today's dynamic and mobile workforce. By integrating user firewall policies, administrators can permit or restrict network access of employees, contractors, partners, and other users based on the roles they are assigned. User role firewalls enable greater threat mitigation, provide more informative forensic resources, improve record archiving for regulatory compliance, and enhance routine access provisioning.

User role firewalls trigger two actions:

- Retrieve user and role information associated with the traffic
- Determine the action to take based on six match criteria within the context of the zone pair

The source-identity field distinguishes a user role firewall from other types of firewalls. If the source identity is specified in any policy for a particular zone pair, it is a user role firewall. The user and role information must be retrieved before policy lookup occurs. If the source identity is not specified in any policy, user and role lookup is not required.

To retrieve user and role information, authentication tables are searched for an entry with an IP address corresponding to the traffic. If an entry is found, the user is classified as an authenticated user. If not found, the user is classified as an unauthenticated user.

Copyright © 2016, Juniper Networks, Inc.

The username and roles associated with an authenticated user are retrieved for policy matching. Both the authentication classification and the retrieved user and role information are used to match the source-identity field.

Characteristics of the traffic are matched to the policy specifications. Within the zone context, the first policy that matches the user or role and the five standard match criteria determines the action to be applied to the traffic.

The following sections describe the interaction of user and role retrieval and the policy lookup process, methods for acquiring user and role assignments, techniques for configuring user role firewall policies, and an example of configuring user role firewall policies.

Related Documentation

- User Role Retrieval and the Policy Lookup Process on page 1128
- Understanding the User Identification Table on page 1130
- Configuring a User Role Firewall For Captive Portal Redirection on page 1137
- Configuring Resource Policies Using UAC on page 1145
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138

User Role Retrieval and the Policy Lookup Process

For policy lookup, firewall policies are grouped by zone pair (the from zone and to zone). Within the context of the zone pair, IP-based firewall policies are matched to traffic based on five criteria—source IP, source port, destination IP, destination port, and protocol.

User role firewall policies include a sixth match criteria—source identity. The source-identity field specifies the users and roles to which the policy applies. When the source-identity field is specified in any policy within the zone pair, user and role information must be retrieved before policy lookup can proceed. (If all policies in the zone pair are set to **any** or have no entry in the source-identity field, user and role information is not required and the five standard match criteria are used for policy lookup.)

The user identification table (UIT) provides user and role information for an active user who has already been authenticated. Each entry in the table maps an IP address to an authenticated user and any roles associated with that user.

When traffic requires user and role data, each registered UIT is searched for an entry with the same IP address. If a user has not been authenticated, there is no entry for that IP address in the table. If no UIT entry exists, the user is considered an unauthenticated user.

Policy lookup resumes after the user and role information has been retrieved. The characteristics of the traffic are matched against the match criteria in the policies. The source-identity field of a policy can specify one or more users or roles, and the following keywords:

authenticated-user—Users that have been authenticated.

unauthenticated-user—Users that have not been authenticated.

Copyright © 2016, Juniper Networks, Inc.

any—All users regardless of authentication. If the source-identity field is not configured or is set to any in all of the policies for the zone pair, only five criteria are matched.

unknown-user—Users unable to be authenticated due to an authentication server disconnection, such as a power outage.

For example, consider user-c who is assigned to the mgmt role. When traffic from the trust zone to the untrust zone is received from user-c at IP address 198.51.100.3, policy lookup is initiated. Table 74 represents three policies in a user role firewall for the trust to untrust zone pair.

Table 74: Trust Zone to Untrust Zone Policy Sequence

src-zone	src-zone	dest-zone	src-IP	dest-IP	source-identity	Application	Action	Services
P1	trust	untrust	192.0.2.0	203.0.113.0	any	http	deny	–
P2	trust	untrust	any	any	mgmt	any	permit	–
P3	trust	untrust	198.51.100.3	any	employee	http	deny	–

All policies for the zone pair are checked first for a source-identity option. If any of the policies specifies a user, a role, or a keyword, user and role retrieval must occur before policy lookup continues. Table 74 shows that policy P2 specifies mgmt as the source identity, making this a user role firewall. User and roles must be retrieved before policy lookup can continue.

NOTE: User and role retrieval would not be performed if the keyword any or if no source identity was specified in all of the policies in the zone context. In such cases, only the five remaining values are matched to the policy criteria.

The UIT represented in Table 75 is checked for the IP address. Because the address is found, the username user-c, all roles listed for user-c (in this case, mgmt and employee), and the keyword authenticated-user become data used to match the traffic to the **source-identity** field of a policy.

Table 75: UIT Authentication Details

Source IP Address	Username	Roles
192.0.2.4	user-a	employee
198.51.100.3	user-c	mgmt, employee
203.0.113.2	user-s	contractor

Policy lookup resumes and compares the match criteria in each policy in Table 74 to the incoming traffic. Assuming all other criteria match, the first policy that specifies user-c, mgmt, employee, authenticated-user, or any in the source-identity field could be a match

Copyright © 2016, Juniper Networks, Inc.

for this traffic. Policy P1 matches one of the retrieved roles for user-c, but the source IP address does not match; therefore policy lookup continues. For policy P2, all criteria match the traffic; therefore the policy action is followed and the traffic is permitted. Note that the traffic also matches policy P3, but user firewall policies are terminal—policy lookup ends when the first policy match is found. Because policy P2 matches all criteria, policy lookup ends and policy P3 is not checked.

Policies can also be based on the classification assigned to a user from the user and role retrieval results. Consider a different set of policies for the same zone pair represented by Table 76. If traffic is received from user-q at IP 198.51.100.5, user and role retrieval is required because the source-identity field is specified in at least one of the policies.

Table 76: Trust Zone to Untrust Zone Policy Sequence

policy-name	src-zone	dest-zone	src-IP	dest-IP	source-identity	application	action	Services
P1	trust	untrust	any	any	unauthenticated user	http	deny	–
P2	trust	untrust	any	any	mgmt	any	permit	–
P3	trust	untrust	198.51.100.3	any	employee	http	deny	–

When the UIT entries in Table 75 are checked, no entry is found for IP address 198.51.100.5. Therefore, the user is considered an unauthenticated user. When policy lookup resumes, the traffic matches policy P1 and the traffic is denied.

Related Documentation

- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130
- Configuring a User Role Firewall For Captive Portal Redirection on page 1137
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138
- Configuring Resource Policies Using UAC on page 1145

Understanding the User Identification Table

On the SRX Series device, the user identification table (UIT) contains the IP address, username, and role information for each authenticated user. Entries are ordered by IP address. When username and role information is required by a security policy, all UITs are checked. Finding the IP address in an entry in one of the UITs means that the user at that address has already been successfully authenticated.

Each authentication source maintains its own UIT independently and provides query functions for accessing data. Three types of UITs are supported—the local authentication table, the Unified Access Control (UAC) authentication table, and the firewall authentication table.

Local authentication table—A static UIT created on the SRX Series device either manually or programmatically using CLI commands. All users included in the local authentication table are considered authenticated users. When a matching IP address

Copyright © 2016, Juniper Networks, Inc.

is found, user and role information is retrieved from the table entry and associated with the traffic. User and role information can be created on the device manually or ported from a third-party authentication server, but the data in the local authentication table is not updated in real time.

UAC authentication table—A dynamic UIT pushed from the Junos Pulse Access Control Service to the SRX Series device. The UAC authentication table of a Junos Pulse Access Control Service contains an entry for each authenticated user. The data in this table is updated and pushed to the SRX Series device whenever its authentication table is updated. Depending on the device configuration, authentication could occur on the Junos Pulse Access Control Service itself or on a third-party authentication server. If the Access Control Service is relaying data from a third-party server, the data is restructured by the Access Control Service to match the file format of its authentication table and pushed to the SRX Series device.

Firewall authentication table—A dynamic UIT created on the SRX when **user-firewall** is specified as the firewall authentication type in a security policy. This UIT provides an alternative user role source to UAC when firewall authentication is already in use on your SRX Series device. In this way, users defined for pass-through authentication can also be used as a source for usernames and roles when the **user-firewall** option is specified as the firewall authentication type in a policy.

The **user-firewall** authentication type initiates firewall authentication to verify the user by using either local authentication information or external authentication servers supporting RADIUS, LDAP, or SecureID authentication methods. When this type is specified for firewall authentication, the username and associated groups (roles) from the authentication source are mapped to the IP address and added to the firewall authentication UIT.

- Local Authentication Table on page 1131
- UAC Authentication Table on page 1133
- Firewall Authentication Table on page 1134
- Policy Provisioning With Users and Roles on page 1135

Local Authentication Table

The local authentication table is managed with CLI commands that insert or delete entries. A local authentication table can be used as a backup solution when a dynamic UIT is not available, or to assign user and role information to devices that cannot authenticate to the network, such as printers or file servers. The local authentication table can be used for testing or to demonstrate how a user role firewall works without firewall authentication or the Access Control Service configured.

The IP addresses, user names, and roles from a third-party authentication source can be downloaded and added to the local authentication table programmatically using CLI commands. If an authentication source defines users and groups, the groups can be configured as roles and associated with the user as usual.

To be compliant with the UAC authentication table, user names are limited to 65 characters and role names are limited to 64 characters. The local authentication table has a maximum of 10,000 authentication entries on SRX1500 devices and above, 5000

Copyright © 2016, Juniper Networks, Inc.

authentication entries on SRX550 devices and below, and 5120 authentication entries on the vSRX. Each authentication entry can be associated with up to 200 roles. The maximum capacity is based on an average of 10 roles assigned to each user. This is the same capacity specified for a UAC authentication table.

Use the following command to add an entry to a local authentication table. Note that each entry is keyed by IP address.

user@host> **request security user-identification local-authentication-table add user** ***user-name*** **ip-address** ***ip-address*** **role [***role-name role-name* **]**

The role option in a single CLI command accepts up to 40 roles. To associate more than 40 roles with a single user, you need to enter multiple commands. Keep the following characteristics in mind when adding or modifying authentication user and role entries.

- Role names cannot be the same as usernames.
- Using the **add** option with an existing IP address and username aggregates the role entries. The table can support up to 200 roles per user.
- Using the **add** option with an existing IP address and a new username overwrites the existing username for that IP address.
- Role aggregation does not affect existing sessions.
- To change the role list of an existing entry, you need to delete the existing entry and add an entry with the new role list.
- To change the IP address of an existing entry, you need to delete the existing entry and add an entry with the new IP address.

An entry can be deleted by IP address or by username.

user@host> **request security user-identification local-authentication-table delete (***ip-address* **|** ***user-name*****)**

The local authentication table can be cleared with the following command:

user@host> **clear security user-identification local-authentication-table**

To display the content of the local authentication table, use the following **show...** command:

user@host> **show security user-identification local-authentication-table all (brief | extensive)**

The **brief** option (the default) displays information in a tabular format sequenced by IP address. User names and role lists are truncated to fit the format.

user@host> **show security user-identification local-authentication-table all**

```
Total entries: 2
Source IP       Username            Roles
198.51.100.1       user1            role1
203.0.113.2        user2            role2, role3
```

The **extensive** option displays the full content for each field. Other options limit the display to a single username, IP address, or role.

Copyright © 2016, Juniper Networks, Inc.

user@host> **show security user-identification local-authentication-table all extensive**

```
Total entries: 3
Ip-address: 198.51.100.2
Username: user1
Roles: role1

Ip-address: 203.0.113.2
Username: user1
Roles: role2

Ip-address: 192.0.2.3
Username: user3
Roles: role1, role2
```

UAC Authentication Table

An SRX Series device can act as an enforcer for a Junos Pulse Access Control Service. In this implementation, the SRX Series device acts as a Layer 3 enforcement point and controls access to resources with IP-based resource policies that have been pushed down from the Access Control Service.

When implemented as a user role firewall, the SRX Series device can access the UAC network in a similar way for user role retrieval. In this instance, user and role information for all authenticated users is pushed from the Access Control Service.

The SRX Series device configuration is similar to that of an enforcer. To establish communication, both devices require configuration and password settings to recognize the other. From the SRX Series device, connect the Access Control Service as an infranet controller.

[edit]
user@host# **set services unified-access-control infranet-controller** *ic-name* **address** *ip-address*
user@host# **set services unified-access-control infranet-controller** *ic-name* **interface** *interface-name*
user@host# **set services unified-access-control infranet-controller** *ic-name* **password** *password*

From the Access Control Service, define the SRX Series device as a New Enforcer. Use the same password specified on the SRX Series device.

Users and passwords are defined on the Access Control Service as in a standard authentication configuration. One or more roles can also be associated with users. When a user is authenticated, an entry containing the IP address, username, and associated roles is added to the UAC authentication table on the Access Control Service.

The UAC authentication table is pushed from the Access Control Service to the SRX Series device when the connection between the two devices is initialized. Whenever an entry is added, removed, or updated on the Access Control Service, the updated UAC authentication table is pushed to the SRX Series device.

Resource access policies are not necessary on the Access Control Service for a user role firewall implementation. The access behavior is provided in the policy configurations on the SRX Series device. If resource access policies are defined on the Access Control

Copyright © 2016, Juniper Networks, Inc.

Service, they are pushed to the SRX Series device, but they are not used unless a specific firewall policy implements UAC policies in the policy's action field.

NOTE: On SRX1500, SRX5400, SRX5600, and SRX5800 devices, the default authentication table capacity is 45,000; the administrator can increase the capacity to a maximum of 50,000.

On SRX300, SRX320, SRX340, SRX345, and SRX550 devices, the default authentication table capacity is 10,000; the administrator can increase the capacity to a maximum of 15,000.

The following **show services** command displays the content of the UAC authentication table on the SRX Series device, confirming that the table has been pushed from the Access Control Service successfully:

```
user@host> show services unified-access-control authentication-table extended

Id    Source IP       Username     Age   Role name
3     192.0.2.1          april        60    Users
6     192.0.2.2          june         60    Employeees
Total: 2
```

The SRX Series device monitors connections and detects if communication to the Access Control Service has been lost. Based on the UAC configuration, the SRX Series device waits for a response for a configured interval before issuing another request. If a response is received, the Access Control Service is considered functional. If no response is received after a specified timeout period, communication is considered lost and the timeout action is applied. The following UAC command syntax configures the interval, timeout, and timeout action:

```
user@host# set services unified-access-control interval seconds
user@host# set services unified-access-control timeout seconds
user@host# set services unified-access-control timeout-action (close | no-change | open)
```

During a disconnection, if user and role lookup is attempted for the disconnected device, it returns a failure code regardless of the timeout action. If access to all authentication sources is lost, the keyword unknown-user is associated with the IP address. When policy lookup resumes, a policy with unknown-user as the source identity would match the traffic. By implementing a specific policy for unknown-user, you can create a method for handling the loss of authentication sources.

Firewall Authentication Table

Firewall authentication requires users to authenticate to the SRX firewall before permitting access between zones and devices. When traffic is received, the user is prompted for a username and password, and verified against a specified profile of valid users. Depending on the device configuration, firewall authentication verifies that telnet, HTTP, HTTPS (for high-end SRX devices), and FTP traffic has been authenticated locally or by a RADIUS, LDAP, or SecureID authentication server.

If firewall authentication is in use on a device, the authentication process can also provide the username and role information needed for user role firewall match criteria. In this

Copyright © 2016, Juniper Networks, Inc.

case, the information is collected and maintained in a UIT called the firewall authentication table. One or more access policies in the **edit access** hierarchy define authentication methods to be used for firewall authentication.

The firewall authentication table must be enabled as the authentication source for user role information retrieval. The **priority** option specifies the sequence in which all UITs will be checked.

```
user@host# set security user-identification authentication-source firewall-authentication
  priority priority
```

In a firewall policy for a given zone pair, the **firewall-authentication** service specified for the **permit** action initiates authentication of matching traffic. The **user-firewall** authentication type generates the UIT entry for the authenticated user. The name specified in the **access-profile** option identifies the profile to be used to authenticate valid users.

```
[edit security policies from-zone zone to-zone zone policy policy-name]
user@host# set match source-identity unauthenticated-user
user@host# set then permit firewall-authentication user-firewall access-profile
  profile-name
```

The UIT table entry contains the IP address of the traffic mapped to the authenticated user and the user's associated groups. When the user is no longer active, the entry is removed from the table. Because entries are continuously added and removed as the traffic and authenticated users change, the firewall authentication table is considered dynamic.

When policies within the same zone pair specify the **source-identity** field as part of its match criteria, all enabled UITs are searched for an entry corresponding to the IP address of the traffic. If found, the associated username and groups are retrieved for source-identity matching. (User authentication group names are considered role names for source-identity matching.)

Policy Provisioning With Users and Roles

All users and roles, whether defined on the SRX Series device or on the Access Control Service, are maintained in a user role file on the SRX Series device. To display all users and roles available for provisioning, use the following **show security...** commands.

NOTE: Usernames and roles in the firewall authentication table are not included in the following displays.

- To display all of the roles that are available for provisioning, use the **show security user-identification role-provision all** command. Note that the roles from all UITs are listed together.
- To display all of the users that are available for provisioning, use the **show security user-identification user-provision all** command.
- To display all of the users and roles that are available for provisioning, use the **show security user-identification source-identity-provision all** command.

Copyright © 2016, Juniper Networks, Inc.

When a policy configuration is committed, the user role file is checked to determine if all users and roles specified in the policy are available for provisioning. If a user or role is not found, a warning identifies the missing user or role so that you can define it later.

NOTE: **The policy is committed even if a user or role is not yet defined.**

Related Documentation

- Understanding User Role Firewalls on page 1127
- User Role Retrieval and the Policy Lookup Process on page 1128
- Configuring a User Role Firewall For Captive Portal Redirection on page 1137
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138
- Configuring Resource Policies Using UAC on page 1145
- Acquiring User Role Information from an Active Directory Authentication Server on page 5269

Obtaining Username and Role Information Through Firewall Authentication

User role firewall policies can be integrated with firewall authentication both to authenticate users and to retrieve username and role information. The information is mapped to the IP address of the traffic, stored in the firewall authentication table, and used for user role firewall policy enforcement.

The following CLI statements configure firewall authentication for user role firewall enforcement.

1. If not already established, define the access profile to be used for firewall authentication. You can skip this step if an existing access profile provides the client data needed for your implementation.

   The access profile is configured in the **[edit access profile]** hierarchy as with other firewall authentication types. It defines clients as firewall users and the passwords that provide them access. Use the following command to define a profile and add client names and passwords for firewall authentication.

   set access profile *profile-name* client *client-name* firewall-user password *pwd*

2. If HTTPS traffic is expected, define the access profile to be used for SSL termination services. You can skip this step if an existing SSL termination profile provides the services needed for your implementation.

   The SSL termination profile is configured in the **[edit services ssl]** hierarchy.

   set services ssl termination profile *ssl-profile-name* server-certificate *certificate-type*

3. Enable the firewall authentication table as an authentication source.

   set security user-identification authentication-source firewall-authentication priority *priority*

Copyright © 2016, Juniper Networks, Inc.

The priority value determines the sequence in which authentication sources are checked. The default value is 150 for the firewall authentication table. (It is 100 for the local authentication table and 200 for the Unified Access Control (UAC) authentication table.) By default, the local authentication table is checked first, the firewall authentication table is next, and the UAC authentication table is third if it is enabled. You can change this sequence by changing the priority value of one or more of the tables.

4. Configure policies that permit traffic for user firewall authentication.

```
edit security policies from-zone zone to-zone zone policy policy-name
set match source-identity unauthenticated-user
set then permit firewall-authentication user-firewall access-profile profile-name
  ssl-termination-profile profile-name
```

When unauthenticated traffic is permitted for firewall authentication, the user is authenticated based on the access profile configured in this statement. The **ssl-termination-profile** option is needed only for HTTPS traffic.

By specifying the authentication type **user-firewall**, the firewall authentication table is propagated with the IP address, the username, and any group names associated with the authenticated user. (Group names from firewall authentication are interpreted as roles by the user role firewall.) Any further traffic from this IP address will match the IP address in the firewall authentication table, and not require authentication. The associated username and roles are retrieved from the table for use as potential match criteria in subsequent security policies.

Related Documentation

- Understanding the User Identification Table on page 1130

Configuring a User Role Firewall For Captive Portal Redirection

To automatically redirect unauthenticated users to the Access Control Service, use the UAC captive portal feature. The following syntax defines the profile for the captive portal:

```
set services unified-access-control captive-portal profile-name redirect-traffic
  [unauthenticated | all]
set services unified-access-control captive-portal profile-name redirect-url host-url
```

The Kerberos protocol, used for authentication encryption, identifies the Access Control Service only by its service principal name (SPN). The protocol does not accept an IP address. Therefore, the format for the redirect URL must be

service://hostname/options

In this implementation, the service is HTTP and the hostname is the FQDN of the Access Control Service. Options specified after the hostname pass additional information to the Access Control Service directing the user back to the original destination, to the SRX Series device, or to the policy that originated the redirection. You can configure the options using the following keyword and variable pairs:

?target=%dest-url%—Specifies the protected resource which the user is trying to access.

Copyright © 2016, Juniper Networks, Inc.

&enforcer=%enforcer-id%—Specifies the ID assigned to the SRX Series device when it is configured as an enforcer by the Access Control Service.

&policy=%policy-id%—Specifies the encrypted policy ID for the security policy that redirected the traffic.

The following statements define the profile of the captive portal named auth-redirect. The captive portal redirects unauthenticated users to the URL of the Access Control Service for authentication. After successful authentication, the traffic will be directed back to the SRX Series device.

```
[edit]
user@host# set services unified-access-control captive-portal auth-redirect redirect-traffic
  unauthenticated
user@host# set services unified-access-control captive-portal auth-redirect redirect-url
  “http://ic6000.example.com/?target=%dest-url%&policy=%policy-id%”
```

A defined captive-portal profile is displayed as part of the UAC configuration.

```
...
services unified-access-control captive-portal auth-redirect {
     redirect-traffic unauthenticated;
     redirect-url
“http://ic6000.example.com/?target=%dest-url%&policy=%policy-id%”;
}
```

After the profile is defined, a policy can apply the captive portal as an application service when certain criteria are matched. The following example defines policy P1 to apply the auth-redirect captive portal profile to any HTTP traffic from the trust to untrust zones whenever the role is unauthenticated-user:

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy P1 match
  application http
user@host# set security policies from-zone trust to-zone untrust policy P1 match
  source-identity unauthenticated-user
user@host# set security policies from-zone trust to-zone untrust policy P1 then permit
  application-services uac-policy captive-portal auth-redirect
```

Related Documentation

- Understanding User Role Firewalls on page 1127
- User Role Retrieval and the Policy Lookup Process on page 1128
- Understanding the User Identification Table on page 1130
- Configuring Resource Policies Using UAC on page 1145
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138

Example: Configuring a User Role Firewall on an SRX Series Device

The following example configures a user role firewall on an SRX Series device. The firewall controls access from the trust zone to the untrust zone based on active, authenticated

Copyright © 2016, Juniper Networks, Inc.

users or their associated roles. User role firewall policies establish the following restrictions:

- Only authenticated users are permitted from the trust zone to the untrust zone.

  Unauthenticated users are redirected to an Access Control Service for authentication.

- Traffic from IP 192.0.2.0 to IP 203.0.113.0 within the zone context is restricted. Only the traffic from users with the dev-abc, http-juniper-accessible, or ftp-accessible role is permitted. Permitted traffic is further evaluated by AppFW rules.
  - Permitted traffic identified as junos:FACEBOOK-ACCESS, junos:GOOGLE-TALK, or junos:MEEBO application traffic is denied.
  - Permitted traffic for any other application is permitted.
- All other traffic from the trust zone to the untrust zone is permitted.

- Requirements on page 1139
- Overview on page 1139
- Configuration on page 1140

Requirements

Before you begin, ensure that the SRX Series device with Junos OS Release 12.1 or later is configured and initialized.

In this example, user and role information associated with the IP address of the traffic is provided by an Access Control Service. For instructions on configuring the Access Control Server, see "Acquiring User Role Information from an Active Directory Authentication Server" on page 5269.

Overview

Table 77 outlines a firewall that meets the requirements for this example. The user role firewall consists of four policies.

Table 77: User Role Firewall Policies

policy-name	src-zone	dest-zone	src-IP	dest-IP	source-identity	application	action	Services
user-role-fw1	trust	untrust	any	any	unauthenticated-user	http	permit	UAC captive portal
user-role-fw2	trust	untrust	192.0.2.0	203.0.113.0	dev-abc http-juniper-accessible ftp-accessible	http	permit	AppFW ruleset RS1
user-role-fw3	trust	untrust	192.0.2.0	203.0.113.0	any	http	deny	
user-role-fw4	trust	untrust	any	any	any	http	permit	

Copyright © 2016, Juniper Networks, Inc.

Because the **source-identity** field is specified for at least one of the policies in this firewall, user and role information must be retrieved before policy lookup is conducted. The source IP of the traffic is compared to the items in the UIT. If the source IP address is found, the keyword **authenticated**, the username, and any roles associated with this user are stored for later use in policy lookup. If a matching entry for the IP address is not found in the UIT, the keyword **unauthenticated-user** is stored for policy lookup.

After retrieving the username, roles, and keywords, policy lookup begins. Characteristics of the incoming traffic are compared to each policy's match criteria. If a match is found, the action specified in that policy is taken.

A policy match is a terminal event, and no policies after the match are checked. Policy sequence influences the action to be taken for matching traffic. In this example, policies are applied in the following sequence:

user-role-fw1—Applies the UAC captive portal service to matching HTTP traffic with the unauthenticated-user keyword, and redirects it to the Access Control Service for authentication. A UAC profile must also be configured to identify the captive portal specifications.

user-role-fw2—Applies an AppFW rule set to any HTTP traffic from address 192.0.2.0 to address 203.0.113.0 that has a matching username or role. An application firewall must also be configured to define the rule set.

user-role-fw3—Denies all remaining HTTP traffic from address 192.0.2.0 to address 203.0.113.0 for this zone pair.

user-role-fw4—Permits all remaining HTTP traffic for this zone pair.

Configuration

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

Configuring Redirection For Unauthenticated Users

Step-by-Step Procedure

When an IP address is not listed in the UIT, the unauthenticated-user keyword is used in policy lookup. Instead of denying access to this traffic, a policy can redirect the traffic to a UAC captive portal for authentication.

NOTE: It is important to position a redirection policy for unauthenticated-user before a policy for "any" user so that UAC authentication is not shadowed by a policy intended for authenticated users.

To configure redirection from the SRX Series device to the Access Control Service:

1. From configuration mode, configure the UAC profile for the captive portal acs-device.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set services unified-access-control captive-portal acs-device
  redirect-traffic unauthenticated-user
```

2. Configure the redirection URL for the Access Control Service or a default URL for the captive portal.

```
[edit]
user@host# set services unified-access-control captive-portal acs-device
redirect-url “https://%ic-url%/?target=%dest-url%&enforcer=%enforcer-id%”
```

This policy specifies the default target and enforcer variables to be used by the Access Control Service to direct the user back after authentication. This ensures that changes to system specifications will not affect configuration results.

NOTE: When variables, such as ?target=, are included in the command line, you must enclose the URL and variables in quotation marks.

3. Configure a user role firewall policy that redirects HTTP traffic from zone trust to zone untrust if the source-identity is unauthenticated-user. The captive portal profile name is specified as the action to be taken for traffic matching this policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw1
  match source-address any
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw1
  match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw1
  match application http
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw1
  match source-identity unauthenticated-user
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw1
  then permit application-services uac-policy captive-portal acs-device
```

4. If you are done configuring the policies, commit the changes.

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

Results From configuration mode, confirm your configuration by entering the **show services** and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show services
...
unified-access-control {
 captive-portal acs-device {
  redirect-traffic unauthenticated;
  redirect-url “https://%ic-ip%/?target=%dest-url%&enforcer=%enforcer-id%”
...
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security policies
...
from-zone trust to-zone untrust {
 policy user-role-fw1 {
  match {
   source-address any;
   destination-address any;
   application http;
   source-identity unauthenticated-user
  }
  then {
   permit {
    application-services {
     uac-policy {
      captive-portal acs-device;
     }
    }
   }
  }
 }
}
...
```

Creating a User Role Policy With an Application Firewall

Step-by-Step Procedure

This policy restricts traffic from IP 192.0.2.0 to IP 203.0.113.0 based on its user and roles, and also its application. The configuration defines an application rule set and applies it to matching user role traffic.

1. Configure the AppFW rule set rs1. The following rule set denies junos:FACEBOOK-ACCESS, junos:GOOGLE-TALK, or junos:MEEBO application traffic. It applies the default setting, permit, to the remaining traffic.

   ```
   [edit]
   user@host# set security application-firewall rule-sets rs1
   [edit security application-firewall rule-sets rs1]
   user@host# set rule r1 match dynamic-application [junos:FACEBOOK-ACCESS
     junos:GOOGLE-TALK junos:MEEBO]
   user@host# set rule r1 then deny
   user@host# set default-rule permit
   ```

2. Configure a policy to apply the rs1 application firewall rule set to traffic from IP 192.0.2.0 to IP 203.0.113.0 with the dev-abc, http-mgmt-accessible, or ftp-accessible user role.

   ```
   [edit]
   user@host# set security policies from-zone trust to-zone untrust policy user-role-fw2
     match source-address 192.0.2.0
   user@host# set security policies from-zone trust to-zone untrust policy user-role-fw2
     match destination-address 203.0.113.0
   user@host# set security policies from-zone trust to-zone untrust policy user-role-fw2
     match application http
   user@host# set security policies from-zone trust to-zone untrust policy user-role-fw2
     match source-identity [dev-abc http-mgmt-accessible ftp-accessible]
   user@host# set security policies from-zone trust to-zone untrust policy user-role-fw2
     then permit application-services application-firewall rule-set rs1
   ```


Copyright © 2016, Juniper Networks, Inc.

3. If you are done configuring the policy, commit the changes.

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

Results Verify that the AppFW rule set is configured properly. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security application-firewall
...
rule-sets rs1 {
  rule r1 {
   match {
    dynamic-application [junos:FACEBOOK-ACCESS junos:GOOGLE-TALK junos:MEEBO]
     }
   then {
    deny;
   }
  }
  default-rule {
   permit;
  }
}
```

Creating Remaining Security Policies Based on User and Role

Step-by-Step Procedure The following procedure configures policies for the remaining traffic.

1. Configure a policy to deny traffic with the same source and destination address but with different user and role criteria than specified in the user-role-fw2 policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw3
  match source-address 192.0.2.0
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw3
  match destination-address 203.0.113.0
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw3
  match application http
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw3
  match source-identity any
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw3
  then deny
```

2. Configure a security policy to permit all other HTTP traffic from zone trust to zone untrust.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw4
  match source-address any
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw4
  match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw4
  match application http
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw4
  match source-identity any
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security policies from-zone trust to-zone untrust policy user-role-fw4
  then permit
```

Results Verify the content and sequence of the user role firewall policies. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security policies

...
from-zone trust to-zone untrust {
 policy user-role-fw1 {
  match {
   source-address any;
   destination-address any;
   application http;
   source-identity unauthenticated-user
  }
  then {
   permit {
    application-services {
     uac-policy {
      captive-portal acs-device;
     }
    }
   }
  }
 }
}
from-zone trust to-zone untrust {
  policy user-role-fw2 {
   match {
    source-address 192.0.2.0;
    destination-address 203.0.113.0;
    application http;
    source-identity [dev-abc http-juniper-accessible ftp-accessible]
   }
   then {
    permit {
     application-services {
      application-firewall {
       rule-set rs1
      }
     }
    }
   }
  }
}
from-zone untrust to-zone trust {
  policy user-role-fw3 {
   match {
    source-address 192.0.2.0;
    destination-address 203.0.113.0;
    application http;
    source-identity any
   }
   then {
    deny
   }
```


Copyright © 2016, Juniper Networks, Inc.

```
  }
}
from-zone trust to-zone untrust {
 policy user-role-fw4 {
  match {
   source-address any;
   destination-address any;
   application http;
   source-identity any
  }
  then {
   permit
  }
 }
}
```

Related Documentation

- Understanding User Role Firewalls on page 1127
- User Role Retrieval and the Policy Lookup Process on page 1128
- Understanding the User Identification Table on page 1130

Configuring Resource Policies Using UAC

When using the user role firewall feature, resource policies are not necessary on the Access Control Service. If, however, resource policies exist, they are pushed to the SRX Series device at connection. You can create policies that use these resource policies by applying the UAC application service in the policy configuration. Table 78 shows three firewall policies that use the UAC resource policies exclusively:

Table 78: User Role Firewall Usage

policy-name	src-zone	dest-zone	src-IP	dest-IP	source-identity	application	action	Services
P1	zone1	zone2	any	192.0.2.1	any	http	permit	UTM
P2	zone1	zone2	any	net2	any	http	permit	IDP
P3	zone1	zone2	any	any	any	any	permit	UAC

The policies for traffic from zone1 to zone2 do not initiate user and role retrieval because any is specified in the source-identity field of every policy. In this example, traffic to the IP address 192.0.2.1 is permitted, but must meet processing requirements for the specified application service, in this case, UTM. Traffic to net2 is permitted and processed by the IDP processing requirements. Any remaining traffic is permitted and processed by the UAC processing requirements.

The configuration for this firewall policy would be as follows:

```
[edit]
user@host# show security policies
```

```
from-zone zone1 to-zone zone2 {
 policy P1 {
```

```
  match {
   source-address any;
   destination-address 192.0.2.1;
   source-identity any;
   application http;
  }
  then {
    permit {
    application-services {
     idp;
    }
   }
  }
 }
}
from-zone zone1 to-zone zone2 {
 policy P2 {
  match {
   source-address any;
   destination-address net2;
   source-identity any;
   application http;
  }
  then {
    permit {
    application-services {
     utm;
    }
   }
  }
 }
}
from-zone zone1 to-zone zone2 {
 policy P3 {
  match {
   source-address any;
   destination-address any;
   source-identity any;
   application any;
  }
  then {
    permit {
    application-services {
     uac-policy;
    }
   }
  }
 }
}
...
```

In this sample configuration, the action fields in P1 and P2 apply any requirements that have been configured for IDP and UTM respectively. By specifying the uac-policy option, the resource policies pushed to the SRX Series device determine whether the destination is accessible.

A user role firewall can implement both user role policies and the resource policies pushed from the Access Control Service. Table 79 shows the policies for three zone pairs.

Copyright © 2016, Juniper Networks, Inc.

Table 79: User Role Firewall Usage

policy-name	src-zone	dest-zone	src-IP	dest-IP	source-identity	application	action	Services
P1	zone1	zone2	any	any	unauthenticated-user	any	permit	UAC captive portal
P2	zone1	zone2	any	192.0.2.1	role2	http	permit	IDP
P3	zone1	zone2	any	net2	authenticated-user	http	permit	UTM
P4	zone1	zone2	any	any	any	any	permit	
P5	zone1	zone3	any	any	any	any	permit	UAC
P6	zone2	zone3	any	any	any	any	permit	UAC

Traffic from zone1 to zone2 is subject to one of four user role policies. The first of these policies uses the UAC captive portal to redirect unauthenticated users to the Access Control Service for authentication.

The access of traffic from zone1 to zone3 and from zone2 to zone3 is controlled by the resource policies pushed from the Access Control Service.

Related Documentation

- Understanding User Role Firewalls on page 1127
- User Role Retrieval and the Policy Lookup Process on page 1128
- Understanding the User Identification Table on page 1130
- Configuring a User Role Firewall For Captive Portal Redirection on page 1137

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 56

Setting Security Policy Reorder

- Understanding Security Policy Ordering on page 1149
- Example: Reordering the Policies on page 1151

Understanding Security Policy Ordering

Junos OS offers a tool for verifying that the order of policies in the policy list is valid.

It is possible for one policy to eclipse, or *shadow*, another policy. Consider the following examples:

Example 1

```
[edit]
user@host# set security zones security-zone trust interfaces ge-0/0/2 host-inbound-traffic
    system-services all
user@host# set security zones security-zone untrust interfaces ge-0/0/1
    host-inbound-traffic system-services all
user@host# set security policies from-zone trust to-zone untrust policy permit-all match
    source-address any
user@host# set security policies from-zone trust to-zone untrust match
    destination-address any
user@host# set security policies from-zone trust to-zone untrust match application any
user@host# set security policies from-zone trust to-zone untrust set then permit
user@host# set security policies from-zone untrust to-zone trust policy deny-all match
    source-address any
user@host# set security policies from-zone untrust to-zone trust policy deny-all match
    destination-address any
user@host# set security policies from-zone untrust to-zone trust policy deny-all match
    application any
user@host# set security policies from-zone untrust to-zone trust policy deny-all then
    deny
```

Example 2

```
[edit]
user@host# set security zones security-zone trust interfaces ge-0/0/2.0
    host-inbound-traffic system-services all
user@host# set security zones security-zone untrust interfaces ge-0/0/1.0
    host-inbound-traffic system-services all
user@host# set security address-book book1 address mail-untrust 192.0.2.1/24
user@host# set security address-book book1 attach zone untrust
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security address-book book2 address mail-trust 192.168.1.1/24
user@host# set security address-book book2 attach zone trust
user@host# set security policies from-zone trust to-zone untrust policy permit-mail match
    source-address mail-trust
user@host# set security policies from-zone trust to-zone untrust policy permit-mail match
    destination-address mail-untrust
user@host# set security policies from-zone trust to-zone untrust policy permit-mail match
    application junos-mail
user@host# set security policies from-zone trust to-zone untrust policy permit-mail then
    permit
```

In examples 1 and 2, where policy **permit-mail** is configured after policy **permit-all** from zone **trust** to zone **untrust**. All traffic coming from zone **untrust** matches the first policy **permit-all** and is allowed by default. No traffic matches policy **permit-mail**.

Because Junos OS performs a policy lookup starting from the top of the list, when it finds a match for traffic received, it does not look any lower in the policy list. To correct the previous example, you can simply reverse the order of the policies, putting the more specific one first:

```
[edit]
user@host# insert security policies from-zone trust to-zone untrust policy permit-mail
    before policy permit-all
```

In cases where there are dozens or hundreds of policies, the eclipsing of one policy by another might not be so easy to detect. To check if policies are being shadowed, enter any of the following commands:

```
[edit]
user@host# run show security shadow-policies logical-system lsys-name from-zone
    from-zone-name to-zone to-zone-name

[edit]
user@host# run show security shadow-policies logical-system lsys-name global
```

This command reports the shadowing and shadowed policies. It is then the administrator's responsibility to correct the situation.

NOTE: The concept of policy *shadowing* refers to the situation where a policy higher in the policy list always takes effect before a subsequent policy. Because the policy lookup always uses the first policy it finds that matches the five-part tuple of the source and destination zone, source and destination address, and application type, if another policy applies to the same tuple (or a subset of the tuple), the policy lookup uses the first policy in the list and never reaches the second one.

Related Documentation

- Security Policies Configuration Overview on page 1093
- Example: Configuring a Security Policy to Permit or Deny All Traffic on page 1094
- Example: Configuring a Security Policy to Permit or Deny Selected Traffic on page 1098

Copyright © 2016, Juniper Networks, Inc.

Example: Reordering the Policies

This example shows show how to move policies around after they have been created.

- Requirements on page 1151
- Overview on page 1151
- Configuration on page 1151
- Verification on page 1151

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure the address book and create addresses for use in the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.

Overview

To reorder policies to correct shadowing, you can simply reverse the order of the policies, putting the more specific one first.

Configuration

Step-by-Step Procedure

To reorder existing policies:

1. Reorder two existing policies by entering the following command:

   ```
   [edit]
   user@host# insert security policies from-zone trust to-zone untrust policy
     permit-mail before policy permit-all
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security policies** command.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Ordering on page 1149

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 57

Monitoring and Troubleshooting Security Policies

- Matching Security Policies on page 1153
- Tracking Policy Hit Counts on page 1155
- Best Practices for Defining Policies on High-End SRX Series Devices on page 1155
- Checking Memory Status on page 1157
- Synchronizing a Security Policy on SRX Series Devices on page 1159
- Verifying Scheduled Policies on page 1159
- Verifying Shadow Policies on page 1160
- Monitoring Policy Statistics on page 1162
- Troubleshooting Security Policies on page 1162

Matching Security Policies

The **show security match-policies** command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, the **show security match-policies** command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The **result-count** option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are "shadowed" by the first and are never encountered by matching traffic.

NOTE: The show security match-policies command is applicable only to security policies; IDP policies are not supported.

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1, to-zone z2 source-ip 10.10.10.1
destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
```

Copyright © 2016, Juniper Networks, Inc.

```
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 203.0.113.1/25
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.129/25
    d3: 192.0.2.1/24
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
    Source port range: [0-0]
    Destination port range: [21-21]
```

Example 2: Using the result-count Option

By default, the output list contains the policy that will be applied to traffic with the specified characteristics. To list more than one policy that match the criteria, use the **result-count** option. The first policy listed is always the policy that will be applied to matching traffic. If the **result-count** value is from 2 to 16, the output includes all policies that match the criteria up to the specified **result-count**. All policies listed after the first are "shadowed" by the first policy and are never applied to matching traffic.

Use this option to test the positioning of a new policy or to troubleshoot a policy that is not applied as expected for particular traffic.

In the following example, the traffic criteria matches two policies. The first policy listed, **p1**, contains the action applied to the traffic. Policy **p15** is shadowed by the first policy, and its action, therefore, will not be applied to matching traffic.

```
user@host> show security match-policies source-ip 10.10.10.1 destination-ip 192.0.2.5 source_port
1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
    Source port range: [1000-1030]
    Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
    Source port range: [1000-1030]
    Destination port range: [80-80]
```

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Tracking Policy Hit Counts

Use the **show security policies hit-count** command to display the utility rate of security policies according to the number of hits they receive. You can use this feature to determine which policies are being used on the device, and how frequently they are used. Depending on the command options that you choose, the number of hits can be listed without an order or sorted in either ascending or descending order, and they can be restricted to the number of hits that fall above or below a specific count or within a range. Data is shown for all zones associated with the policies or named zones.

Related Documentation

- Security Policies Overview on page 1085
- Troubleshooting Security Policies on page 1162
- Monitoring Policy Statistics on page 1162
- Matching Security Policies on page 1153

Best Practices for Defining Policies on High-End SRX Series Devices

A secure network is vital to a business. To secure a network, a network administrator must create a security policy that outlines all of the network resources within that business and the required security level for those resources. The security policy applies the security rules to the transit traffic within a context (from-zone to to-zone) and each policy is uniquely identified by its name. The traffic is classified by matching the source and destination zones, the source and destination addresses, and the application that the traffic carries in its protocol headers with the policy database in the data plane.

Table 80 provides the policy limitations for high-end SRX Series devices.

Table 80: Policy Limitations for High-End SRX Series Devices

Policy Limitations	SRX1500	SRX5400 SRX5600 SRX5800
Address objects	1024	1024
Application objects	3072	3072
Security policies	40,000	80,000
Policy contexts (zone pairs)	4096	8192
Policies per context	10,000	10,000
Policies with counting enabled	1024	1024

Copyright © 2016, Juniper Networks, Inc.

NOTE: The number of source and destination address objects allowed per firewall rule is 1024. The systemwide maximum allowed is 32,000 address objects.

Therefore, as you increase the number of addresses and applications in each rule, the amount of memory that is used by the policy definition increases, and sometimes the system runs out of memory with fewer than 80,000 policies.

To get the actual memory utilization of a policy on the Packet Forwarding Engine (PFE) and the Routing Engine (RE), you need to take various components of the memory tree into consideration. The memory tree includes the following two components:

- Policy context–Used to organize all policies in this context. Policy context includes variables such as source and destination zones.
- Policy entity–Used to hold the policy data. Policy entity calculates memory using parameters such as policy name, IP addresses, address count, applications, firewall authentication, WebAuth, IPsec, count, application services, and Junos Services Framework (JSF).

Additionally, the data structures used to store policies, rule sets, and other components use different memory on the Packet Forwarding Engine and on the Routing Engine. For example, address names for each address in the policy are stored on the Routing Engine, but no memory is allocated at the Packet Forwarding Engine level. Similarly, port ranges are expanded to prefix and mask pairs and are stored on the Packet Forwarding Engine, but no such memory is allocated on the Routing Engine.

Accordingly, depending on the policy configuration, the policy contributors to the Routing Engine are different from those to the Packet Forwarding Engine, and memory is allocated dynamically.

Memory is also consumed by the “deferred delete” state. In the deferred delete state, when an SRX Series device applies a policy change, there is transitory peak usage whereby both the old and new policies are present. So for a brief period, both old and new policies exist on the Packet Forwarding Engine, taking up twice the memory requirements.

Therefore, there is no definitive way to infer clearly how much memory is used by either component (Packet Forwarding Engine or Routing Engine) at any given point in time, because memory requirements are dependent on specific configurations of policies, and memory is allocated dynamically.

The following best practices for policy implementation enable you to better use system memory and to optimize policy configuration:

- Use single prefixes for source and destination addresses. For example, instead of using /32 addresses and adding each address separately, use a large subnet that covers most of the IP addresses you require.
- Use application “any” whenever possible. Each time you define an individual application in the policy, you can use an additional 52 bytes.

Copyright © 2016, Juniper Networks, Inc.

- Use fewer IPv6 addresses because IPv6 addresses consume more memory.
- Use fewer zone pairs in policy configurations. Each source or destination zone uses about 16,048 bytes of memory.
- The following parameters can change how memory is consumed by the bytes as specified:
  - Firewall authentication–About 16 bytes or more (unfixed)
  - Web authentication–About 16 bytes or more (unfixed)
  - IPsec–12 bytes
  - Application services–28 bytes
  - Count–64 bytes
- Check memory utilization before and after compiling policies.

NOTE: The memory requirement for each device is different. Some devices support 512,000 sessions by default and the bootup memory is usually at 72 to 73 percent. Other devices can have up to 1 million sessions and the bootup memory can be up to 83 to 84 percent. In the worst-case scenario, to support about 80,000 policies in the SPU, the SPU should boot with a flowd kernel memory consumption of up to 82 percent, and with at least 170 megabytes of memory available.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Global Address Books on page 1071
- Global Policy Overview on page 1115
- Example: Configuring a Global Policy with No Zone Restrictions on page 1117
- Checking Memory Status on page 1157

Checking Memory Status

Memory for flow entities (for example, policies, zones, addresses) on SRX1500, SRX5400, SRX5600, and SRX5800 devices is dynamically allocated. However, certain practices can help monitor the current memory usage on the device and optimize parameters to better size system configuration, especially during policy implementation.

You can isolate memory issues by comparing memory values before and after policy configurations.

To check memory usage:

Copyright © 2016, Juniper Networks, Inc.

- Use the **show chassis routing-engine** command to check overall Routing Engine (RE) memory usage. The following output from this command shows memory utilization at 39 percent:

```
user@host# show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
    DRAM                        1024 MB
    Memory utilization            39 percent
    CPU utilization:
      User                         0 percent
      Background                   0 percent
      Kernel                       2 percent
      Interrupt                    0 percent
      Idle                        97 percent
    Model                          RE-PPC-1200-A
    Start time                     2011-07-09 19:19:49 PDT
    Uptime                         37 days, 15 hours, 44 minutes, 13 seconds
    Last reboot reason             0x3:power cycle/failure watchdog
    Load averages:                 1 minute   5 minute  15 minute
                                       0.22       0.16       0.07
```

- Use the **show system processes extensive** command to acquire information on the processes running on the Routing Engine.

  Use the **find nsd** option in the **show system processes extensive** command to see direct usage on the Network Security Daemon (NSD) with its total memory in use as 10 megabytes and CPU utilization of 0 percent.

```
user@host# show system processes extensive | find nsd
 1182 root         1  96    0 10976K  5676K select   2:08  0.00% nsd
 1191 root         4   4    0  8724K  3764K select   1:57  0.00% slbd
 1169 root         1  96    0  8096K  3520K select   1:51  0.00% jsrpd
 1200 root         1   4    0     0K    16K peer_s   1:10  0.00% peer proxy
 1144 root         1  96    0  9616K  3528K select   1:08  0.00% lacpd
 1138 root         1  96    0  6488K  2932K select   1:02  0.00% ppmd
 1130 root         1  96    0  7204K  2208K select   1:02  0.00% craftd
 1163 root         1  96    0 16928K  5188K select   0:58  0.00% cosd
 1196 root         1   4    0     0K    16K peer_s   0:54  0.00% peer proxy
   47 root         1 -16    0     0K    16K sdflus   0:54  0.00% softdepflush
 1151 root         1  96    0 15516K  9580K select   0:53  0.00% appidd
  900 root         1  96    0  5984K  2876K select   0:41  0.00% eventd
```

- Check the configuration file size. Save your configuration file with a unique name before exiting the CLI. Then, enter the **ls -l filename** command from the shell prompt in the UNIX-level shell to check the file size as shown in the following sample output:

```
user@host> start shell
% ls -l config
-rw-r--r--  1 remote  staff  12681 Feb 15 00:43 config
```

Related Documentation

- Best Practices for Defining Policies on High-End SRX Series Devices on page 1155
- Security Policies Overview on page 1085
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

- Global Policy Overview on page 1115
- Example: Configuring a Global Policy with No Zone Restrictions on page 1117

Synchronizing a Security Policy on SRX Series Devices

Security policies are stored in both the Routing Engine (RE), and the Packet Forwarding Engine (PFE). When you modify the policies on the Routing Engine side, the policies are synchronized to the Packet Forwarding Engine side when you commit the configuration.

The policies in the Routing Engine and Packet Forwarding Engine must always be in sync for the configuration to commit successfully. Under certain circumstances, policies in the Routing Engine and the Packet Forwarding Engine might be out of sync resulting in generation of system core files upon commit completion.

When the policy configuration is modified and policies are out of sync, the following error message is displayed when you attempt to commit a configuration:

```
Policy is out of sync between RE and PFE <SPU-name(s)>. Please resync before
commit.

error: configuration check-out failed
```

To synchronize policies between the Routing Engine and the Packet Forwarding Engine, you must:

- Reboot the device (device in standalone mode)
- Reboot both devices (devices in a chassis cluster mode)

Related Documentation

- Security Policies Overview on page 1085
- Verifying a Security Policy Commit on page 1163
- Debugging Policy Lookup on page 1163
- Monitoring Policy Statistics on page 1162

Verifying Scheduled Policies

Purpose Display information about address books and zones.

Action Use the **show schedulers** CLI command to display information about schedulers configured on the system. If a specific scheduler is identified, detailed information is displayed for that scheduler only.

```
user@host# show schedulers
scheduler sche1 {
    /* This is sched1 */
    start-date 2006-11-02.12:12 stop-date 2007-11-02.12:11;
}
scheduler sche2 {
    daily {
        all-day;
```

```
    }
    sunday {
        start-time 16:00 stop-time 17:00;
    }
    friday {
        exclude;
    }
}
scheduler sche3 {
    start-date 2006-11-02.12:12 stop-date 2007-11-02.12:11;
    daily {
        start-time 10:00 stop-time 17:00
    }
    sunday {
        start-time 12:00 stop-time 14:00;
        start-time 16:00 stop-time 17:00;
    }
    monday {
        all-day;
    }
    friday {
        exclude;
    }
}
```

Meaning The output displays information about schedulers configured on the system. Verify the following information:

- Daily (recurrent) and one-time only (nonrecurrent) schedulers are configured correctly.
- Schedulers are active if policies are associated.

Related Documentation

- Security Policies Overview on page 1085
- Example: Configuring Schedulers for a Daily Schedule Excluding One Day on page 1124

Verifying Shadow Policies

- Verifying All Shadow Policies on page 1160
- Verifying a Policy Shadows One or More Policies on page 1161
- Verifying a Policy Is Shadowed by One or More Policies on page 1161

Verifying All Shadow Policies

Purpose Verify all the policies that shadows one or more policies.

Action From the operational mode, enter the following commands:

- For logical systems, enter the **show security shadow-policies logical-system** *lsys-name* **from-zone** *from-zone-name* **to-zone** *to-zone-name* command.

Copyright © 2016, Juniper Networks, Inc.

- For global policies, enter the **show security shadow-policies logical-system** *lsys-name* **global** command.

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b
    Policies              Shadowed policies
    P1                    P3
    P1                    P4
    P2                    P5
```

Meaning The output displays the list of all policies that shadows other policies. In this example, P1 policy shadows P3 and P4 policies and P2 policy shadows P5 policy.

Verifying a Policy Shadows One or More Policies

Purpose Verify if a given policy shadows one or more policies positioned after it.

Action From the operational mode, enter the following commands:

- For logical systems, enter the **show security shadow-policies logical-system** *lsys-name* **from-zone** *from-zone-name* **to-zone** *to-zone-name* **policy** *policy-name* command.
- For global policies, enter the **show security shadow-policies logical-system** *lsys-name* **global policy** *policy-name* command.

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P1
   Policies              Shadowed policies
    P1                    P3
    P1                    P4
```

Meaning The output displays all the policies that are shadowed by the given policy. In this example, P1 policy shadows P3 and P4 policies.

Verifying a Policy Is Shadowed by One or More Policies

Purpose Verify if a given policy is shadowed by one or more positioned before it.

Action From the operational mode, enter the following commands:

- For logical systems, enter the **show security shadow-policies logical-system** *lsys-name* **from-zone** *from-zone-name* **to-zone** *to-zone-name* **policy** *policy-name* **reverse** command.
- For global policies, enter the **show security shadow-policies logical-system** *lsys-name* **global policy** *policy-name* **reverse** command.

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse
   Policies              Shadowed policies
    P1                    P4
```

Meaning The output displays the given policy shadowed by one or more policies. In this example, P4 policy is shadowed by P1 policy.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Security Policy Ordering on page 1149
- Example: Reordering the Policies on page 1151

Monitoring Policy Statistics

Purpose Monitor and record traffic that Junos OS permits or denies based on previously configured policies.

Action To monitor traffic, enable the count and log options.

Count—Configurable in an individual policy. If count is enabled, statistics are collected for sessions that enter the device for a given policy, and for the number of packets and bytes that pass through the device in both directions for a given policy. For counts (only for packets and bytes), you can specify that alarms be generated whenever the traffic exceeds specified thresholds. See **count (Security Policies)**.

Log—Logging capability can be enabled with security policies during session initialization (**session-init**) or session close (**session-close**) stage. See **log (Security Policies)**.

- To view logs from denied connections, enable log on **session-init**.
- To log sessions after their conclusion/tear-down, enable log on **session-close**.

NOTE: **Session log is enabled at real time in the flow code which impacts the user performance. If both session-close and session-init are enabled, performance is further degraded as compared to enabling session-init only.**

For details about information collected for session logs, see "Information Provided in Session Log Entries for SRX Series Services Gateways" on page 1787.

Related Documentation

- Security Policies Overview on page 1085
- Troubleshooting Security Policies on page 1162
- Checking a Security Policy Commit Failure on page 1163
- Verifying a Security Policy Commit on page 1163
- Debugging Policy Lookup on page 1163

Troubleshooting Security Policies

- Checking a Security Policy Commit Failure on page 1163
- Verifying a Security Policy Commit on page 1163
- Debugging Policy Lookup on page 1163

Copyright © 2016, Juniper Networks, Inc.

Checking a Security Policy Commit Failure

Problem **Description:** Most policy configuration failures occur during a commit or runtime.

Commit failures are reported directly on the CLI when you execute the CLI command **commit-check** in configuration mode. These errors are configuration errors, and you cannot commit the configuration without fixing these errors.

Solution To fix these errors, do the following:

1. Review your configuration data.
2. Open the file /var/log/nsd_chk_only. This file is overwritten each time you perform a commit check and contains detailed failure information.

Verifying a Security Policy Commit

Problem **Description:** Upon performing a policy configuration commit, if you notice that the system behavior is incorrect, use the following steps to troubleshoot this problem:

Solution

1. Operational **show** Commands—Execute the operational commands for security policies and verify that the information shown in the output is consistent with what you expected. If not, the configuration needs to be changed appropriately.
2. Traceoptions—Set the **traceoptions** command in your policy configuration. The flags under this hierarchy can be selected as per user analysis of the **show** command output. If you cannot determine what flag to use, the flag option **all** can be used to capture all trace logs.

   user@host# set security policies traceoptions <flag all>

   You can also configure an optional filename to capture the logs.

   user@host# set security policies traceoptions <*filename*>

   If you specified a filename in the trace options, you can look in the /var/log/<*filename*> for the log file to ascertain if any errors were reported in the file. (If you did not specify a filename, the default filename is eventd.) The error messages indicate the place of failure and the appropriate reason.

   After configuring the trace options, you must recommit the configuration change that caused the incorrect system behavior.

Debugging Policy Lookup

Problem **Description:** When you have the correct configuration, but some traffic was incorrectly dropped or permitted, you can enable the **lookup** flag in the security policies traceoptions. The **lookup** flag logs the lookup related traces in the trace file.

Solution **user@host# set security policies traceoptions <flag lookup>**

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *Synchronizing Policies Between Routing Engine and Packet Forwarding Engine*
- Checking a Security Policy Commit Failure on page 1163
- Verifying a Security Policy Commit on page 1163
- Debugging Policy Lookup on page 1163
- Monitoring Policy Statistics on page 1162

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 58

Handling Security Policy Violations

- Understanding Searching and Sorting Audit Log on page 1165
- Understanding Packet Flow Alarms and Auditing on page 1166
- Example: Generating a Security Alarm in Response to Policy Violations on page 1167

Understanding Searching and Sorting Audit Log

An audit administrator analyzes the audit trail, reviews the audit record, and deletes the audit trail for maintenance purposes. The search and sort capability provides an efficient mechanism to the audit administrator for viewing pertinent audit information. This helps the audit administrator to identify potential security violations and take action against possible security breaches. The audit log can be viewed by all the administrators (such as Audit, Cryptographic, Security, and IDS administrators) . An IDS audit log can be viewed only by IDS audit administrator.

The security administrator can configure audit events and set thresholds that could indicate a potential security violation. The device monitors the occurrences of these events and notifies the administrator after an event has occurred or a set threshold has been met.

The audit administrator can search or group the audit log data based on the following:

- Destination subject identity
- Source subject identity
- Range of date, time, user identities, subject service identifiers, or Transport Layer protocol
- Rule identity
- User identity
- Network interface
- Success of auditable security events
- Failure of auditable security events

Copyright © 2016, Juniper Networks, Inc.

NOTE:

- The device sends an alarm to the console or the security alarm system when the in-memory audit event log exceeds the limit configured by the security administrator. The device then overwrites the oldest log messages with the new audit event log messages.
- During system reboot the device does a commit of the existing configuration including login classes. Therefore, there will be audit log entries for all user-defined classes indicating that they have been modified.

Related Documentation

- Understanding Packet Flow Alarms and Auditing on page 1166
- Example: Generating a Security Alarm in Response to Policy Violations on page 1167

Understanding Packet Flow Alarms and Auditing

Alarms are triggered when packets are dropped because of a policy violation. A policy violation occurs when a packet matches a reject or deny policy. A policy violation alarm is generated when the system monitors any of the following audited events:

- Number of policy violations by a source network identifier within a specified period
- Number of policy violations to a destination network identifier within a specified period
- Number of policy violations to an application within a specified period
- Policy rule or group of rule violations within a specified period

There are four types of alarms corresponding to these four events. The alarms are based on source IP, destination IP, application, and policy.

When a packet encounters a reject or deny policy, the policy violation counters for all enabled types of alarm are increased. When any counter reaches the specified threshold within a specified period, an alarm is generated. After a specified period, the policy violation counter is reset and reused to start another counting cycle.

To view the alarm information, run the **show security alarms** command. The violation count and the alarm do not persist across system reboots. After a reboot, the violation count resets to zero and the alarm is cleared from the alarm queue.

After taking appropriate actions, you can clear the alarm. The alarm remains in the queue until you clear it (or until you reboot the device). To clear the alarm, run the **clear security alarms** command. After you clear the alarm, a subsequent series of flow policy violations can cause a new alarm to be raised.

Related Documentation

- Example: Setting an Audible Alert as Notification of a Security Alarm on page 6779
- Example: Generating a Security Alarm in Response to Policy Violations on page 1167

Copyright © 2016, Juniper Networks, Inc.

Example: Generating a Security Alarm in Response to Policy Violations

This example shows how to configure the device to generate a system alarm when a policy violation occurs. By default, no alarm is raised when a policy violation occurs.

- Requirements on page 1167
- Overview on page 1167
- Configuration on page 1167
- Verification on page 1168

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an alarm to be raised when:

- The application size is 10240 units.
- The source IP violation exceeds 1000 within 20 seconds.
- The destination IP violations exceeds 1000 within 10 seconds.
- The policy match violation exceeds 100, with a size of 100 units.

Configuration

CLI Quick Configuration

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

```
set security alarms potential-violation policy application size 10240
set security alarms potential-violation policy source-ip threshold 1000 duration 20
set security alarms potential-violation policy destination-ip threshold 1000 duration 10
set security alarms potential-violation policy policy-match threshold 100 size 100
```

Step-by-Step Procedure

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

To configure policy violation alarms:

1. Enable security alarms.

   ```
   [edit]
   user@host# edit security alarms
   ```

2. Specify that an alarm should be raised when an application violation occurs.

   ```
   [edit security alarms potential-violation policy]
   user@host# set application size 10240
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Specify that an alarm should be raised when a source IP violation occurs.

   ```
   [edit security alarms potential-violation policy]
   user@host# set source-ip threshold 1000 duration 20
   ```

4. Specify that an alarm should be raised when a destination IP violation occurs.

   ```
   [edit security alarms potential-violation policy]
   user@host# set destination-ip threshold 1000 duration 10
   ```

5. Specify that an alarm should be raised when a policy match violation occurs.

   ```
   [edit security alarms potential-violation policy]
   user@host# set policy-match threshold 100 size 100
   ```

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

```
policy {
    source-ip {
        threshold 1000;
        duration 20;
    }
    destination-ip {
        threshold 1000;
        duration 10;
    }
    application {
        size 10240;
    }
    policy-match {
        threshold 100;
        size 100;
    }
}
```

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

Verification

To confirm that the configuration is working properly, from operational mode, enter the **show security alarms** command.

Related Documentation

- Understanding Packet Flow Alarms and Auditing on page 1166

Copyright © 2016, Juniper Networks, Inc.

PART 14

Configuring Security Policy Applications

- Configuring Applications and Application Sets on page 1171
- Configuring Custom Policy Applications on page 1175
- Setting Policy Application Timeout on page 1181
- Understanding Predefined Policy Applications on page 1185

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 59

Configuring Applications and Application Sets

- Security Policy Applications Overview on page 1171
- Policy Application Sets Overview on page 1172
- Example: Configuring Applications and Application Sets on page 1172

Security Policy Applications Overview

Applications are types of traffic for which protocol standards exist. Each application has a transport protocol and destination port number(s) associated with it, such as TCP/port 21 for FTP and TCP/port 23 for Telnet. When you create a policy, you must specify an application for it.

You can select one of the predefined applications from the application book, or a custom application or application set that you created. You can see which application you can use in a policy by using the **show applications** CLI command.

NOTE: Each predefined application has a source port range of 1–65535, which includes the entire set of valid port numbers. This prevents potential attackers from gaining access by using a source port outside of the range. If you need to use a different source port range for any predefined application, create a custom application. For information, see "Understanding Custom Policy Applications" on page 1175.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

Policy Application Sets Overview

When you create a policy, you must specify an application, or service, for it to indicate that the policy applies to traffic of that type. Sometimes the same applications or a subset of them can be present in multiple policies, making it difficult to manage. Junos OS allows you to create groups of applications called application sets. Application sets simplify the process by allowing you to manage a small number of application sets, rather than a large number of individual application entries.

The application (or application set) is referred to by security policies as match criteria for packets initiating sessions. If the packet matches the application type specified by the policy and all other criteria match, then the policy action is applied to the packet.

You can specify the name of an application set in a policy. In this case, if all of the other criteria match, any one of the applications in the application set serves as valid matching criteria; **any** is the default application name that indicates all possible applications.

Applications are created in the **.../applications/application/*application-name*** directory. You do not need to configure an application for any of the services that are predefined by the system.

In addition to predefined services, you can configure a custom service. After you create a custom service, you can refer to it in a policy.

Related Documentation

- Security Policy Applications Overview on page 1171
- Custom Application Mappings on page 1175
- Understanding Policy Application Timeout Configuration and Lookup on page 1181
- Example: Configuring Applications and Application Sets on page 1172

Example: Configuring Applications and Application Sets

This example shows how to configure applications and application sets.

- Requirements on page 1172
- Overview on page 1172
- Configuration on page 1173
- Verification on page 1173

Requirements

Before you begin, configure the required applications. See "Policy Application Sets Overview" on page 1172.

Overview

Rather than creating or adding multiple individual application names to a policy, you can create an application set and refer to the name of the set in a policy. For example, for a

Copyright © 2016, Juniper Networks, Inc.

group of employees, you can create an application set that contains all the approved applications.

In this example, you create an application set that are used to log into the servers in the ABC (intranet) zone, to access the database, and to transfer files.

- Define the applications in the configured application set.
- Managers in zone A and managers in zone B use these services. Therefore, give the application set a generic name, such as MgrAppSet.
- Create an application set for the applications that are used for e-mail and Web-based applications that are delivered by the two servers in the external zone.

Configuration

Step-by-Step Procedure

To configure an application and application set:

1. Create an application set for managers.

   ```
   [edit applications]
   user@host# set application-set MgrAppSet application junos-ssh
   user@host# set application-set MgrAppSet application junos-telnet
   ```

2. Create another application set for e-mail and Web-based applications.

   ```
   [edit applications]
   user@host# set application-set WebMailApps application junos-smtp
   user@host# set application-set WebMailApps application junos-pop3
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show applications** command.

Verifying the Configuration

Related Documentation

- Security Policies Overview on page 1085
- Security Policy Applications Overview on page 1171

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 60

Configuring Custom Policy Applications

- Understanding Custom Policy Applications on page 1175
- Custom Application Mappings on page 1175
- Example: Adding and Modifying Custom Policy Applications on page 1176
- Example: Configuring Custom Policy Application Term Options on page 1177

Understanding Custom Policy Applications

If you do not want to use predefined applications in your policy, you can easily create custom applications.

You can assign each custom application the following attributes:

- Name
- Transport protocol
- Source and destination port numbers for applications using TCP or UDP
- Type and code values for applications using ICMP
- Timeout value

Related Documentation

- Custom Application Mappings on page 1175
- Understanding Policy Application Timeout Configuration and Lookup on page 1181
- Understanding Policy Application Timeouts Contingencies on page 1182
- Example: Adding and Modifying Custom Policy Applications on page 1176

Custom Application Mappings

The application option specifies the Layer 7 application that maps to the Layer 4 application that you reference in a policy. A predefined application already has a mapping to a Layer 7 application. However, for custom applications, you must link the application to an application explicitly, especially if you want the policy to apply an Application Layer Gateway (ALG) or deep inspection to the custom application.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Junos OS supports ALGs for numerous applications, including DNS, FTP, H.323, HTTP, RSH, SIP, Telnet, and TFTP.

Applying an ALG to a custom application involves the following two steps:

- Define a custom application with a name, timeout value, transport protocol, and source and destination ports.
- When configuring a policy, reference that application and the application type for the ALG that you want to apply.

Related Documentation

- Understanding Custom Policy Applications on page 1175
- Understanding Policy Application Timeout Configuration and Lookup on page 1181
- Understanding Policy Application Timeouts Contingencies on page 1182
- Example: Adding and Modifying Custom Policy Applications on page 1176

Example: Adding and Modifying Custom Policy Applications

This example shows how to add and modify custom policy applications.

- Requirements on page 1176
- Overview on page 1176
- Configuration on page 1177
- Verification on page 1177

Requirements

Before you begin, create addresses and security zones. See "Example: Creating Security Zones" on page 1051.

Overview

In this example, you create a custom application using the following information:

- A name for the application: **cust-telnet**.
- A range of source port numbers: **1** through **65535**.
- A destination port number: 23000.
- The protocol used by the application: TCP.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

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

To add and modify a custom policy application:

1. Configure TCP and specify the source port and destination port.

   ```
   [edit applications application cust-telnet]
   user@host# set protocol tcp source-port 1–65535 destination-port 23000
   ```

2. Specify the length of time that the application is inactive.

   ```
   [edit applications application cust-telnet]
   user@host# set inactivity-timeout 1800
   ```

3. Modify a custom policy application.

   ```
   [edit applications application cust-telnet]
   user@host# delete protocol tcp
   user@host# set application-protocol ftp
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show applications application** command.

NOTE: The timeout value is in seconds. If you do not set it, the timeout value of a custom application is 1800 seconds. If you do not want an application to time out, type never.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Custom Policy Applications on page 1175
- Example: Defining a Custom ICMP Application on page 1199

Example: Configuring Custom Policy Application Term Options

This example shows how to configure applications properties and term options for application protocols.

- Requirements on page 1178
- Overview on page 1178

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1178
- Verification on page 1180

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- A PC

Before you begin:

- Configure the required applications. See "Example: Adding and Modifying Custom Policy Applications" on page 1176 .

Overview

In this example, you create an application name, app-name, and a term called custom-options to define your custom policy application term options.

You configure Domain Name Service (DNS) as the Application Layer Gateway (ALG) type and UDP as the networking protocol type. You set the source port to 24000 and the destination port to 23000. Then you set the Internet Control Message Protocol (ICMP) packet type value to 5 and the ICMP code value to 0. You set the remote procedure call (RPC) program number value to 50 and the Universal Unique Identifier (UUID) value to 1be617c0-31a5-11cf-a7d8-00805f48a135. Finally, you set the inactivity-timeout value to 60.

Configuration

CLI Quick Configuration

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

```
user@host# set applications application app-name term custom-options
user@host# set applications application app-name term custom-options alg dns
user@host#set applications application app-name term custom-options protocol udp
user@host#set applications application app-name term custom-options source-port
  24000
user@host#set applications application app-name term custom-options destination-port
  23000
user@host#set applications application app-name term custom-options icmp-type 5
user@host#set applications application app-name term custom-options icmp-code 0
user@host#set applications application app-name term custom-options
  rpc-program-number 50
user@host#set applications application app-name term custom-options uuid
  1be617c0-31a5-11cf-a7d8-00805f48a135
user@host#set applications application app-name term custom-options inactivity-timeout
  60
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To configure custom policy application term options:

1. Configure the term name.

   ```
   [edit applications]
   user@host# set application app-name term custom-options
   ```

2. Configure the ALG type.

   ```
   [edit applications]
   user@host# set application app-name term custom-options alg dns
   ```

3. Configure the networking protocol type.

   ```
   [edit applications]
   user@host# set application app-name term custom-options protocol udp
   ```

4. Configure the source port number.

   ```
   [edit applications]
   user@host#set application app-name term custom-options source-port 24000
   ```

5. Configure the TCP or UDP destination port number.

   ```
   [edit applications]
   user@host# set application app-name term custom-options destination-port 23000
   ```

6. Specify the application type value.

   ```
   [edit applications]
   user@host# set application app-name term custom-options icmp-type 5
   ```

7. Specify the application code value.

   ```
   [edit applications]
   user@host# set application app-name term custom-options icmp-code 0
   ```

8. Specify the RPC program number.

   ```
   [edit applications]
   user@host#  set application app-name term custom-options rpc-program-number
     50
   ```

9. Specify the UUID value.

   ```
   [edit applications]
   user@host# set application app-name term custom-options uuid
     1be617c0-31a5-11cf-a7d8-00805f48a135
   ```

10. Specify the inactivity timeout value.

    ```
    [edit applications]
    user@host# set application app-name term custom-options inactivity-timeout 60
    ```

Results

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

```
[edit]
  user@host# show applications
  application app-name {
```

```
    term custom-options alg dns protocol udp source-port 24000 icmp-type 5 icmp-code
      0 rpc-program-number 50 uuid 1be617c0-31a5-11cf-a7d8-00805f48a135
      inactivity-timeout 60;
}
```

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

Verification

Verifying the Configuration

Purpose Verify that the configuration is correct.

Action From operational mode, enter the **show applications** command.

```
user@host> show applications
  application app-name {
  term custom-options alg dns protocol udp source-port 24000 icmp-type 5 icmp-code
    0 rpc-program-number 50 uuid 1be617c0-31a5-11cf-a7d8-00805f48a135
    inactivity-timeout 60;
}
```

Related Documentation

- Security Policy Applications Overview on page 1171
- Understanding Custom Policy Applications on page 1175

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 61

Setting Policy Application Timeout

- Understanding Policy Application Timeout Configuration and Lookup on page 1181
- Understanding Policy Application Timeouts Contingencies on page 1182
- Example: Setting a Policy Application Timeout on page 1182

Understanding Policy Application Timeout Configuration and Lookup

The application timeout value you set for an application determines the session timeout. You can set the timeout threshold for a predefined or custom application; you can use the application default timeout, specify a custom timeout, or use no timeout at all.

Application timeout values are stored in the root TCP and UDP port-based timeout table and in the protocol-based default timeout table. When you set an application timeout value, Junos OS updates these tables with the new value. There are also default timeout values in the applications entry database, which are taken from predefined applications. You can set a timeout, but you cannot alter a default value.

Each custom application can be configured with its own custom application timeout. If multiple custom applications are configured with custom timeouts, then each application will have its own custom application timeout.

If the application that is matched for the traffic has a timeout value, that timeout value is used. Otherwise, the lookup proceeds in the following order until an application timeout value is found:

1. The root TCP and UDP port-based timeout table is searched for a timeout value.
2. The protocol-based default timeout table is searched for a timeout value. See Table 81.

Table 81: Protocol-Based Default Timeout

Protocol	Default Timeout (seconds)
TCP	1800
UDP	60
ICMP	60
OSPF	60

Copyright © 2016, Juniper Networks, Inc.

Table 81: Protocol-Based Default Timeout *(continued)*

Protocol	Default Timeout (seconds)
Other	1800

Related Documentation

- Understanding Custom Policy Applications on page 1175
- Understanding Policy Application Timeouts Contingencies on page 1182
- Custom Application Mappings on page 1175
- Example: Adding and Modifying Custom Policy Applications on page 1176

Understanding Policy Application Timeouts Contingencies

When setting timeouts, be aware of the following contingencies:

- If an application contains several application rule entries, all rule entries share the same timeout. You need to define the application timeout only once. For example, if you create an application with two rules, the following commands will set the timeout to **20** seconds for both rules:

  ```
  user@host# set applications application test protocol tcp destination-port 1035-1035
  inactivity-timeout 20
  user@host# set applications application test term test protocol udp
  user@host# set applications application test term test source-port 1-65535
  user@host# set applications application test term test destination-port 1111-1111
  ```

- If multiple custom applications are configured with custom timeouts, then each application will have its own custom application timeout. For example:

  ```
  user@host# set applications application ftp-1 protocol tcp source-port 0-65535 destination-port
  2121-2121 inactivity-timeout 10
  user@host# set applications application telnet-1 protocol tcp source-port 0-65535
  destination-port 2300-2348 inactivity-timeout 20
  ```

  With this configuration, Junos OS applies a 10-second timeout for destination port **2121** and a 20-second timeout for destination port **2300** in an application group.

Related Documentation

- Understanding Custom Policy Applications on page 1175
- Custom Application Mappings on page 1175
- Understanding Policy Application Timeout Configuration and Lookup on page 1181
- Example: Adding and Modifying Custom Policy Applications on page 1176

Example: Setting a Policy Application Timeout

This example shows how to set a policy application timeout value.

- Requirements on page 1183
- Overview on page 1183

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1183
- Verification on page 1183

Requirements

Before you begin, understand policy application timeouts. See "Understanding Policy Application Timeout Configuration and Lookup" on page 1181.

Overview

Application timeout values are stored in the application entry database and in the corresponding vsys TCP and UDP port-based timeout tables. In this example, you set the device for a policy application timeout to 75 minutes for the FTP predefined application.

When you set an application timeout value, Junos OS updates these tables with the new value.

Configuration

Step-by-Step Procedure

To set a policy application timeout:

1. Set the inactivity timeout value.

   ```
   [edit applications application ftp]
   user@host# set inactivity-timeout 75
   ```

2. Commit the configuration if you are done configuring the device.

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

Verification

To verify the configuration is working properly, enter the **show applications** command.

Related Documentation

- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 62

Understanding Predefined Policy Applications

- Understanding Internet-Related Predefined Policy Applications on page 1185
- Understanding Microsoft Predefined Policy Applications on page 1187
- Understanding Dynamic Routing Protocols Predefined Policy Applications on page 1188
- Understanding Streaming Video Predefined Policy Applications on page 1188
- Understanding Sun RPC Predefined Policy Applications on page 1189
- Understanding Security and Tunnel Predefined Policy Applications on page 1190
- Understanding IP-Related Predefined Policy Applications on page 1191
- Understanding Instant Messaging Predefined Policy Applications on page 1191
- Understanding Management Predefined Policy Applications on page 1192
- Understanding Mail Predefined Policy Applications on page 1193
- Understanding UNIX Predefined Policy Applications on page 1194
- Understanding Miscellaneous Predefined Policy Applications on page 1194
- Understanding the ICMP Predefined Policy Application on page 1195
- Example: Defining a Custom ICMP Application on page 1199
- Default Behavior of ICMP Unreachable Errors on page 1201

Understanding Internet-Related Predefined Policy Applications

When you create a policy, you can specify predefined Internet-related applications for the policy.

Table 82 lists Internet-related predefined applications. Depending on your network requirements, you can choose to permit or deny any or all of these applications. Each entry lists the application name, default receiving port, and application description.

Table 82: Predefined Applications

Application Name	Port(s)	Application Description
AOL	5190-5193	America Online Internet service provider (ISP) provides Internet, chat, and instant messaging applications.

Copyright © 2016, Juniper Networks, Inc.

Table 82: Predefined Applications *(continued)*

Application Name	Port(s)	Application Description
DHCP relay	67 (default)	Dynamic Host Configuration Protocol.
DHCP	68 client 67 server	Dynamic Host Configuration Protocol allocates network addresses and delivers configuration parameters from server to hosts.
DNS	53	Domain Name System translates domain names into IP addresses.
FTP	20 data 21 control	File Transfer Protocol (FTP) allows the sending and receiving of files between machines. You can choose to deny or permit ANY or to selectively permit or deny. We recommend denying FTP applications from untrusted sources (Internet).
Gopher	70	Gopher organizes and displays Internet servers' contents as a hierarchically structured list of files. We recommend denying Gopher access to avoid exposing your network structure.
HTTP	80	HyperText Transfer Protocol is the underlying protocol used by the World Wide Web (WWW). Denying HTTP application disables your users from viewing the Internet. Permitting HTTP application allows your trusted hosts to view the Internet.
HTTP-EXT	—	Hypertext Transfer Protocol with extended nonstandard ports
HTTPS	443	Hypertext Transfer Protocol with Secure Sockets Layer (SSL) is a protocol for transmitting private documents through the Internet. Denying HTTPS disables your users from shopping on the Internet and from accessing certain online resources that require secure password exchange. Permitting HTTPS allows your trusted hosts to participate in password exchange, shop online, and visit various protected online resources that require user login.
Internet Locator Service	—	Internet Locator Service includes LDAP, User Locator Service, and LDAP over TSL/SSL.
IRC	6665-6669	Internet Relay Chat (IRC) allows people connected to the Internet to join live discussions.
LDAP	389	Lightweight Directory Access Protocol is a set of protocols used to access information directories.
PC-Anywhere	—	PC-Anywhere is a remote control and file transfer software.
TFTP	69	Trivial File transfer Protocol (TFTP) is a protocol for simple file transfer.

Copyright © 2016, Juniper Networks, Inc.

Table 82: Predefined Applications *(continued)*

Application Name	Port(s)	Application Description
WAIS	—	Wide Area Information Server is a program that finds documents on the Internet.

Related Documentation

- Understanding Dynamic Routing Protocols Predefined Policy Applications on page 1188
- Example: Configuring Applications and Application Sets on page 1172

Understanding Microsoft Predefined Policy Applications

When you create a policy, you can specify predefined Microsoft applications for the policy.

Table 83 lists predefined Microsoft applications, parameters associated with each application, and a brief description of each application. Parameters include universal unique identifiers (UUIDs) and TCP/UDP source and destination ports. A UUID is a 128-bit unique number generated from a hardware address, a timestamp, and seed values.

Table 83: Predefined Microsoft Applications

Application	Parameter/UUID	Description
Junos MS-RPC-EPM	135 e1af8308-5d1f-11c9-91a4-08002b14a0fa	Microsoft remote procedure call (RPC) Endpoint Mapper (EPM) Protocol.
Junos MS-RPC	—	Any Microsoft remote procedure call (RPC) applications.
Junos MS-RPC-MSEXCHANGE	3 members	Microsoft Exchange application group includes: • Junos-MS-RPC-MSEXCHANGE-DATABASE • Junos-MS-RPC-MSEXCHANGE-DIRECTORY • Junos-MS-RPC-MSEXCHANGE-INFO-STORE
Junos-MS-RPC-MSEXCHANGE-DATABASE	1a190310-bb9c-11cd-90f8-00aa00466520	Microsoft Exchange Database application.
Junos-MS-RPC-MSEXCHANGE-DIRECTORY	f5cc5a18-4264-101a-8c59-08002b2f8426 f5cc5a7c-4264-101a-8c59-08002b2f8426 f5cc59b4-4264-101a-8c59-08002b2f8426	Microsoft Exchange Directory application.
Junos-MS-RPC-MSEXCHANGE-INFO-STORE	0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde 1453c42c-0fa6-11d2-a910-00c04f990f3b 10f24e8e-0fa6-11d2-a910-00c04f990f3b 1544f5e0-613c-11d1-93df-00c04fd7bd09	Microsoft Exchange Information Store application.

Copyright © 2016, Juniper Networks, Inc.

Table 83: Predefined Microsoft Applications *(continued)*

Application	Parameter/UUID	Description
Junos-MS-RPC-TCP	—	Microsoft Transmission Control Protocol (TCP) application.
Junos-MS-RPC-UDP	—	Microsoft User Datagram Protocol (UDP) application.
Junos-MS-SQL	—	Microsoft Structured Query Language (SQL).
Junos-MSN	—	Microsoft Network Messenger application.

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172
- Policy Application Sets Overview on page 1172

Understanding Dynamic Routing Protocols Predefined Policy Applications

When you create a policy, you can specify predefined dynamic routing protocol applications for the policy.

Depending on your network requirements, you can choose to permit or deny messages generated from these dynamic routing protocols and packets of these dynamic routing protocols. Table 84 lists each supported dynamic routing protocol by name, port, and description.

Table 84: Dynamic Routing Protocols

Dynamic Routing Protocol	Port	Description
RIP	520	RIP is a common distance-vector routing protocol.
OSPF	89	OSPF is a common link-state routing protocol.
BGP	179	BGP is an exterior/interdomain routing protocol.

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172

Understanding Streaming Video Predefined Policy Applications

When you create a policy, you can specify predefined streaming video applications for the policy.

Table 85 lists each supported streaming video application by name and includes the default port and description. Depending on your network requirements, you can choose to permit or deny any or all of these applications.

Copyright © 2016, Juniper Networks, Inc.

Table 85: Supported Streaming Video Applications

Application	Port	Description
H.323	TCP source 1-65535; TCP destination 1720, 1503, 389, 522, 1731 UDP source 1-65535; UDP source 1719	H.323 is a standard approved by the International Telecommunication Union (ITU) that defines how audiovisual conference data is transmitted across networks.
NetMeeting	TCP source 1-65535; TCP destination 1720, 1503, 389, 522 UDP source 1719	Microsoft NetMeeting uses TCP to provide teleconferencing (video and audio) applications over the Internet.
Real media	TCP source 1-65535; TCP destination 7070	Real Media is streaming video and audio technology.
RTSP	554	Real-Time Streaming Protocol (RTSP) is for streaming media applications
SIP	5056	Session Initiation Protocol (SIP) is an Application-Layer control protocol for creating, modifying, and terminating sessions.
VDO Live	TCP source 1-65535; TCP destination 7000-7010	VDOLive is a scalable, video streaming technology.

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172

Understanding Sun RPC Predefined Policy Applications

When you create a policy, you can specify predefined Sun RPC applications for the policy.

Table 86 lists each Sun remote procedure call Application Layer Gateway (RPC ALG) application name, parameters, and full name.

Table 86: RPC ALG Applications

Application	Program Numbers	Full Name
SUN-RPC-PORTMAPPER	111100000	Sun RPC Portmapper protocol
SUN-RPC-ANY	ANY	Any Sun RPC applications
SUN-RPC-PROGRAM-MOUNTD	100005	Sun RPC Mount Daemon
SUN-RPC-PROGRAM-NFS	100003 100227	Sun RPC Network File System
SUN-RPC-PROGRAM-NLOCKMGR	100021	Sun RPC Network Lock Manager

Copyright © 2016, Juniper Networks, Inc.

Table 86: RPC ALG Applications *(continued)*

Application	Program Numbers	Full Name
SUN-RPC-PROGRAM-RQUOTAD	100011	Sun RPC Remote Quota Daemon
SUN-RPC-PROGRAM-RSTATD	100001	Sun RPC Remote Status Daemon
SUN-RPC-PROGRAM-RUSERD	100002	Sun RPC Remote User Daemon
SUN-RPC-PROGRAM-SADMIND	100232	Sun RPC System Administration Daemon
SUN-RPC-PROGRAM-SPRAYD	100012	Sun RPC Spray Daemon
SUN-RPC-PROGRAM-STATUS	100024	Sun RPC Status
SUN-RPC-PROGRAM-WALLD	100008	Sun RPC Wall Daemon
SUN-RPC-PROGRAM-YPBIND	100007	SUN RPC Yellow Page Bind application

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172

Understanding Security and Tunnel Predefined Policy Applications

When you create a policy, you can specify predefined security and tunnel applications for the policy.

Table 87 lists each supported application and gives the default port(s) and a description of each entry.

Table 87: Supported Applications

Application	Port	Description
IKE	UDP source 1-65535; UDP destination 500	Internet Key Exchange is the protocol that sets up a security association in the IPsec protocol suite. Internet Key protocol (IKE) is a protocol to obtain authenticated keying material for use with ISAKMP.
IKE-NAT	4500	IKE-Network Address Translation (NAT) performs Layer 3 NAT for S2C IKE traffic.
L2TP	1701	L2TP combines PPTP with Layer 2 Forwarding (L2F) for remote access.
PPTP	1723	Point-to-Point Tunneling Protocol allows corporations to extend their own private network through private *tunnels* over the public Internet.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Example: Configuring Applications and Application Sets on page 1172

Understanding IP-Related Predefined Policy Applications

When you create a policy, you can specify predefined IP-related applications for the policy.

Table 88 lists the predefined IP-related applications. Each entry includes port numbers and a description of the application.

TCP-ANY means any application that is using TCP, so there is no default port for it. The same is true for UDP-ANY.

Table 88: Predefined IP-Related Applications

Application	Port	Description
Any	—	Any application
TCP-ANY	0-65,535	Any protocol using the TCP
UDP-ANY	0-65,535	Any protocol using the UDP

Related Documentation
- Example: Configuring Applications and Application Sets on page 1172

Understanding Instant Messaging Predefined Policy Applications

When you create a policy, you can specify predefined instant messaging applications for the policy.

Table 89 lists predefined Internet-messaging applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 89: Predefined Internet-Messaging Applications

Application	Port	Description
Gnutella	6346 (default)	Gnutella is a public domain file sharing protocol that operates over a distributed network. You can assign any port, but the default is 6346.
MSN	1863	Microsoft Network Messenger is a utility that allows you to send instant messages and talk online.
NNTP	119	Network News Transport Protocol is a protocol used to post, distribute, and retrieve USENET messages.
SMB	445	Server Message Block (SMB) over IP is a protocol that allows you to read and write files to a server on a network.

Copyright © 2016, Juniper Networks, Inc.

Table 89: Predefined Internet-Messaging Applications *(continued)*

Application	Port	Description
YMSG	5010	Yahoo! Messenger is a utility that allows you to check when others are online, send instant messages, and talk online.

Related Documentation

- Understanding Management Predefined Policy Applications on page 1192
- Example: Configuring Applications and Application Sets on page 1172

Understanding Management Predefined Policy Applications

When you create a policy, you can specify predefined management applications for the policy.

Table 90 lists the predefined management applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 90: Predefined Management Applications

Application	Port	Description
NBNAME	137	NetBIOS Name application displays all NetBIOS name packets sent on UDP port 137.
NDBDS	138	NetBIOS Datagram application, published by IBM, provides connectionless (datagram) applications to PCs connected with a broadcast medium to locate resources, initiate sessions, and terminate sessions. It is unreliable and the packets are not sequenced.
NFS	—	Network File System uses UDP to allow network users to access shared files stored on computers of different types. SUN RPC is a building block of NFS.
NS Global	—	NS-Global is the central management protocol for Juniper Networks Firewall/VPN devices.
NS Global PRO	—	NS Global-PRO is the scalable monitoring system for the Juniper Networks Firewall/VPN device family.
NSM	—	Network and Security Manager
NTP	123	Network Time Protocol provides a way for computers to synchronize to a time reference.
RLOGIN	513	RLOGIN starts a terminal session on a remote host.
RSH	514	RSH executes a shell command on a remote host.
SNMP	161	Simple Network Management Protocol is a set of protocols for managing complex networks.

Copyright © 2016, Juniper Networks, Inc.

Table 90: Predefined Management Applications *(continued)*

Application	Port	Description
SQL*Net V1	66	SQL*Net Version 1 is a database language that allows for the creation, access, modification, and protection of data.
SQL*Net V2	66	SQL*Net Version 2 is a database language that allows for the creation, access, modification, and protection of data.
MSSQL	1433 (default instance)	Microsoft SQL is a proprietary database server tool that allows for the creation, access, modification, and protection of data.
SSH	22	SSH is a program to log into another computer over a network through strong authentication and secure communications on an unsecure channel.
SYSLOG	514	Syslog is a UNIX program that sends messages to the system logger.
Talk	517-518	Talk is a visual communication program that copies lines from your terminal to that of another user.
Telnet	23	Telnet is a UNIX program that provides a standard method of interfacing terminal devices and terminal-oriented processes to each other.
WinFrame	—	WinFrame is a technology that allows users on non-Windows machines to run Windows applications.
X-Windows	—	X-Windows is the windowing and graphics system that Motif and OpenLook are based on.

Related Documentation
- Example: Configuring Applications and Application Sets on page 1172

Understanding Mail Predefined Policy Applications

When you create a policy, you can specify predefined mail applications for the policy.

Table 91 lists the predefined mail applications. Each includes the name of the application, the default or assigned port number, and a description of the application.

Table 91: Predefined Mail Applications

Application	Port	Description
IMAP	143	Internet Message Access Protocol is used for retrieving messages.
Mail (SMTP)	25	Simple Mail Transfer Protocol is used to send messages between servers.
POP3	110	Post Office Protocol is used for retrieving e-mail.

Related Documentation
- Example: Configuring Applications and Application Sets on page 1172

Copyright © 2016, Juniper Networks, Inc.

Understanding UNIX Predefined Policy Applications

When you create a policy, you can specify predefined UNIX applications for the policy.

Table 92 lists the predefined UNIX applications. Each entry includes the name of the application, the default or assigned port, and a description of the application.

Table 92: Predefined UNIX Applications

Application	Port	Description
FINGER	79	Finger is a UNIX program that provides information about the users.
UUCP	117	UNIX-to-UNIX Copy Protocol (UUCP) is a UNIX utility that enables file transfers between two computers over a direct serial or modem connection.

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172

Understanding Miscellaneous Predefined Policy Applications

When you create a policy, you can specify miscellaneous predefined applications for the policy.

Table 93 lists predefined miscellaneous applications. Each entry includes the application name, default or assigned port, and a description of the application.

Table 93: Predefined Miscellaneous Applications

Application	Port	Description
CHARGEN	19	Character Generator Protocol is a UDP- or TCP-based debugging and measurement tool.
DISCARD	9	Discard protocol is an Application Layer protocol that describes a process for discarding TCP or UDP data sent to port 9.
IDENT	113	Identification protocol is a TCP/IP Application Layer protocol used for TCP client authentication.
LPR	515 listen; 721-731 source range (inclusive)	Line Printer Daemon protocol is a TCP-based protocol used for printing applications.
RADIUS	1812	Remote Authentication Dial-In User Service application is a server program used for authentication and accounting purposes.
RADIUS Accounting	1813	A RADIUS Accounting server receives statistical data about users logging in to or out of a LAN.

Copyright © 2016, Juniper Networks, Inc.

Table 93: Predefined Miscellaneous Applications *(continued)*

Application	Port	Description
SQLMON	1434 (SQL Monitor Port)	SQL monitor (Microsoft)
VNC	5800	Virtual Network Computing facilitates viewing and interacting with another computer or mobile Juniper Networks device connected to the Internet.
WHOIS	43	Network Directory Application Protocol is a way to look up domain names.
SCCP	2000	Cisco Station Call Control Protocol (SCCP) uses the signaling connection control port to provide high availability and flow control.

Related Documentation

- Example: Configuring Applications and Application Sets on page 1172

Understanding the ICMP Predefined Policy Application

When you create a policy, you can specify the ICMP predefined application for the policy.

Internet Control Message Protocol (ICMP) is a part of IP and provides a way to query a network (ICMP query messages) and to receive feedback from the network for error patterns (ICMP error messages). ICMP does not, however, guarantee error message delivery or report all lost datagrams; and it is not a reliable protocol. ICMP codes and type codes describe ICMP query messages and ICMP error messages.

You can choose to permit or deny any or specific types of ICMP messages to improve network security. Some types of ICMP messages can be exploited to gain information about your network that might compromise security. For example, ICMP, TCP, or UDP packets can be constructed to return ICMP error messages that contain information about a network, such as its topology, and access list filtering characteristics. Table 94 lists ICMP message names, the corresponding code, type, and description.

Table 94: ICMP Messages

ICMP Message Name	Type	Code	Description
ICMP-ANY	all	all	ICMP-ANY affects any protocol using ICMP. Denying ICMP-ANY impairs any attempt to ping or monitor a network using ICMP. Permitting ICMP-ANY allows all ICMP messages.

Copyright © 2016, Juniper Networks, Inc.

Table 94: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-ADDRESS-MASK • Request • Reply	17 18	0 0	ICMP address mask query is used for systems that need the local subnet mask from a bootstrap server. Denying ICMP address mask request messages can adversely affect diskless systems. Permitting ICMP address mask request messages might allow others to fingerprint the operating system of a host in your network.
ICMP-DEST-UNREACH	3	0	ICMP destination unreachable error message indicates that the destination host is configured to reject the packets. Codes 0, 1, 4, or 5 can be from a gateway. Codes 2 or 3 can be from a host (RFC 792). Denying ICMP destination unreachable error messages can remove the assumption that a host is up and running behind an SRX Series device. Permitting ICMP destination unreachable error messages can allow some assumptions, such as security filtering, to be made about the network.
ICMP Fragment Needed	3	4	ICMP fragmentation error message indicates that fragmentation is needed but the don't fragment flag is set. We recommend denying these messages from the Internet to an internal network.
ICMP FragmentReassembly	11	1	ICMP fragment reassembly time exceeded error indicates that a host reassembling a fragmented message ran out of time and dropped the packet. This message is sometimes sent. We recommend denying these messages from the Internet (external) to the trusted (internal) network.
ICMP-HOST-UNREACH	3	1	ICMP host unreachable error messages indicate that routing table entries do not list or list as infinity a particular host. Sometimes this error is sent by gateways that cannot fragment when a packet requiring fragmentation is received. We recommend denying these messages from the Internet to a trusted network. Permitting these messages allows others to be able to determine your internal host's IP addresses by a process of elimination or make assumptions about gateways and fragmentation.

Copyright © 2016, Juniper Networks, Inc.

Table 94: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-INFO • Request • Reply	15 16	0 0	ICMP-INFO query messages allow diskless host systems to query the network and self-configure. Denying ICMP address mask request messages can adversely affect diskless systems. Permitting ICMP address mask request messages might allow others to broadcast information queries to a network segment to determine computer type.
ICMP-PARAMETER-PROBLEM	12	0	ICMP parameter problem error messages notify you when incorrect header parameters are present and have caused a packet to be discarded We recommend denying these messages from the Internet to a trusted network. Permitting ICMP parameter problem error messages allows others to make assumptions about your network.
ICMP-PORT-UNREACH	3	3	ICMP port unreachable error messages indicate that gateways processing datagrams requesting certain ports are unavailable or unsupported in the network. We recommend denying these messages from the Internet to a trusted network. Permitting ICMP port unreachable error messages can allow others to determine which ports you use for certain protocols.
ICMP-PROTOCOL-UNREACH	3	2	ICMP protocol unreachable error messages indicate that gateways processing datagrams requesting certain protocols are unavailable or unsupported in the network. We recommend denying these messages from the Internet to a trusted network. Permitting ICMP protocol unreachable error messages can allow others to determine what protocols your network is running.
ICMP-REDIRECT	5	0	ICMP redirect network error messages are sent by an SRX Series device. We recommend denying these messages from the Internet to a trusted network.
ICMP-REDIRECT-HOST	5	1	ICMP redirect messages indicate datagrams destined for the specified host to be sent along another path.

Copyright © 2016, Juniper Networks, Inc.

Table 94: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
ICMP-REDIRECT-TOS-HOST	5	3	ICMP redirect type of service (TOS) and host error is a type of message.
ICMP-REDIRECT-TOS-NET	5	2	ICMP redirect TOS and network error is a type of message.
ICMP-SOURCE-QUENCH	4	0	ICMP source quench error message indicates that a device does not have the buffer space available to accept, queue, and send the packets on to the next hop. Denying these messages will not help or impair internal network performance. Permitting these messages can allow others to know that a device is congested, making it a viable attack target.
ICMP-SOURCE-ROUTE-FAIL	3	5	ICMP source route failed error message We recommend denying these messages from the Internet (external).
ICMP-TIME-EXCEEDED	11	0	ICMP time-to-live (TTL) exceeded error message indicates that a packet's TTL setting reached zero before the packet reached its destination. This ensures that older packets are discarded before resent ones are processed. We recommend denying these messages from a trusted network out to the Internet.
ICMP-TIMESTAMP • Request • Reply	13 14	0 0	ICMP-TIMESTAMP query messages provide the mechanism to synchronize time and coordinate time distribution in a large, diverse network.
Ping (ICMP ECHO)	8	0	Ping is a utility to determine whether a specific host is accessible by its IP address. Denying ping functionality removes your ability to check to see if a host is active. Permitting ping can allow others to execute a denial-of-service (DoS) or Smurf attack.
ICMP-ECHO-FRAGMENT-ASSEMBLY-EXPIRE	11	1	ICMP fragment echo reassembly time expired error message indicates that the reassembly time was exceeded. We recommend denying these messages.

Copyright © 2016, Juniper Networks, Inc.

Table 94: ICMP Messages *(continued)*

ICMP Message Name	Type	Code	Description
Traceroute • Forward • Discard	30 30	0 1	Traceroute is a utility to indicate the path to access a specific host. We recommend denying this utility from the Internet (external) to your trusted network (internal).

Related Documentation

- Default Behavior of ICMP Unreachable Errors on page 1201
- Example: Configuring Applications and Application Sets on page 1172

Example: Defining a Custom ICMP Application

This example shows how to define a custom ICMP application.

- Requirements on page 1200
- Overview on page 1200
- Configuration on page 1201
- Verification on page 1201

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Understand custom policy application. See "Understanding Custom Policy Applications" on page 1175.
- Understand the ICMP predefined policy application. See "Understanding the ICMP Predefined Policy Application" on page 1195.

Overview

Junos OS supports ICMP—as well as several ICMP messages—as predefined or custom applications. When configuring a custom ICMP application, you define a type and code.

- There are different message types within ICMP. For example:
  - type 0 = Echo Request message
  - type 3 = Destination Unreachable message
- An ICMP message type can also have a message code. The code provides more specific information about the message, as shown in Table 95.

Table 95: Message Descriptions

Message Type	Message Code
5 = Redirect	0 = Redirect datagram for the network (or subnet)
	1 = Redirect datagram for the host
	2 = Redirect datagram for the type of application and network
	3 = Redirect datagram for the type of application and host
11 = Time Exceeded Codes	0 = Time to live exceeded in transit
	1 = Fragment reassembly time exceeded

Junos OS supports any type or code within the range of **0** through **55** .

In this example, you define a custom application named host-unreachable using ICMP as the transport protocol. The type is 3 (for destination unreachable) and the code is 1 (for host unreachable). You set the timeout value at 4 minutes.

NOTE: For more information about ICMP types and codes, refer to RFC 792, *Internet Control Message Protocol*.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

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

To define a custom ICMP application:

1. Set the application type and code.

   ```
   [edit applications application host-unreachable]
   user@host# set icmp-type 5 icmp-code 0
   ```

2. Set the inactivity timeout value.

   ```
   [edit applications application host-unreachable]
   user@host# set inactivity-timeout 4
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show applications** command.

Related Documentation

- Security Policies Overview on page 1085

Default Behavior of ICMP Unreachable Errors

For different levels of security, the default behavior for ICMP unreachable errors is handled as follows:

- Sessions are closed for ICMP type-3, code-0, code-1, code-2, and code-3 messages only when the following conditions are met:
  - The ICMP unreachable message is received in the server–to-client direction.
  - No normal packet is received in the server-to-client direction.

  Otherwise, sessions do not close.
- Sessions do not close for ICMP type-3, code-4 messages.

Related Documentation

- Understanding the ICMP Predefined Policy Application on page 1195
- Example: Configuring Applications and Application Sets on page 1172

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 15

Configuration Statements and Operational Commands

- Configuration Statements on page 1205
- Operational Commands on page 1339

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 63

Configuration Statements

- Applications Configuration Statement Hierarchy on page 1208
- Security Configuration Statement Hierarchy on page 1209
- [edit security address-book] Hierarchy Level on page 1210
- [edit security policies] Hierarchy Level on page 1211
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security zones] Hierarchy Level on page 1216
- address (Security Address Book) on page 1218
- address-book on page 1219
- address-set on page 1220
- alarms (Security) on page 1221
- alarm-threshold on page 1222
- alarm-without-drop on page 1223
- application (Applications) on page 1224
- application (Security Alarms) on page 1227
- application (Security Policies) on page 1228
- application-protocol (Applications) on page 1229
- application-services (Security Policies) on page 1230
- application-tracking (Security Zones) on page 1231
- application-traffic-control (Application Services) on page 1231
- attach on page 1232
- audible (Security Alarms) on page 1232
- authentication (Security Alarms) on page 1233
- authentication-source on page 1234
- captive-portal (Services UAC Policy) on page 1234
- count (Security Policies) on page 1235
- default-policy on page 1235
- deny (Security Policies) on page 1236
- description (Applications) on page 1236

Copyright © 2016, Juniper Networks, Inc.

- description (Security Address Book) on page 1237
- description (Security Policies) on page 1238
- description (Security Zone) on page 1239
- destination-address (Security Policies) on page 1240
- destination-address (Security Policies Flag) on page 1241
- destination-address-excluded on page 1242
- destination-ip (Security Alarms) on page 1243
- destination-port (Applications) on page 1244
- dns-proxy on page 1248
- dynamic-dns on page 1249
- exclude (Schedulers) on page 1250
- firewall-authentication (Security Policies) on page 1251
- firewall-authentication (User Identification) on page 1252
- forward-only (DNS) on page 1252
- from-zone (Security Policies) on page 1253
- from-zone (Security Policies Global) on page 1255
- functional-zone on page 1256
- global (Security Policies) on page 1257
- host-inbound-traffic on page 1259
- icmp-code (Applications) on page 1260
- icmp-type (Applications) on page 1260
- inactivity-timeout (Applications) on page 1261
- interfaces (Security Zones) on page 1262
- initial-tcp-mss on page 1263
- ipsec-group-vpn (Security Policies) on page 1264
- ipsec-vpn (Security Policies) on page 1264
- local-authentication-table on page 1265
- log (Security Policies) on page 1265
- management (Security Zones) on page 1266
- match (Security Policies) on page 1267
- match (Security Policies Global) on page 1268
- no-policy-cold-synchronization on page 1269
- pair-policy on page 1270
- pass-through on page 1271
- permit (Security Policies) on page 1272
- policies on page 1274
- policy (Security Alarms) on page 1279

Copyright © 2016, Juniper Networks, Inc.

- policy (Security Policies) on page 1280
- policy-match on page 1282
- policy-rematch on page 1283
- policy-stats on page 1284
- potential-violation on page 1285
- protocol (Applications) on page 1287
- protocols (Security Zones Host Inbound Traffic) on page 1288
- protocols (Security Zones Interfaces) on page 1290
- range-address on page 1291
- redirect-wx (Application Services) on page 1291
- reject (Security) on page 1292
- reverse-tcp-mss on page 1293
- rpc-program-number (Applications) on page 1294
- scheduler (Security Policies) on page 1295
- scheduler-name on page 1296
- schedulers (Security Policies) on page 1296
- screen (Security Zones) on page 1297
- secure-domains on page 1297
- secure-neighbor-discovery on page 1298
- security-zone on page 1299
- sequence-check-required on page 1300
- services-offload (Security) on page 1300
- session-close on page 1301
- session-init on page 1301
- simple-mail-client-service on page 1302
- source-address (Security Policies) on page 1303
- source-address-excluded on page 1303
- source-identity on page 1304
- source-ip (Security Alarms) on page 1305
- source-port (Applications) on page 1306
- ssl-proxy (Application Services) on page 1306
- ssl-termination-profile on page 1307
- start-date on page 1307
- start-time (Schedulers) on page 1308
- stop-date on page 1309
- stop-time on page 1310
- syn-check-required on page 1310

Copyright © 2016, Juniper Networks, Inc.

- system-services (Security Zones Host Inbound Traffic) on page 1311
- system-services (Security Zones Interfaces) on page 1313
- tcp-options (Security Policies) on page 1315
- tcp-rst on page 1316
- term (Applications) on page 1316
- then (Security Policies) on page 1317
- to-zone (Security Policies) on page 1319
- to-zone (Security Policies Global) on page 1321
- traceoptions (Security Policies) on page 1322
- traceoptions (Security User Identification) on page 1324
- traceoptions (System Services DNS) on page 1326
- tunnel (Security Policies) on page 1328
- uac-policy (Application Services) on page 1328
- unified-access-control (Security) on page 1329
- user-firewall on page 1330
- user-identification on page 1331
- utm-policy on page 1332
- uuid (Applications) on page 1333
- vrrp on page 1334
- web-authentication on page 1335
- web-redirect on page 1335
- zones on page 1336

Applications Configuration Statement Hierarchy

Use the statements in the **applications** configuration hierarchy to configure applications properties and group applications objects.

```
applications {
  application application-name {
    application-protocol (dns | ftp |gprs-gtp-c| gprs-gtp-u | gprs-gtp-v0 |gprs-sctp | http
      | ignore | ike-esp-nat | mgcp-ca | mgcp-ua | ms-rpc | q931 | ras | realaudio | rsh | rtsp
      | sccp | sip | sqlnet-v2 | sun-rpc | talk | tftp);
    description text;
    destination-port port-identifier;
    do-not-translate-A-query-to-AAAA-query;
    do-not-translate-AAAA-query-to-A-query;
    ether-type hex-value;
    icmp-code value;
    icmp-type value;
    icmp6-code value;
    icmp6-type value;
    inactivity-timeout (seconds | never);
    protocol number;
    rpc-program-number number;
```


Copyright © 2016, Juniper Networks, Inc.

```
    source-port port-number;
    term term-name {
      alg application;
      destination-port port-identifier;
      icmp-code value;
      icmp-type value;
      icmp6-code value;
      icmp6-type value;
      inactivity-timeout (seconds | never);
      protocol number;
      rpc-program-number number;
      source-port port-number;
      uuid hex-value;
    }
    uuid hex-value;
  }
  application-set application-set-name {
    application application-name;
    application-set application-set-name;
    description text;
  }
}
```

Related Documentation

- application (Applications) on page 1224

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526

Copyright © 2016, Juniper Networks, Inc.

- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

[edit security address-book] Hierarchy Level

```
security {
  address-book (book-name | global) {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
      }
      address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
      }
      attach {
        zone zone-name;
      }
      description text;
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Address Books on page 1069

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          destination-address-excluded;
          source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          source-address-excluded;
          source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
          }
        }
        scheduler-name scheduler-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
then {
    count {
        alarm {
            per-minute-threshold number;
            per-second-threshold number;
        }
    }
    deny;
    log {
        session-close;
        session-init;
    }
    permit {
        application-services {
            application-firewall {
                rule-set rule-set-name;
            }
            application-traffic-control {
                rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
                profile-name profile-name;
            }
            uac-policy {
                captive-portal captive-portal;
            }
            utm-policy policy-name;
        }
        destination-address {
            drop-translated;
            drop-untranslated;
        }
        firewall-authentication {
            pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
            }
            user-firewall {
                access-profile profile-name;
                domain domain-name
                ssl-termination-profile profile-name;
            }
            web-authentication {
                client-match user-or-group-name;
            }
        }
        services-offload;
        tcp-options {
            sequence-check-required;
```


Copyright © 2016, Juniper Networks, Inc.

```
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
  global {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        from-zone {
          [zone-name];
          any;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
        to-zone {
          [zone-name];
          any;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
```

```
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
          captive-portal captive-portal;
        }
        utm-policy policy-name;
      }
      destination-address {
        drop-translated;
        drop-untranslated;
      }
      firewall-authentication {
        pass-through {
          access-profile profile-name;
          client-match user-or-group-name;
          ssl-termination-profile profile-name;
          web-redirect;
          web-redirect-to-https;
        }
        user-firewall {
          access-profile profile-name;
          domain domain-name
          ssl-termination-profile profile-name;
        }
        web-authentication {
          client-match user-or-group-name;
        }
      }
      services-offload;
      tcp-options {
        initial-tcp-mss mss-value;
        reverse-tcp-mss mss-value;
        sequence-check-required;
        syn-check-required;
      }
    }
    reject;
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      policy-rematch;
      policy-stats {
        system-wide (disable | enable);
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

[edit security user-identification] Hierarchy Level

```
security {
  user-identification {
    authentication-source {
      active-directory-authentication-table priority priority;
      firewall-authentication priority priority;
      local-authentication-table priority priority;
      unified-access-control priority priority;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- authentication-source on page 1234
- Security Configuration Statement Hierarchy on page 603

[edit security zones] Hierarchy Level

```
security {
  zones {
    functional-zone {
      management {
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
        interfaces interface-name  {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
      }
    }
    security-zone zone-name {
      address-book {
        address address-name {
          ip-prefix {
            description text;
          }
          description text;
          dns-name domain-name {
            ipv4-only;
            ipv6-only;
          }
          range-address lower-limit to upper-limit;
          wildcard-address ipv4-address/wildcard-mask;
        }
        address-set address-set-name {
          address address-name;
          address-set address-set-name;
          description text;
        }
      }
      application-tracking;
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
            system-services service-name {
                except;
            }
        }
        interfaces interface-name {
            host-inbound-traffic {
                protocols protocol-name {
                    except;
                }
                system-services service-name {
                    except;
                }
            }
        }
        screen screen-name;
        tcp-rst;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Security Zones and Interfaces Overview on page 1049

Copyright © 2016, Juniper Networks, Inc.

address (Security Address Book)

Syntax

```
address address-name {
  ip-prefix {
    description text;
  }
  description text;
  dns-name domain-name {
    ipv4-only;
    ipv6-only;
  }
  range-address lower-limit to upper-limit;
  wildcard-address ipv4-address/wildcard-mask;
}
```

Hierarchy Level [edit security address-book *book-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for address range added in Junos OS Release 12.1. The **description** option added in Junos OS Release 12.1.

Description Add an entry containing an IP address or DNS hostname, or wildcard address to the address book. An address book contains entries for addressable entities in security zones, policies, and NAT rules. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, and DNS names.

Options

- **address *address-name*—**Name of an address entry.
- **description *text*—**Descriptive text about an address entry.
- dns-address ***domain-name*—**DNS address name.
- ***ip-prefix*—**IP address with prefix.
- **range-address *lower-limit* to *upper-limit*—**Address range for an address book.
- **wildcard-address *ipv4-address/wildcard-mask*—**IPv4 wildcard address in the form of A.B.C.D/wildcard-mask.

NOTE: **IPv6 wildcard address configuration is not supported in this release.**

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security address-book] Hierarchy Level on page 643

Copyright © 2016, Juniper Networks, Inc.

address-book

Syntax

```
address-book (book-name | global) {
  address address-name {
    ip-prefix {
      description text;
    }
    description text;
    dns-name domain-name {
      ipv4-only;
      ipv6-only;
    }
    range-address lower-limit to upper-limit;
    wildcard-address ipv4-address/wildcard-mask;
  }
  address-set address-set-name {
    address address-name;
    address-set address-set-name;
    description text;
  }
  attach {
    zone zone-name;
  }
  description text;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OSRelease 11.1. Statement moved under the security hierarchy in Junos OS Release 11.2. Support for address range added in Junos OS Release 12.1. The **description** option added in Junos OS Release 12.1.

Description Define entries in the address book. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use them when configuring different features, such as security policies and NAT.

NOTE: IPv6 wildcard address configuration is not supported in this release.

Options

- **address-book *book-name*—**Name of the address book.
- **global—**An address book that is available by default. You can add any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, DNS names, or address range to the global address book. You do not need to attach the global address book to a security zone; entries in the global address book are available to all security zones that are not attached to address books.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

address-set

Syntax

```
address-set address-set-name {
    address address-name;
    address-set address-set-name;
    description text;
}
```

Hierarchy Level
[edit security address-book *book-name*]

Release Information
Statement introduced in Junos OS Release 8.5. Support for nested address sets introduced in Release 11.2 of Junos OS. The **description** option added in Junos OS Release 12.1.

Description
Specify a collection of addresses, as defined in the **address (Address Book)** statement. Using address sets, you can organize addresses in logical groups and use them to easily configure other features, such as policies and NAT rules. Using this statement, you can also include a description for an address set.

You can also define address sets within address sets.

Options
address-set-name—Name of the address set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security address-book] Hierarchy Level on page 643

Copyright © 2016, Juniper Networks, Inc.

alarms (Security)

Syntax

```
alarms {
  audible {
    continuous;
  }
  potential-violation {
    authentication failures;
    cryptographic-self-test;
    decryption-failures {
      threshold value;
    }
    encryption-failures {
      threshold value;
    }
    idp;
    ike-phase1-failures {
      threshold value;
    }
    ike-phase2-failures {
      threshold value;
    }
    key-generation-self-test;
    non-cryptographic-self-test;
    policy {
      application {
        duration interval;
        size count;
        threshold value;
      }
      destination-ip {
        duration interval;
        size count;
        threshold value;
      }
      policy match {
        duration interval;
        size count;
        threshold value;
      }
      source-ip {
        duration interval;
        size count;
        threshold value;
      }
    }
    replay-attacks {
      threshold value;
    }
    security-log-percent-full percentage;
  }
}
```

Hierarchy Level [edit security]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 11.2.

Description Configures security alarms.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085

alarm-threshold

Syntax alarm-threshold *number*;

Hierarchy Level [edit security screen ids-option *screen-name* tcp syn-flood]

Release Information Statement modified in Junos OS Release 9.2.

Description Define the number of half-complete proxy connections per second at which the device makes entries in the event alarm log.

Options *number* —Threshold value.

Range: 1 through 500,000 per second

Default: 1024 per second

NOTE: For SRX Series devices the applicable range is 1 through 1,000,000 per second.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

alarm-without-drop

Syntax alarm-without-drop;

Hierarchy Level [edit security screen ids-option *screen-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Direct the device to generate an alarm when detecting an attack but not block the attack. Use this statement to allow an attack to enter a segment of your network that you have previously prepared to receive it (for example, a honeynet, which is a decoy network with extensive monitoring capabilities).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application (Applications)

Syntax

```
application application-name {
  application-protocol (dns | ftp |gprs-gtp-c| gprs-gtp-u | gprs-gtp-v0 |gprs-sctp | http |
    ignore | ike-esp-nat | mgcp-ca | mgcp-ua | ms-rpc | q931 | ras | realaudio | rsh | rtsp | sccp
    | sip | sqlnet-v2 | sun-rpc | talk | tftp);
  description text;
  destination-port port-identifier;
  do-not-translate-A-query-to-AAAA-query;
  do-not-translate-AAAA-query-to-A-query;
  ether-type hex-value;
  icmp-code value;
  icmp-type value;
  icmp6-code value;
  icmp6-type value;
  inactivity-timeout (seconds | never);
  protocol number;
  rpc-program-number number;
  source-port port-number;
  term term-name {
    alg application;
    destination-port port-identifier;
    icmp-code value;
    icmp-type value;
    icmp6-code value;
    icmp6-type value;
    inactivity-timeout (seconds | never);
    protocol number;
    rpc-program-number number;
    source-port port-number;
    uuid hex-value;
  }
  uuid hex-value;
}
```

Hierarchy Level [edit applications]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure application properties at the [applications] hierarchy level.

Options

- **application-protocol *protocol-name***—Specify the name of the application protocol.
- **description *text***—Describe the application.
- **destination-port *port-identifier***—Specify a TCP or UDP destination port number.
- **do-not-translate-A-query-to-AAAA-query**—Set the statement to control the translation of A query to AAAA query.
- **do-not-translate-AAAA-query-to-A-query**—Set the statement to control the translation of AAAA query to A query.
- **ether-type *value***—Specify the Ethernet packet type value.

Copyright © 2016, Juniper Networks, Inc.

- **icmp-code *value***—Specify the Internet Control Message Protocol (ICMP) code value.

  Range: 0 through 255.

- **icmp-type *value***—Specify the ICMP packet type value.

  Range: 0 through 255.

- **icmp6-code *value***—Specify the ICMP6 code value.

  Range: 0 through 255.

- **icmp6-type *value***—Specify the ICMP6 packet type value.

  Range: 0 through 255.

- **inactivity-timeout (*seconds* | never)**—Specify the amount of time the application is inactive before it times out in seconds.

  Range: 4 through 129,600 seconds.

  Default: For TCP, 1800 seconds; for UDP, 60 seconds.

- **protocol *value***—Specify the networking protocol name or number.

- **rpc-program-number *value***—Specify the remote procedure call (RPC) or Distributed Computing Equipment (DCE) value.

- **source-port *port-number***—Specify a TCP or UDP source port number.

Copyright © 2016, Juniper Networks, Inc.

- **term** ***term-name***—Specify the individual application protocol.
  - **alg** ***application***—Specify the name of the application protocol.
  - **destination-port** ***port-identifier***—Specify a TCP or UDP destination port number.
  - **icmp-code** ***value***—Specify the ICMP code value.

    Range: 0 through 255.
  - **icmp-type** ***value***—Specify the ICMP packet type value.

    Range: 0 through 255.
  - **icmp6-code** ***value***—Specify the ICMP6 code value.

    Range: 0 through 255.
  - **icmp6-type** ***value***—Specify the ICMP6 packet type value.

    Range: 0 through 255.
  - **inactivity-timeout (*seconds* | never)**—Specify the amount of time the application is inactive before it times out in seconds.

    Range: 4 through 129,600 seconds.

    Default: For TCP, 1800 seconds; for UDP, 60 seconds.
  - **protocol** ***number***—Specify the networking protocol name or number.
  - **rpc-program-number** ***number***—Specify the RPC or DCE value.
  - **source-port** ***port-number***—Specify a TCP or UDP source port number.
  - **uuid** ***hex-vale***—Specify the universal unique identifier (UUID) for objects.
- **uuid** ***hex-value***—Specify the UUID for objects.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- alg (Applications) on page 332
- application-protocol (Applications) on page 335
- destination-port (Applications) on page 347

Copyright © 2016, Juniper Networks, Inc.

application (Security Alarms)

Syntax

```
application {
  duration interval;
  size count;
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation policy]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for a number of policy violations to an application within a specified time period.

Options

- **duration *interval***—Indicate the duration of counters.

  Range: 1 through 3600 seconds.
  Default: 1 second.

- **size *count***—Indicate the number of applications for which policy violation checks can be done concurrently.

  Range: 1 through 10240.
  Default: 1024.

- **threshold *value***—Indicate the maximum number of application matches required to raise an alarm.

  Range: 1 through 4294967295.
  Default: 1000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application (Security Policies)

Syntax

```
application {
  [application];
  any;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP or remote procedure call (RPC) application or set of applications to be used as match criteria.

Options

application-name-or-set—Name of the predefined or custom application or application set used as match criteria.

any—Any predefined or custom applications or application sets.

NOTE: A custom application that does not use a well-known destination port for the application will not be included in the any option, and must be named explicitly.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application-protocol (Applications)

Syntax application-protocol (dns | ftp | http | https | ignore | ike-esp-nat | mgcp-ca | mgcp-ua | ms-rpc | pptp | q931 | ras | realaudio | rsh | rtsp | sccp | sip | smtp | sqlnet-v2 | ssh | sun-rpc | talk | telnet | tftp);

Hierarchy Level [edit applications **application** *application-name*]

Release Information Statement modified in Junos OS Release 8.5. The **ike-esp-nat** option introduced in Junos OS Release 10.2.

Description Identify the application protocol name. The following protocols are supported:

- **dns** —Domain Name Service
- **ftp** —File Transfer Protocol
- **http** —Hypertext Transfer Protocol
- **https** —Hypertext Transfer Protocol
- **ignore** —Ignore application type
- **ike-esp-nat** —IKE ESP NAT application protocol
- **mgcp-ca** —Media Gateway Control Protocol with Call Agent
- **mgcp-ua** —MGCP with User Agent
- **ms-rpc** —Microsoft RPC
- **pptp** —Point-to-Point Tunneling Protocol
- **q931** —ISDN connection control protocol (Q.931)
- **ras** —Remote Access Service
- **realaudio** —RealAudio
- **rsh** —UNIX remote shell services
- **rtsp**—Real-Time Streaming Protocol
- **sccp** —Skinny Client Control Protocol
- **sip**—Session Initiation Protocol
- **smtp** —Simple Mail Transfer Protocol
- **sqlnet-v2** —Oracle SQLNET v2
- **ssh** —Secure Shell Protocol
- **sun-rpc** —Sun Microsystems RPC
- **talk**—TALK program
- **telnet** —Telnet Protocol
- **tftp** —Trivial File Transfer Protocol

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

application-services (Security Policies)

Syntax

```
application-services {
  application-firewall {
    rule-set rule-set-name;
  }
  application-traffic-control {
    rule-set rule-set-name;
  }
  gprs-gtp-profile profile-name;
  gprs-sctp-profile profile-name;
  idp;
  redirect-wx | reverse-redirect-wx;
  ssl-proxy {
    profile-name profile-name;
  }
  uac-policy {
    captive-portal captive-portal;
  }
  utm-policy policy-name;
}
```

Hierarchy Level
[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information
Statement modified in Junos OS Release 11.1.

Description
Enable application services within a security policy.

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-tracking (Security Zones)

Syntax application-tracking;

Hierarchy Level [edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable application tracking support for the zone.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppTrack on page 574
- [edit security zones] Hierarchy Level on page 324

application-traffic-control (Application Services)

Syntax
```
application-traffic-control {
  rule-set rule-set-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enables AppQoS, application-aware quality of service, as specified in the rules of the specified rule set.

Options
- **rule-set *rule-set-name***—Name of the rule set that contains application-aware traffic control specification rules.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppQoS on page 589
- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

attach

Syntax

```
attach {
  zone zone-name;
}
```

Hierarchy Level [edit security address-book *book-name*]

Release Information Statement introduced in Release 11.2 of Junos OS.

Description Attach a security zone to an address book. You do not need to attach a security zone to the global address book. The global address book is available by default.

Options **zone *zone-name*—**Name of a security zone to be attached to the address book.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

audible (Security Alarms)

Syntax

```
audible {
  continuous;
}
```

Hierarchy Level [edit security alarms]

Release Information Statement modified in Junos OS Release 11.2.

Description Configure alarm to beep when a new security alarm is enabled.

Options **continuous**—Specify alarm to keep beeping until all security alarms are cleared.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

authentication (Security Alarms)

Syntax authentication *failures*;

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm when the device or switch detects a specified number of authentication failures (bad password attempts) before an alarm is raised.

This value must be equal to or less than the **tries-before-disconnect** setting at the **[edit system login retry-options]** hierarchy level; otherwise, the login session ends before the user reaches the alarmable threshold.

Default Multiple authentication failures do not cause an alarm to be raised.

Options *failures*—Number of authentication failures that causes an alarm to be raised.
Range: 1 through 10.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

authentication-source

Syntax

```
authentication-source {
    active-directory-authentication-table priority priority;
    firewall-authentication priority priority;
    local-authentication-table priority priority;
    unified-access-control priority priority;
}
```

Hierarchy Level [edit security user-identification]

Release Information Statement introduced in Junos OS Release 12.1. Statement updated in Junos OS Release 12.1-X45-D10. Support for the **active-directory-authentication-table priority** command statement added in Junos OS Release 12.1X47-D10.

Description Identifies one or more tables to be used as the source for user role information. Tables are searched in sequence based on lowest to highest priority.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- local-authentication-table on page 1265
- firewall-authentication (User Identification) on page 1252
- unified-access-control (Security) on page 1329
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

captive-portal (Services UAC Policy)

Syntax captive-portal *captive-portal-policy-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services uac-policy]

Release Information Statement introduced in Junos OS Release 10.2.

Description Create the captive portal policy in the UAC security policy. You use the captive portal policy to configure the captive portal feature on the Junos OS Enforcer. By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

count (Security Policies)

Syntax count;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 12.1X47-D10.

Description Enable a count, in bytes or kilobytes, of all network traffic the policy allows to pass through the device in both directions: the originating traffic from the client to the server (from the from-zone to the to-zone), and the return traffic from the server to the originating client.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- show security policies on page 802
- Security Policies Overview on page 1085

default-policy

Syntax default-policy (deny-all | permit-all);

Hierarchy Level [edit security policies]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the default security policy that defines the actions the device takes on a packet that does not match any user-defined policy.

Options **deny-all**—Deny all traffic. Packets are dropped. This is the default.

permit-all—Permit all traffic that does not match a policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

deny (Security Policies)

Syntax deny;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5.

Description Block the service at the firewall. The device drops the packets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085

description (Applications)

Syntax description *text*;

Hierarchy Level [edit applications application *application-name*],
[edit applications application-set *application-set-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for an application or an application set.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options *text*—Descriptive text about an application or an application set.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

description (Security Address Book)

Syntax description *text*;

Hierarchy Level [edit security address-book *book-name*]

Release Information Statement introduced in Release 12.1 of Junos OS.

Description Specify descriptive text for an address book.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options ***text***—Descriptive text about an address book.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

Copyright © 2016, Juniper Networks, Inc.

description (Security Policies)

Syntax description *description*;

Hierarchy Level
[edit security group-vpn member ike policy *policy-name*]
[edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike policy *policy-name*]
[edit security group-vpn server ipsec proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]
[edit security ike policy *policy-name*],
[edit security ike proposal *proposal-name*],
[edit security ipsec policy *policy-name*],
[edit security ipsec proposal *proposal-name*]
[edit security polices from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2. Support for the **security policies** hierarchy added in Junos OS Release 12.1.

Description Specify descriptive text for an IKE policy, an IPsec policy, an IKE proposal, an IPsec proposal, or a security policy.

Options ***description*** —Descriptive text about an IKE policy, an IPsec policy, an IKE proposal, an IPsec proposal, or a security policy.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security policies] Hierarchy Level on page 320
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

description (Security Zone)

Syntax description *text*;

Hierarchy Level [edit security zones functional-zone management]
[edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for a security zone.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options *text*—Descriptive text about a security zone.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [edit security zones] Hierarchy Level on page 324

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Policies)

Syntax

```
destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options

address—IP address (**any**, **any-ipv4**, **any-ipv6**), IP address set, or address book entry, or wildcard address (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Policies Flag)

Syntax

```
destination-address {
  drop-translated;
  drop-untranslated;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify whether the traffic permitted by the security policy is limited to packets where the destination IP address has been translated by means of a destination NAT rule or to packets where the destination IP address has not been translated.

On Juniper Networks security devices, destination NAT rules are processed before security policy lookup. Therefore, it is possible for a security policy to permit traffic from a source **S** to a destination **D** (where no destination NAT is performed) and also to permit traffic from the source **S** to the destination **d** (where **d** has been translated to **D**).

Options

- **drop-translated**—Drop packets with translated destination IP addresses. Traffic permitted by the security policy is limited to packets where the destination IP address has not been translated.
- **drop-untranslated**—Drop packets without translated destination IP addresses. Traffic permitted by the security policy is limited to packets where the destination IP address has been translated by means of a destination NAT rule.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

destination-address-excluded

Syntax destination-address-excluded;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Exclude the destination address(es) from the policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

destination-ip (Security Alarms)

Syntax

```
destination-ip {
  duration interval;
  size count;
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation policy]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for a number of policy violations to a destination network identifier within a specified time period.

Options

- **duration *interval***—Indicate the duration of counters.

  Range: 1 through 3600 seconds.
  Default: 1 second.

- **size *count***—Indicate the number of destination IP addresses for which policy violation checks can be done concurrently.

  Range: 1 through 10240.
  Default: 1024.

- **threshold *value***—Indicate the maximum number of destination IP address matches required to raise an alarm.

  Range: 1 through 4294967295.
  Default: 1000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

destination-port (Applications)

Syntax destination-port *port-identifier*;

Hierarchy Level [edit applications application *application-name*],
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify a TCP or UDP destination port number.

Options ***port-identifier*** —Range of ports. You can use a numeric value or one of the text synonyms listed in Table 14.

Copyright © 2016, Juniper Networks, Inc.

Table 96: Port Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760

Copyright © 2016, Juniper Networks, Inc.

Table 96: Port Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162

Copyright © 2016, Juniper Networks, Inc.

Table 96: Port Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

dns-proxy

Syntax

```
dns-proxy {
  cache hostname inet ip-address;
  default-domain domain-name {
    forwarders ip-address;
  }
  interface interface-name;
  propogate-setting (enable | disable);
  view view-name {
    domain domain-name {
      forward-only;
      forwarders ip-address;
    }
    match-clients subnet-address;
  }
}
```

Hierarchy Level [edit system services dns dns-proxy]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Configure the device as a DNS proxy server by enabling DNS proxy on a logical interface.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *DNS Proxy Overview*
- *Configuring the Device as a DNS Proxy*

Copyright © 2016, Juniper Networks, Inc.

dynamic-dns

Syntax

```
dynamic-dns {
  client hostname {
    agent agent-name;
    interface interface-name;
    password server-password;
    server server-name;
    username user-name;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Configure the device as a dynamic DNS server that maintains the list of the changed addresses and their associated domain names registered with it. The device updates these DDNS servers with this information periodically or whenever there is a change in IP addresses.

Options

- **client** —Specifies the hostname of the registered client.
- **agent**—Specifies the name of the dynamic DNS agent.
- **interface** —Specifies the interface whose IP address is mapped to the registered domain name.
- **password**—Specifies the password.
- **server**—Specifies the name of the dynamic DNS server that allows dynamic DNS clients to update the IP address changes associated with the registered hostname.
- **username**—Specifies the dynamic DNS username.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

exclude (Schedulers)

Syntax exclude;

Hierarchy Level
[edit schedulers scheduler *scheduler-name* daily],
[edit schedulers scheduler *scheduler-name* friday],
[edit schedulers scheduler *scheduler-name* monday],
[edit schedulers scheduler *scheduler-name* saturday],
[edit schedulers scheduler *scheduler-name* sunday],
[edit schedulers scheduler *scheduler-name* tuesday],
[edit schedulers scheduler *scheduler-name* thursday],
[edit schedulers scheduler *scheduler-name* wednesday]

Release Information Statement introduced in Junos OS Release 8.5.

Description Exclude a specified day from the schedule.

Use the **exclude** statement to exclude a day from a daily schedule created with the **daily** statement. You cannot use the **exclude** statement for a particular day unless it is in conjunction with the **daily** statement in a schedule.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policy Schedulers Overview on page 1123

Copyright © 2016, Juniper Networks, Inc.

firewall-authentication (Security Policies)

Syntax

```
firewall-authentication {
  pass-through {
    access-profile profile-name;
    client-match user-or-group-name;
    ssl-termination-profile profile-name;
    web-redirect;
    web-redirect-to-https;
  }
  user-firewall {
    access-profile profile-name;
    domain domain-name
    ssl-termination-profile profile-name;
  }
  web-authentication {
    client-match user-or-group-name;
  }
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement introduced in Junos OS Release 8.5.
Support added for the **user-firewall** option in Junos OS Release 12.1X45-D10.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Configure firewall authentication methods.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127

firewall-authentication (User Identification)

Syntax firewall-authentication priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Junos OS Release 12.1X45-D10. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description Enables the firewall authentication table as an authentication source. The priority of this table among other authentication tables establishes the search sequence used to identify user and role values.

Options priority *priority*—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the firewall authentication table is 150.

Setting the priority value of the firewall authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127

forward-only (DNS)

Syntax forward-only;

Hierarchy Level [edit system services dns dns-proxy view *view-name* domain *domain-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify that the server to forward only DNS queries. This configuration prevents the device from acquiring public IP addresses, in case the IP address specified in **forwarders** option is not reachable, by terminating the DNS query.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

from-zone (Security Policies)

Syntax

```
from-zone zone-name to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```

```
                redirect-wx | reverse-redirect-wx;
                ssl-proxy {
                  profile-name profile-name;
                }
                uac-policy {
                  captive-portal captive-portal;
                }
                utm-policy policy-name;
              }
              destination-address {
                drop-translated;
                drop-untranslated;
              }
              firewall-authentication {
                pass-through {
                  access-profile profile-name;
                  client-match user-or-group-name;
                  ssl-termination-profile profile-name;
                  web-redirect;
                  web-redirect-to-https;
                }
                user-firewall {
                  access-profile profile-name;
                  domain domain-name
                  ssl-termination-profile profile-name;
                }
                web-authentication {
                  client-match user-or-group-name;
                }
              }
              services-offload;
              tcp-options {
                initial-tcp-mss mss-value;
                reverse-tcp-mss mss-value;
                sequence-check-required;
                sequence-check-required;
                syn-check-required;
              }
              tunnel {
                ipsec-group-vpn group-vpn;
                ipsec-vpn vpn-name;
                pair-policy pair-policy;
              }
            }
            reject;
          }
        }
      }
```

Hierarchy Level [edit security policies]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **description** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Specify a source zone and destination zone to be associated with the security policy.

Options

- **from-zone** *zone-name*—Name of the source zone.
- **to-zone** *zone-name*—Name of the destination zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

from-zone (Security Policies Global)

Syntax

```
from-zone {
    [zone-name];
    any;
}
```

Hierarchy Level [edit security policies global policy *policy-name* match]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Identify a single source zone or multiple source zones to be used as a match criteria for a policy. You must configure specific zones or default to any zone, but you cannot have both in a global policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

functional-zone

Syntax

```
functional-zone {
  management {
    description text;
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
    interfaces interface-name  {
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
    }
    screen screen-name;
  }
}
```

Hierarchy Level [edit security zones]

Release Information Statement introduced in Junos OS Release 8.5. The **description** option added in Junos OS Release 12.1.

Description Configure a functional zone.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Functional Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

global (Security Policies)

Syntax

```
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      from-zone {
        [zone-name];
        any;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
        to-zone {
          [zone-name];
          any;
        }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
```

```
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
          redirect-wx | reverse-redirect-wx;
          ssl-proxy {
            profile-name profile-name;
          }
          uac-policy {
            captive-portal captive-portal;
          }
          utm-policy policy-name;
        }
        destination-address {
          drop-translated;
          drop-untranslated;
        }
        firewall-authentication {
          pass-through {
            access-profile profile-name;
            client-match user-or-group-name;
            web-redirect;
          }
          web-authentication {
            client-match user-or-group-name;
          }
        }
        services-offload;
        tcp-options {
          sequence-check-required;
          syn-check-required;
        }
      }
      reject;
    }
  }
}
```

Hierarchy Level [edit security policies]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated with **from-zone** and **to-zone** policy match options in Junos OS Release 12.1X47-D10.

Description Configure a global policy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Global Policy Overview on page 1115

host-inbound-traffic

Syntax

```
host-inbound-traffic {
  protocols protocol-name {
    except;
  }
  system-services service-name {
    except;
  }
}
```

Hierarchy Level

[edit security zones functional-zone management],
[edit security zones functional-zone management interfaces *interface-name*],
[edit security zones security-zone *zone-name*],
[edit security zones security-zone *zone-name* interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Control the type of traffic that can reach the device from interfaces bound to the zone.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058

Copyright © 2016, Juniper Networks, Inc.

icmp-code (Applications)

Syntax icmp-code *value*;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the Internet Control Message Protocol (ICMP) code value.

Options *value* — Specify the Internet Control Message Protocol (ICMP) code value such as **host-unreachable** or **host-unreachable-for-tos**.

Range: 0 through 255.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

icmp-type (Applications)

Syntax icmp-type *value*;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the ICMP packet type value.

Options *value*—ICMP type value, such as **echo** or **echo-reply**.

Range: 0 through 255.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

inactivity-timeout (Applications)

Syntax inactivity-timeout (*seconds* | never) ;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Inactivity timeout period, in seconds.

Options ***seconds***—Specify the amount of time the application is inactive before it times out in seconds.

Range: 4 through 129,600 seconds.

Default: For TCP, 1800 seconds; for UDP, 60 seconds.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

interfaces (Security Zones)

Syntax

```
interfaces interface-name {
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
}
```

Hierarchy Level [edit security zones functional-zone management],
[edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the set of interfaces that are part of the zone.

Options ***interface-name*** —Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Security Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

initial-tcp-mss

Syntax intial-tcp-mss *mss-value*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tcp-options]

Release Information Statement introduced in Junos OS Release 12.3X48-D20.

Description Configure the TCP maximum segment size (MSS) for packets that arrive at the ingress interface (initial direction), match a specific policy, and for which a session is created. The value you configure overrides the TCP MSS value in the incoming packet when the value in the packet is higher than the one you specify.

The **initial-tcp-mss** value per policy takes precedence over a global **tcp-mss** value (**all-tcp**, **ipsec-vpn**, **gre-in**, **gre-out**), if one is configured. However, when the **syn-flood-protection-mode syn-proxy** statement at the **[edit security flow]** hierarchy level is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overriden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure a TCP MSS value for the reverse session, use the **reverse-tcp-mss** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- reverse-tcp-mss on page 1293
- tcp-mss (Security Flow) on page 1885
- syn-flood-protection-mode on page 1008

Copyright © 2016, Juniper Networks, Inc.

ipsec-group-vpn (Security Policies)

Syntax ipsec-group-vpn *group-vpn*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tunnel]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify group VPN tunnel for the traffic configured by the scope policy on a group member.

Options ***group-vpn*** —Name of the group VPN tunnel configured on the group member.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Group VPNv2 Overview on page 6649
- Understanding Security Building Blocks for Security Devices on page 1045

ipsec-vpn (Security Policies)

Syntax ipsec-vpn *vpn-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tunnel]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define IPsec name for VPN.

Options ***vpn-name*** —Name of the IPsec.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

local-authentication-table

Syntax local-authentication-table priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Release 12.1 of Junos OS. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description An authentication table created on the SRX Series device using the **request security user-identification local-authentication-table add** command.

Options **priority *priority***—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the local authentication table is 100.

Setting the priority value of the local authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

log (Security Policies)

Syntax log (session-close | session-init);

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5.

Description Log traffic information for a specific policy. Traffic information is logged when a session begins (session-init) or closes (session-close).

Options **session-close**—Start logging traffic when the session ends.

session-init—Start logging traffic when the session begins.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

management (Security Zones)

Syntax

```
management {
  description text;
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
  interfaces interface-name  {
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
  }
  screen screen-name;
}
```

Hierarchy Level [edit security zones functional-zone]

Release Information Statement introduced in Junos OS Release 8.5.The **description** option added in Junos OS Release 12.1.

Description Specify the host for out-of-band management interfaces. You can set firewall options in this zone to protect the management interface from different types of attacks. Because this zone cannot be specified in policies, traffic entering from this zone can only be traffic originating from the device itself and cannot originate from any other zone.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

match (Security Policies)

Syntax

```
match {
  application {
    [application];
    any;
  }
  destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  source-identity {
    [role-name];
    any;
    authenticated-user;
    unauthenticated-user;
    unknown-user;
  }
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated with **source-identity** option in Junos OS Release 12.1.

Description Configure security policy match criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

match (Security Policies Global)

Syntax

```
match {
  application {
    [application];
    any;
  }
  destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  from-zone {
    [zone-name];
    any;
  }
  source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  source-identity {
    [role-name];
    any;
    authenticated-user;
    unauthenticated-user;
    unknown-user;
  }
  to-zone {
    [zone-name];
    any;
  }
}
```

Hierarchy Level [edit security policies global policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5. Statement updated with **source-identity** option in Junos OS Release 12.1. Statement updated with **to-zone** and **from-zone** options in Junos OS Release 12.1X47-D10.

Description Configure security global policy match criteria.

NOTE: We recommend that, for security reasons and to avoid spoofing traffic, when you create a multizone policy you use identical matching criteria (source address, destination address, application) and an identical action. For more information see "Global Policy Overview" on page 1115.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

no-policy-cold-synchronization

Syntax
no-policy-cold-synchronization

Hierarchy Level
[edit security idp sensor-configuration high-availability]

Release Information
Statement introduced in Junos OS Release 11.1.

Description
Disable policy cold synchronization functionality. This prevents the SRX Series devices from waiting for the IPS policy to be loaded on all service PICs, and there by forfeits IPS service protection during the ISSU (In-service software upgrade) operation/window.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

pair-policy

Syntax pair-policy *pair-policy* ;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tunnel]

Release Information Statement introduced in Junos OS Release 8.5.

Description Link the policy that you are configuring with another policy that references the same VPN tunnel so that both policies share one proxy ID and one security association (SA). Policy pairing is useful when you want to allow bidirectional traffic over a policy-based VPN that is using source or destination address translation with a dynamic IP address pool or destination address translation with a mapped IP (MIP) or dynamic IP (DIP) address pool.

Without policy pairing, the device derives a different proxy ID from the outbound and inbound policies. Two proxy IDs causes a problem for the remote peer with a single proxy ID for the VPN tunnel.

Pairing two policies solves the proxy ID problem for the remote peer and conserves SA resources. The single proxy ID is derived from the policy you configured last.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

pass-through

Syntax

```
pass-through {
  access-profile profile-name;
  client-match user-or-group-name;
  ssl-termination-profile profile-name;
  web-redirect;
  web-redirect-to-https;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication]

Release Information

Statement introduced in Junos OS Release 8.5. Support for **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10.

Description

Configure pass-through firewall user authentication. The user needs to use an FTP, Telnet, or HTTP client to access the IP address of the protected resource in another zone. Subsequent traffic from the user or host is allowed or denied based on the result of this authentication. Once authenticated, the firewall proxies the connection.

Options

- **access-profile** *profile-name* —(Optional) Specify the name of the access profile.
- **client-match** *user-or-group* —(Optional) Specify the name of the users or user groups in a profile who are allowed access by this policy. If you do not specify any users or user groups, any user who is successfully authenticated is allowed access.
- **ssl-termination-profile** *profile-name* —(Optional) Specify the SSL termination profile used for SSL offloading.
- **web-redirect**—(Optional) Enable redirecting an HTTP request to the device and redirecting the client system to a webpage for authentication. Including this statement allows users an easier authentication process because they need to know only the name or IP address of the resource they are trying to access.
- **web-redirect-to-https**—(Optional) Redirect unauthenticated HTTP requests to the internal HTTPS Web server of the device.

NOTE: If **web-redirect-to-https** is set, then you must specify the SSL termination profile used for SSL offloading.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

permit (Security Policies)

Syntax

```
permit {
  application-services {
    application-firewall {
      rule-set rule-set-name;
    }
    application-traffic-control {
      rule-set rule-set-name;
    }
    gprs-gtp-profile profile-name;
    gprs-sctp-profile profile-name;
    idp;
    redirect-wx | reverse-redirect-wx;
    ssl-proxy {
      profile-name profile-name;
    }
    uac-policy {
      captive-portal captive-portal;
    }
    utm-policy policy-name;
  }
  destination-address {
    drop-translated;
    drop-untranslated;
  }
  firewall-authentication {
    pass-through {
      access-profile profile-name;
      client-match user-or-group-name;
      ssl-termination-profile profile-name;
      web-redirect;
      web-redirect-to-https;
    }
    user-firewall {
      access-profile profile-name;
      domain domain-name
      ssl-termination-profile profile-name;
    }
    web-authentication {
      client-match user-or-group-name;
    }
  }
  services-offload;
  tcp-options {
    sequence-check-required;
    syn-check-required;
  }
  tunnel {
    ipsec-group-vpn group-vpn;
    ipsec-vpn vpn-name;
    pair-policy pair-policy;
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **tcp-options** added in Junos OS Release 10.4. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.

Description Specify the policy action to perform when packets match the defined criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

policies

Syntax

```
policies {
  default-policy (deny-all | permit-all);
  from-zone zone-name to-zone zone-name {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
        deny;
        log {
          session-close;
          session-init;
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
}
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
```

```
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      from-zone {
        [zone-name];
        any;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
      to-zone {
        [zone-name];
        any;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
          redirect-wx | reverse-redirect-wx;
          ssl-proxy {
            profile-name profile-name;
          }
          uac-policy {
```


Copyright © 2016, Juniper Networks, Inc.

```
                captive-portal captive-portal;
              }
              utm-policy policy-name;
            }
            destination-address {
              drop-translated;
              drop-untranslated;
            }
            firewall-authentication {
              pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
              }
              web-authentication {
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable) ;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Release Information

Statement introduced in Junos OS Release 8.5.
Support for the **services-offload** option added in Junos OS Release 11.4.
Support for the **source-identitiy** option added in Junos OS Release 12.1.
Support for the **description** option added in Junos OS Release 12.1.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.
Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.
Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description

Configure network security policies.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy (Security Alarms)

Syntax

```
policy {
  application {
    duration interval;
    size count;
    threshold value;
  }
  destination-ip {
    duration interval;
    size count;
    threshold value;
  }
  policy match {
    duration interval;
    size count;
    threshold value;
  }
  source-ip {
    duration interval;
    size count;
    threshold value;
  }
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for policy violation, based on source IP, destination IP, application, and policy rule.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy (Security Policies)

Syntax

```
policy policy-name {
  description description;
  match {
    application {
      [application];
      any;
    }
    destination-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
```


Copyright © 2016, Juniper Networks, Inc.

```
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              web-redirect;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. The **services-offload** option added in Junos OS Release 11.4. Statement updated with the **source-identity** option and the **description** option added in Junos OS Release 12.1. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Define a security policy.

Copyright © 2016, Juniper Networks, Inc.

Options ***policy-name***—Name of the security policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320
- *Building Blocks Feature Guide for Security Devices*

policy-match

Syntax
```
policy match {
  duration interval;
  size count;
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation policy]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for a policy rule or a group of rule violations within a specified time period.

Options
- **duration *interval***—Indicate the duration of counters.

  Range: 1 through 3600 seconds.
  Default: 1 second.

- **size *count***—Indicate the number of policies for which policy violation checks can be done concurrently.

  Range: 1 through 10240.
  Default: 1024.

- **threshold *value***—Indicate the number of policies for which policy violation checks can be done concurrently.

  Range: 1 through 4294967295.
  Default: 1000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy-rematch

Syntax policy-rematch <extensive>;

Hierarchy Level [edit security policies]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **extensive** option added in Junos OS Release 15.1X49D20.

Description Enable the device to reevaluate an active session when its associated security policy is changed, renamed, deactivated, or deleted. The session remains open if it matches any policy that allows the session.

The policy rematch feature is disabled by default, so that modified policies do not affect active sessions, and an active session is closed if its associated policy is renamed, deactivated, or deleted.

Options **extensive**—When a security policy associated with an active session is changed, renamed, deactivated, or deleted, the session remains active if it matches another policy that allows the session. If a match occurs, the new policy is applied to the session, along with the new scheduler (if any), but the session retains the timeout value of the old policy. If **policy-rematch** is specified without **extensive**, a rematch is attempted only for changed policies, and a session is closed if it no longer matches the modified policy.

NOTE: The extensive option does not apply to ALG data sessions or to policies that specify a source-identity, application-services, destination-address (drop-untranslated or drop-translated), firewall-authentication, or a tunnel.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy-stats

Syntax

```
policy-stats {
   system-wide (disable | enable) ;
}
```

Hierarchy Level [edit security policies]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Configure systemwide policies statistics. The systemwide policies statistics are disabled by default.

Options **disable**—Disable systemwide policy statistics.

enable—Enable systemwide policy statistics.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- show security policies on page 802

Copyright © 2016, Juniper Networks, Inc.

potential-violation

Syntax

```
potential-violation {
  authentication failures;
  cryptographic-self-test;
  decryption-failures {
    threshold value;
  }
  encryption-failures {
    threshold value;
  }
  idp;
  ike-phase1-failures {
    threshold value;
  }
  ike-phase2-failures {
    threshold value;
  }
  key-generation-self-test;
  non-cryptographic-self-test;
  policy {
    application {
      duration interval;
      size count;
      threshold value;
    }
    destination-ip {
      duration interval;
      size count;
      threshold value;
    }
    policy match {
      duration interval;
      size count;
      threshold value;
    }
    source-ip {
      duration interval;
      size count;
      threshold value;
    }
  }
  replay-attacks {
    threshold value;
  }
  security-log-percent-full percentage;
}
```

Hierarchy Level [edit security alarms]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for potential violation.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

protocol (Applications)

Syntax protocol *number*;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the networking protocol name or number.

Options ***protocol-name***— Networking protocol name. The following text values are supported. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

- **ah**—IP Security Authentication Header
- **egp**—Exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload
- **gre**—Generic routing encapsulation
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP over IP
- **node**—Clear each session that uses the specified IP protocol on a specific node.
- **ospf**—Open Shortest Path First
- **pim**—Protocol Independent Multicast
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Transmission Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

NOTE: Internet Protocol version 6 (IPv6) is not supported as a network protocol in application definitions

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

Copyright © 2016, Juniper Networks, Inc.

protocols (Security Zones Host Inbound Traffic)

Syntax

```
protocols {
    (protocol-name | all <protocol-name except>);
}
```

Hierarchy Level [edit security zones security-zone *zone-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of protocol traffic that can reach the device for all interfaces in a zone. You can do this in one of several ways:

- You can enable traffic from each protocol individually.
- You can enable traffic from all protocols.
- You can enable traffic from all but some protocols.

Options *protocol-name*—Protocol for which traffic is allowed. The following protocols are supported:

- **all**—Enable traffic from all possible protocols available. Use the *except* option to disallow specific protocols.
- **bfd**—Enable incoming Bidirectional Forwarding Detection (BFD) protocol traffic.
- **bgp**—Enable incoming BGP traffic.
- **dvmrp**—Enable incoming Distance Vector Multicast Routing Protocol (DVMRP) traffic.
- **igmp**—Enable incoming Internet Group Management Protocol (IGMP) traffic.
- **ldp**—Enable incoming Label Distribution Protocol (LDP) traffic (UDP and TCP port 646).
- **msdp**—Enable incoming Multicast Source Discovery Protocol (MSDP) traffic.
- **nhrp**—Enable incoming Next Hop Resolution Protocol (NHRP) traffic.
- **ospf**—Enable incoming OSPF traffic.
- **ospf3**—Enable incoming OSPF version 3 traffic.
- **pgm**—Enable incoming Pragmatic General Multicast (PGM) protocol traffic (IP protocol number 113).
- **pim**—Enable incoming Protocol Independent Multicast (PIM) traffic.
- **rip**—Enable incoming RIP traffic.
- **ripng**—Enable incoming RIP next generation traffic.
- **router-discovery**—Enable incoming router discovery traffic.
- **rsvp**—Enable incoming Resource Reservation Protocol (RSVP) traffic (IP protocol number 46).

Copyright © 2016, Juniper Networks, Inc.

- **sap**— Enable incoming Session Announcement Protocol (SAP) traffic. SAP always listens on **224.2.127.254:9875**. New addresses and ports can be added dynamically. This information must be propagated to the Packet Forwarding Engine (PFE).
- **vrrp**—Enable incoming Virtual Router Redundancy Protocol (VRRP) traffic.

except—(Optional) Disable specific incoming protocol traffic, but only when the *all* option has been defined . For example, to enable all but BGP and VRRP protocol traffic:

```
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust host-inbound-traffic protocols bgp except
set security zones security-zone trust host-inbound-traffic protocols vrrp except
```

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

protocols (Security Zones Interfaces)

Syntax

```
protocols protocol-name {
    except;
}
```

Hierarchy Level [edit security zones security-zone *zone-name* interfaces *interface-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of routing protocol traffic that can reach the device on a per-interface basis.

Options

- ***protocol-name*** —Protocol for which traffic is allowed. The following protocols are supported:
  - **all**—Enable traffic from all possible protocols available.
  - **bfd**—Enable incoming Bidirectional Forwarding Detection (BFD) Protocol traffic.
  - **bgp**—Enable incoming BGP traffic.
  - **dvmrp**—Enable incoming Distance Vector Multicast Routing Protocol (DVMRP) traffic.
  - **igmp**—Enable incoming Internet Group Management Protocol (IGMP) traffic.
  - **ldp**—Enable incoming Label Distribution Protocol (LDP) traffic (UDP and TCP port 646).
  - **msdp**—Enable incoming Multicast Source Discovery Protocol (MSDP) traffic.
  - **nhrp**—Enable incoming Next Hop Resolution Protocol (NHRP) traffic.
  - **ospf**—Enable incoming OSPF traffic.
  - **ospf3**—Enable incoming OSPF version 3 traffic.
  - **pgm**—Enable incoming Pragmatic General Multicast (PGM) protocol traffic (IP protocol number 113).
  - **pim**—Enable incoming Protocol Independent Multicast (PIM) traffic.
  - **rip**—Enable incoming RIP traffic.
  - **ripng**—Enable incoming RIP next generation traffic.
  - **router-discovery**—Enable incoming router discovery traffic.
  - **rsvp**—Enable incoming Resource Reolution Protocol (RSVP) traffic (IP protocol number 46).
  - **sap**— Enable incoming Session Announcement Protocol (SAP) traffic. SAP always listens on **224.2.127.254:9875**.
  - **vrrp**—Enable incoming Virtual Router Redundancy Protocol (VRRP) traffic.

except—(Optional) except can only be used if all has been defined.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Understanding Functional Zones on page 1050

range-address

Syntax
range-address *lower-limit* to *upper-limit*;

Hierarchy Level
[edit security address-book address *address-name*]

Release Information
Statement introduced in Release 12.1 of Junos OS.

Description
Configure the address range for an address book.

Options
- **lower-limit**—Lower limit of an address range.
- **upper-limit**—Upper limit of an address range.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

redirect-wx (Application Services)

Syntax
redirect-wx;

Hierarchy Level
[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Define the acceleration zone security policy for WX redirection of packets to the WXC Integrated Service Module (ISM 200) for WAN acceleration. During the redirection process, the direction of the WX packet and its type determine further processing of the packet.

Specify the WX redirection needed for the packets that arrive from the LAN. Use the **reverse-redirect-wx** statement to specify the WX redirection needed for the reverse flow of the packets that arrive from the WAN.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

reject (Security)

Syntax reject;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5.

Description Block the service at the firewall. The device drops the packet and sends a TCP reset (RST) segment to the source host for TCP traffic and an ICMP "destination unreachable, port unreachable" message (type 3, code 3) for UDP traffic. For types of traffic other than TCP and UDP, the device drops the packet without notifying the source host, which is also what occurs when the action is **deny**.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

reverse-tcp-mss

Syntax reverse-tcp-mss *mss-value*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tcp-options]

Release Information Statement introduced in Junos OS Release 12.3X48-D20.

Description Configure the TCP maximum segment size (MSS) for packets that match a specific policy and travel in the reverse direction of a session. The value you configure replaces the TCP MSS value when the value in the packet is higher than the one you specify.

The **reverse-tcp-mss** value per policy takes precedence over a global **tcp-mss** value (**all-tcp**, **ipsec-vpn**, **gre-in**, **gre-out**), if one is configured. However, when the **syn-flood-protection-mode syn-proxy** statement at the **[edit security flow] hierarchy level** is used to enable SYN proxy defenses against SYN attacks, the TCP MSS value is not overridden.

Because each policy has two directions, you can configure a value for both directions or for just one direction. To configure the TCP MSS value for the initial session, use the **initial-tcp-mss** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- initial-tcp-mss on page 1263
- tcp-mss (Security Flow) on page 1885
- syn-flood-protection-mode on page 1008

Copyright © 2016, Juniper Networks, Inc.

rpc-program-number (Applications)

Syntax rpc-program-number *number*;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the remote procedure call (RPC) or Distributed Computing Environment (DCE) value.

Options ***number***—RPC or DCE program value.

Range: 0 through 65,535

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Policy Application Sets Overview on page 1172

 Copyright © 2016, Juniper Networks, Inc.

scheduler (Security Policies)

Syntax

```
scheduler scheduler-name {
  daily {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  description text;
  friday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  monday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  saturday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  start-date date-time stop-date date-time;
  sunday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  thursday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  tuesday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
  wednesday {
    (all-day | exclude | start-time hh:mm stop-time hh:mm);
  }
}
```

Hierarchy Level [edit schedulers]

Release Information Statement introduced in Junos OS Release 8.5. The **description** option added in Junos OS Release 12.1.

Description Create or modify a scheduler that defines when security policies are in effect.

You configure a scheduler to start at a specific date and time or start on a recurrent basis.

Options ***scheduler-name*** —Name of the scheduler. The scheduler name must consist of 1 to 63 characters that can be letters, numbers, dashes, and underscores and can begin with a number or letter.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

scheduler-name

Syntax scheduler-name *scheduler-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the schedule (as defined by the **scheduler** ***scheduler-name*** statement) for which the policy is in effect.

Options ***scheduler-name*** —Name of the scheduler.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

schedulers (Security Policies)

Syntax schedulers { ... }

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure schedules for security policies that allow you to control network traffic flow and enforce network security.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

screen (Security Zones)

Syntax screen *screen-name*;

Hierarchy Level [edit security zones functional-zone management],
[edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a security screen for a security zone.

Options ***screen-name*** —Name of the screen.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

secure-domains

Syntax secure-domains [*domain-name*];

Hierarchy Level [edit system services dns dnssec]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure secure domains in the DNS server. The server accepts only signed responses for this domain. For unsigned responses, the server returns SERVFAIL error to the client.

Options **domain-name**—Name of the domain.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

secure-neighbor-discovery

Syntax

```
secure-neighbor-discovery {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Provide support for protecting Secure Neighbor Discovery Protocol (SEND) messages.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the Secure Neighbor Discovery (SEND) protocol process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

security-zone

Syntax

```
security-zone zone-name {
  address-book {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
  }
  application-tracking;
  description text;
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
  interfaces interface-name {
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
  }
  screen screen-name;
  tcp-rst;
}
```

Hierarchy Level [edit security zones]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description Define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Copyright © 2016, Juniper Networks, Inc.

Options *zone-name* —Name of the security zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security zones] Hierarchy Level on page 324
- Security Zones and Interfaces Overview on page 1049
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

sequence-check-required

Syntax sequence-check-required;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tcp-options]

Release Information Statement introduced in Junos OS Release 10.4.

Description Enable sequence check per policy. The sequence-check-required value overrides the global value no-sequence-check.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

services-offload (Security)

Syntax services-offload;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable services offloading within a security policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

session-close

Syntax session-close;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then log]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable traffic to which the policy applies to be logged at the end of a session.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

session-init

Syntax session-init;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then log]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable traffic to which the policy applies to be logged at the beginning of a session.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

simple-mail-client-service

Syntax

```
simple-mail-client-service {
    command binary-file-path;
    disable;
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the SMTP client process.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the SMTP client process.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-address (Security Policies)

Syntax

```
source-address {
  [address];
  any;
  any-ipv4;
  any-ipv6;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options ***address***—IP addresses, address sets, or wildcard addresses (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

source-address-excluded

Syntax

```
source-address-excluded;
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Exclude the source address(es) from the policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

source-identity

Syntax

```
source-identity {
  [user-or-role-name];
  any;
  authenticated-user;
  unauthenticated-user;
  unknown-user;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 12.1. Statement updated in Junos OS Release 12.1X44-D10.

Description

Identifies users and roles to be used as match criteria for a policy. If a value other than **any** is specified as match criteria for a policy within a zone pair, the traffic is matched to table entries to retrieve associated user and roles before policy lookup occurs. Users and roles are retrieved from the local authentication table or from a UIT pushed to the SRX Series device from an access control service when a user is authenticated.

The following entries specify the source identities that match a policy.

user-or-role-name—A list of specific users and roles.

any—Any user or role, as well as the keywords authenticated-user, unauthenticated-user, and unknown-user.

authenticated-user—All users and roles that have been authenticated.

unauthenticated-user—Any user or role that does not have an IP-address mapped to authentication sources and the authentication source is up and running.

unknown-user—Any user or role that does not have an IP address mapped to authentication sources, because the authentication source is disconnected from the SRX Series device. In this case, users are unable to be authenticated due to an authentication server disconnection, such as a power outage.

Unknown-user must be configured for non-domain users to be able to authenticate and log in.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

source-ip (Security Alarms)

Syntax

```
source-ip {
  duration interval;
  size count;
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation policy]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for a number of policy violations by a source network identifier within a specified time period.

Options

- **duration** ***interval***—Indicate the duration of counters.

  Range: 1 through 3600 seconds.
  Default: 1 second.

- **size** ***count***—Indicate the number of source IP addresses for which policy violation checks can be done concurrently.

  Range: 1 through 10240.
  Default: 1024.

- **threshold** ***value***—Indicate the maximum number of source IP address matches required to raise an alarm.

  Range: 1 through 4294967295.
  Default: 1000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

source-port (Applications)

Syntax source-port *port-number*;

Hierarchy Level [edit applications application *application-name*]
[edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the source port identifier.

Options ***port-number***—Identifier for the port. You can use a numeric value or one of the text synonyms listed in **destination-port (Applications)**.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Policy Application Sets Overview on page 1172

ssl-proxy (Application Services)

Syntax
```
ssl-proxy {
  profile-name profile-name
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 12.1.

Description Enable SSL proxy and identify the name of the SSL proxy profile to be used.

Options ***profile-name***—SSL proxy profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

ssl-termination-profile

Syntax ssl-termination-profile *profile-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the SSL termination profile used for SSL offloading.

Options ***profile-name***—Specify the name of the SSL termination profile used to the SSL offload.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

start-date

Syntax start-date *date-time*;

Hierarchy Level [edit schedulers scheduler *scheduler-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the time, day, month, and year that the schedule starts.

Specifying the year is optional. If no year is specified, the schedule applies to the current year and all subsequent years. If the year is specified in either the **start-date** or **stop-date** statement, that year is used for both statements.

Options ***date-time*** —Use the format **[** ***yyyy*** **-]** ***mm*** **-** ***dd*** **.** ***hh*** **.** ***mm*** to specify the year, month, day, hour, and minutes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

start-time (Schedulers)

Syntax start-time *hh:mm*;

Hierarchy Level
[edit schedulers scheduler *scheduler-name* daily],
[edit schedulers scheduler *scheduler-name* friday],
[edit schedulers scheduler *scheduler-name* monday],
[edit schedulers scheduler *scheduler-name* saturday],
[edit schedulers scheduler *scheduler-name* sunday],
[edit schedulers scheduler *scheduler-name* tuesday],
[edit schedulers scheduler *scheduler-name* thursday],
[edit schedulers scheduler *scheduler-name* wednesday]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the time that a schedule starts for a specified day.

If you specify a starting time for a daily schedule with the **daily** statement and also include the **friday**, **monday**, **saturday**, **sunday**, **tuesday**, **thursday**, and **wednesday** statements in the schedule, the starting time specified for a specific day (for example, Friday using the **friday** statement) overrides the starting time set with the **daily** statement.

Options ***time*** —Use the 24-hour format (***hh:mm:ss***) to specify the hours, minutes, and seconds.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

stop-date

Syntax stop-date *date-time*;

Hierarchy Level [edit schedulers scheduler *scheduler-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the time, day, month, and year that the schedule ends.

Specifying the year is optional. If no year is specified, the schedule applies to the current year and all subsequent years. If the year is specified in either the **start-date** or **stop-date** statement, that year is used for both statements.

Options ***date-time*** —Use the format **[** ***yyyy*** **-]** ***mm*** **-** ***dd*** **.** ***hh*** **.** ***mm*** to specify the year, month, day, hour, and minutes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

stop-time

Syntax stop-time *hh:mm*;

Hierarchy Level [edit schedulers scheduler *scheduler-name* daily],
[edit schedulers scheduler *scheduler-name* friday],
[edit schedulers scheduler *scheduler-name* monday],
[edit schedulers scheduler *scheduler-name* saturday],
[edit schedulers scheduler *scheduler-name* sunday],
[edit schedulers scheduler *scheduler-name* tuesday],
[edit schedulers scheduler *scheduler-name* thursday],
[edit schedulers scheduler *scheduler-name* wednesday]

Release Information Statement introduced in Junos OS Release.

Description Specify the time that a schedule stops for a specified day.

If you specify a stop time for a daily schedule with the **daily** statement and also include the the **friday**, **monday**, **saturday**, **sunday**, **tuesday**, **thursday**, and **wednesday** statements in the schedule, the stop time specified for a specific day (for example, Friday using the **friday** statement) overrides the stop time set with the **daily** statement.

Options ***time*** —Use the 24-hour format (***hh:mm:ss***) to specify the hours, minutes, and seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

syn-check-required

Syntax syn-check-required;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit tcp-options]

Release Information Statement introduced in Junos OS Release 10.4.

Description Enable sync check per policy. The syn-check-required value overrides the global value no-syn-check.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849

Copyright © 2016, Juniper Networks, Inc.

system-services (Security Zones Host Inbound Traffic)

Syntax

```
system-services {
  (service-name | all <service-name except>);
}
```

Hierarchy Level [edit security zones security-zone *zone-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of incoming system service traffic that can reach the device for all interfaces in a zone. You can do this in one of several ways:

- You can enable traffic from each system service individually.
- You can enable traffic from all system services.
- You can enable traffic from all but some system services.

Options

- ***service-name*** —System-service for which traffic is allowed. The following system services are supported:
  - **all**—Enable traffic from the defined system services available on the Routing Engine (RE). Use the *except* option to disallow specific system services.
  - **any-service**—Enable all system services on entire port range including the system services that are not defined.
  - **bootp**—Enable traffic destined to BOOTP and DHCP relay agents.
  - **dhcp**—Enable incoming DHCP requests.
  - **dhcpv6**—Enable incoming DHCP requests for IPv6.
  - **dns**—Enable incoming DNS services.
  - **finger**—Enable incoming finger traffic.
  - **ftp**—Enable incoming FTP traffic.
  - **http**—Enable incoming J-Web or clear-text Web authentication traffic.
  - **https**—Enable incoming J-Web or Web authentication traffic over Secure Sockets Layer (SSL).
  - **ident-reset**—Enable the access that has been blocked by an unacknowledged identification request.
  - **ike**—Enable Internet Key Exchange traffic.
  - **lsping**—Enable label switched path ping service.
  - **netconf**—Enable incoming NETCONF service.
  - **ntp**—Enable incoming Network Time Protocol (NTP) traffic.
  - **ping**—Allow the device to respond to ICMP echo requests.
  - **r2cp**—Enable incoming Radio Router Control Protocol traffic.

Copyright © 2016, Juniper Networks, Inc.

- **reverse-ssh**—Reverse SSH traffic.
- **reverse-telnet**—Reverse Telnet traffic.
- **rlogin**—Enable incoming **rlogin** (remote login) traffic.
- **rpm**—Enable incoming Real-time performance monitoring (RPM) traffic.
- **rsh**—Enable incoming Remote Shell (**rsh**) traffic.
- **snmp**—Enable incoming SNMP traffic (UDP port 161).
- **snmp-trap**—Enable incoming SNMP traps (UDP port 162).
- **ssh**—Enable incoming SSH traffic.
- **telnet**—Enable incoming Telnet traffic.
- **tftp**—Enable TFTP services.
- **traceroute**—Enable incoming traceroute traffic (UDP port 33434).
- **xnm-clear-text**—Enable incoming Junos XML protocol traffic for all specified interfaces.
- **xnm-ssl**— Enable incoming Junos XML protocol-over-SSL traffic for all specified interfaces.

- **except**—(Optional) Enable specific incoming system service traffic but only when the *all* option has been defined . For example, to enable all but FTP and HTTP system service traffic:

  ```
  set security zones security-zone trust host-inbound-traffic system-services all
  set security zones security-zone trust host-inbound-traffic system-services ftp except
  set security zones security-zone trust host-inbound-traffic system-services http except
  ```

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

system-services (Security Zones Interfaces)

Syntax

```
system-services service-name {
    except;
}
```

Hierarchy Level [edit security zones security-zone *zone-name* interfaces *interface-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of traffic that can reach the device on a particular interface.

Options

- *service-name* —Service for which traffic is allowed. The following services are supported:
  - **all**—Enable all possible system services available on the Routing Engine (RE).
  - **any-service**—Enable services on entire port range.
  - **bootp**—Enable traffic destined to BOOTP and DHCP relay agents.
  - **dhcp**—Enable incoming DHCP requests.
  - **dhcpv6**—Enable incoming DHCP requests for IPv6.
  - **dns**—Enable incoming DNS services.
  - **finger**—Enable incoming finger traffic.
  - **ftp**—Enable incoming FTP traffic.
  - **http**—Enable incoming J-Web or clear-text Web authentication traffic.
  - **https**—Enable incoming J-Web or Web authentication traffic over Secure Sockets Layer (SSL).
  - **ident-reset**—Enable the access that has been blocked by an unacknowledged identification request.
  - **ike**—Enable Internet Key Exchange traffic.
  - **netconf SSH**—Enable incoming NetScreen Security Manager (NSM) traffic over SSH.
  - **ntp**—Enable incoming Network Time Protocol (NTP) traffic.
  - **ping**—Allow the device to respond to ICMP echo requests.
  - **r2cp**—Enable incoming Radio Router Control Protocol traffic.
  - **reverse-ssh**—Reverse SSH traffic.
  - **reverse-telnet**—Reverse Telnet traffic.
  - **rlogin**—Enable incoming **rlogin** (remote login) traffic.
  - **rpm**—Enable incoming real-time performance monitoring (RPM) traffic.
  - **rsh**—Enable incoming Remote Shell (**rsh**) traffic.
  - **snmp**—Enable incoming SNMP traffic (UDP port 161).

Copyright © 2016, Juniper Networks, Inc.

- **snmp-trap**—Enable incoming SNMP traps (UDP port 162).
- **ssh**—Enable incoming SSH traffic.
- **telnet**—Enable incoming Telnet traffic.
- **tftp**—Enable TFTP services.
- **traceroute**—Enable incoming traceroute traffic (UDP port 33434).
- **xnm-clear-text**—Enable incoming Junos XML protocol traffic for all specified interfaces.
- **xnm-ssl**— Enable incoming Junos XML protocol-over-SSL traffic for all specified interfaces.

- **except**—(Optional) except can only be used if all has been defined.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

tcp-options (Security Policies)

Syntax

```
tcp-options {
  initial-tcp-mss mss-value;
  reverse-tcp-mss mss-value;
  sequence-check-required;
  syn-check-required;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information

Statement introduced in Junos OS Release 10.4. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Specify the TCP options for each policy. You can configure sync and sequence checks for each policy based on your requirements, and, because each policy has two directions, you can configure a TCP MSS value for both directions or for just one direction. To configure per-policy TCP options, you must turn off the respective global options. Otherwise, the commit check will fail.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

tcp-rst

Syntax tcp-rst;

Hierarchy Level [edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable the device to send a TCP segment with the RST (reset) flag set to 1 (one) in response to a TCP segment with any flag other than SYN set and that does not belong to an existing session.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding How to Identify Duplicate Sessions Using the TCP-Reset Parameter on page 1065
- Example: Configuring the TCP-Reset Parameter on page 1065

term (Applications)

Syntax
```
term term-name {
  alg application;
  destination-port port-identifier;
  icmp-code value;
  icmp-type value;
  icmp6-code value;
  icmp6-type value;
  inactivity-timeout (seconds | never);
  protocol number;
  rpc-program-number number;
  source-port port-number;
  uuid hex-value;
}
```

Hierarchy Level [edit applications application *application-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define individual application protocols.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Security Policy Applications Overview on page 1171

Copyright © 2016, Juniper Networks, Inc.

then (Security Policies)

Syntax

```
then {
  count {
    alarm {
      per-minute-threshold number;
      per-second-threshold number;
    }
  }
  deny;
  log {
    session-close;
    session-init;
  }
  permit {
    application-services {
      application-firewall {
        rule-set rule-set-name;
      }
      application-traffic-control {
        rule-set rule-set-name;
      }
      gprs-gtp-profile profile-name;
      gprs-sctp-profile profile-name;
      idp;
      redirect-wx | reverse-redirect-wx;
      ssl-proxy {
        profile-name profile-name;
      }
      uac-policy {
        captive-portal captive-portal;
      }
      utm-policy policy-name;
    }
    destination-address {
      drop-translated;
      drop-untranslated;
    }
    firewall-authentication {
      pass-through {
        access-profile profile-name;
        client-match user-or-group-name;
        ssl-termination-profile profile-name;
        web-redirect;
        web-redirect-to-https;
      }
      user-firewall {
        access-profile profile-name;
        domain domain-name
        ssl-termination-profile profile-name;
      }
      web-authentication {
        client-match user-or-group-name;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    services-offload;
    tcp-options {
      initial-tcp-mss mss-value;
      reverse-tcp-mss mss-value;
      sequence-check-required;
      syn-check-required;
    }
    tunnel {
      ipsec-group-vpn group-vpn;
      ipsec-vpn vpn-name;
      pair-policy pair-policy;
    }
  }
  reject;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Specify the policy action to be performed when packets match the defined criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

to-zone (Security Policies)

Syntax

```
to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```

Copyright © 2016, Juniper Networks, Inc.

```
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

Hierarchy Level [edit security policies from-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** and **junos-host** options added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10.

Description Specify a destination zone to be associated with the security policy.

Options

- **zone-name**—Name of the destination zone object.
- **junos-host**—Default security zone for self-traffic of the device.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

to-zone (Security Policies Global)

Syntax

```
to-zone {
  [zone-name];
  any;
}
```

Hierarchy Level
[edit security policies global policy *policy-name* match]

Release Information
Statement introduced in Junos OS Release 12.1X47-D10.

Description
Identify a single destination zone or multiple destination zones to be used as a match criteria for a policy. You must configure specific zones or default to any zone, but you cannot have both in a global policy.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Global Policy Overview on page 1115
- Example: Configuring a Global Policy with No Zone Restrictions on page 1117
- Example: Configuring a Global Policy with Multiple Zones on page 1119

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Policies)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security policies]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure policy tracing options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
  - **configuration**—Trace configuration events
  - **compilation**—Trace policy compilation events
  - **ipc**—Trace process inter communication events
  - **lookup**—Trace policy lookup events
  - **routing-socket**—Trace routing socket events
  - **rules**—Trace policy rules-related events
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security User Identification)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security user-identification]

Release Information Statement introduced in Release 12.1 of Junos OS.

Description Configure flow tracing options.

Options

- **file**—Configure the trace file options.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files** *number*—Maximum number of trace files. When a trace file named ***trace-file*** its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match** ***regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size** ***maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x*** **K** to specify KB, ***x*** **m** to specify MB, or ***x*** **g** to specify GB

    Range: 10 KB through 1 GB

    Default: 128 KB

Copyright © 2016, Juniper Networks, Inc.

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform.
  - **all**—Trace with all flags enabled
  - **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127

traceoptions (System Services DNS)

Syntax

```
traceoptions {
  category {
    category-type;
  }
  file;
}
```

Hierarchy Level [edit system services dns]

Release Information Statement introduced in Junos OS Release 10.2.

Description Defines tracing options for DNS services.

Options

category—Specifies the logging category. See Table 15 for the different logging categories and their descriptions.

file—Trace file information.

Copyright © 2016, Juniper Networks, Inc.

Table 97: Category Names

Category Name	Description
client	Processing of client requests
config	Configuration file parsing and processing
database	Messages relating to the databases
default	Categories for which there is no specific configuration
delegation-only	Delegation only
dispatch	Dispatching of incoming packets to the server
dnssec	DNSSEC and TSIG protocol processing
edns-disabled	Log query using plain DNS
general	General information
lame-servers	Lame servers
network	Network options
notify	NOTIFY protocol
queries	DNS query resolver
resolver	DNS resolution security
security	Approval and denial of requests
unmatched	Unable to determine the class for messages named
update	Dynamic updates
update-security	Approval and denial of update requests
xfer-in	Zone transfers that the server is receiving xfer-out
xfer-out	Zone transfers that the server is sending

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *DNS Overview*

Copyright © 2016, Juniper Networks, Inc.

tunnel (Security Policies)

Syntax

```
tunnel {
  ipsec-group-vpn group-vpn;
  ipsec-vpn vpn-name;
  pair-policy pair-policy;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement introduced in Junos OS Release 8.5. The **ipsec-group-vpn** option added in Junos OS Release 10.2.

Description Encapsulate outgoing IP packets and decapsulate incoming IP packets.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

uac-policy (Application Services)

Syntax

```
uac-policy {
  captive-portal captive-portal;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement modified in Junos OS Release 9.4.

Description Enable Unified Access Control (UAC) for the security policy. This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a UAC deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series UAC Appliance .

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138

Copyright © 2016, Juniper Networks, Inc.

unified-access-control (Security)

Syntax unified-access-control priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Release 12.1 of Junos OS. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description An authentication table pushed from a configured authentication device, such as the Junos Pulse Access Control Service.

Options priority *priority*—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each authentication table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the unified-access-control authentication table is 200.

Setting the priority value of the unified-access-control authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

user-firewall

Syntax

```
user-firewall {
  access-profile profile-name;
  domain domain-name
  ssl-termination-profile profile-name;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication]

Release Information

Statement introduced in Junos OS Release 12.1X45-D10. Support for the **domain** keyword added in Junos OS Release 12.1X47-D10.

Description

Configure user role firewall authentication, and map the source IP address to the username and its associated roles (groups). The mapped data is written to the firewall authentication table for later retrieval by the user role firewall. The user role firewall uses the username and role information to determine whether to permit or deny a user's session or traffic.

Options

access-profile *profile-name*—Specify the name of the access profile to be used for authentication.

domain *domain-name*—Specify the name of the domain where firewall authentication occurs in the event that the Windows Management Instrumentation client (WMIC) is not available to get IP-to-user mapping for the integrated user firewall feature. The maximum length is 65 bytes.

ssl-termination-profile *profile-name*—For HTTPS traffic, specify the name of the SSL termination profile used for SSL offloading.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Using Firewall Authentication as an Alternative to WMIC on page 5331
- Understanding User Role Firewalls on page 1127
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138

Copyright © 2016, Juniper Networks, Inc.

user-identification

Syntax

```
user-identification {
  authentication-source {
    firewall-authentication (disable | priority priority);
    local-authentication-table (disable | priority priority);
    unified-access-control (disable | priority priority);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Release 12.1 of Junos OS. Statement updated in Release 12.1X45-D10 of Junos OS.

Description Identifies one or more tables to be used as the source for user role information.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

utm-policy

Syntax

```
utm-policy policy-name {
  anti-spam {
    smtp-profile profile-name;
  }
  anti-virus {
    ftp {
      download-profile profile-name;
      upload-profile profile-name;
    }
    http-profile profile-name;
    imap-profile profile-name;
    pop3-profile profile-name;
    smtp-profile profile-name;
  }
  content-filtering {
    ftp {
      download-profile profile-name;
      upload-profile profile-name;
    }
    http-profile profile-name;
    imap-profile profile-name;
    pop3-profile profile-name;
    smtp-profile profile-name;
  }
  traffic-options {
    sessions-per-client {
      limit value;
      over-limit (block | log-and-permit);
    }
  }
  web-filtering {
    http-profile profile-name;
  }
}
```

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a UTM policy for antivirus, antispam, content-filtering, traffic-options, and web-filtering protocols and attach this policy to a security profile to implement it.

Options ***policy-name***—Specify name of the UTM policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Security Configuration Statement Hierarchy on page 603

uuid (Applications)

Syntax uuid *hex-value*;

Hierarchy Level [edit applications application *application-name* term *term-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the Universal Unique Identifier (UUID) for objects. DCE RPC services are mainly used by Microsoft applications. The ALG uses well-known TCP port 135 for port mapping services and uses the Universal Unique Identifier (UUID) instead of the program number to identify protocols. The main application-based DCE RPC is the Microsoft Exchange Protocol.

Support for stateful firewall and NAT services requires that you configure the DCE RPC port map ALG on TCP port 135. The DCE RPC ALG uses the TCP protocol with application-specific UUIDs.

Options **uuid *hex-vale*** —Specify the universal unique identifier (UUID) for objects.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Applications Configuration Statement Hierarchy on page 311

Copyright © 2016, Juniper Networks, Inc.

vrrp

Syntax

```
vrrp {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Virtual Router Redundancy Protocol (VRRP) process.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the VRRP process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

NOTE: On SRX300 and SRX320 devices, you cannot configure the same VRRP group ID on different interfaces of a single device

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

web-authentication

Syntax

```
web-authentication {
  client-match user-or-group-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication]

Release Information Statement introduced in Junos OS Release 8.5.
HTTPS for Web authentication is supported on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Specify that the policy allows access to users who have previously been authenticated by Web authentication. Web authentication must be enabled on one of the addresses on the interface to which the HTTP or HTTPS request is redirected.

Options **client-match** ***user-or-group*** —(Optional) Username or user group name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127

web-redirect

Syntax

```
web-redirect;
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 8.5.

Description Optionally, redirect HTTP requests to the device's internal webserver by sending a redirect HTTP response to the client system to reconnect to the webserver for user authentication. The interface on which the client's request arrived is the interface to which the request is redirected.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

zones

Syntax

```
zones {
  functional-zone {
    management {
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
      interfaces interface-name  {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
    }
  }
  security-zone zone-name {
    address-book {
      address address-name {
        ip-prefix {
          description text;
        }
        description text;
        dns-name domain-name {
          ipv4-only;
          ipv6-only;
        }
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
      }
      address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
      }
    }
    application-tracking;
    description text;
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
```


Copyright © 2016, Juniper Networks, Inc.

```
            except;
          }
        }
        interfaces interface-name {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
        tcp-rst;
      }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description A zone is a collection of interfaces for security purposes. All interfaces in a zone are equivalent from a security point of view. Configure the following zones:

- Functional zone—Special-purpose zone, such as a management zone that can host dedicated management interfaces.
- Security zone—Most common type of zone that is used as a building block in policies.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 64

Operational Commands

- clear security alarms
- clear security policies hit-count
- clear security policies statistics
- clear system services dns dns-proxy
- request security user-identification local-authorization-table add
- request security user-identification local-authentication-table delete
- show security alarms
- show security firewall-authentication users address
- show security firewall-authentication users auth-type
- show security flow session application
- show security match-policies
- show security policies
- show security policies hit-count
- show security policies unknown-source-identity
- show security shadow-policies logical-system
- show security user-identification local-authentication-table
- show security user-identification role-provision all
- show security user-identification source-identity-provision all
- show security user-identification user-provision all
- show security zones
- show security zones type
- show system services dns dns-proxy
- show system services dynamic-dns

Copyright © 2016, Juniper Networks, Inc.

clear security alarms

Syntax clear security alarms
<all>
<alarm-id *id-number*>
<alarm-type [*types*]>
<newer-than *YYYY-MM-DD.HH:MM:SS*>
<older-than *YYYY-MM-DD.HH:MM:SS*>
<process *process*>
<severity *severity*>

Release Information Command introduced in Junos OS Release 11.2.

Description Clear (acknowledge) the alarms that are active on the device.

Options **all**—(Optional) Clear all active alarms.

alarm-id *id-number*—(Optional) Clear the specified alarm.

alarm-type [*types*]—(Optional) Clear the specified alarm type or a set of types.
You can specify one or more of the following alarm types:

- **authentication**
- **cryptographic-self-test**
- **decryption-failures**
- **encryption-failures**
- **ike-phase1-failures**
- **ike-phase2-failures**
- **key-generation-self-test**
- **non-cryptographic-self-test**
- **policy**
- **replay-attacks**

newer-than *YYYY-MM-DD.HH:MM:SS*—(Optional) Clear active alarms that were raised after the specified date and time.

older-than *YYYY-MM-DD.HH:MM:SS*—(Optional) Clear active alarms that were raised before the specified date and time.

process *process*—(Optional) Clear active alarms that were raised by the specified system process.

severity *severity*—(Optional) Clear active alarms of the specified severity.

Copyright © 2016, Juniper Networks, Inc.

You can specify the following severity levels:

- alert
- crit
- debug
- emerg
- err
- info
- notice
- warning

Required Privilege Level security—To view this statement in the configuration.

Related Documentation
- show security alarms on page 1349
- Troubleshooting Security Policies on page 1162

List of Sample Output
clear security alarms all on page 1341
clear security alarms alarm-id <alarm-id> on page 1341
clear security alarms alarm-type authentication on page 1341
clear security alarms newer-than <time> on page 1342
show security alarms older-than <time> on page 1342
show security alarms process <process> on page 1342
show security alarms severity <severity> on page 1342

Output Fields This command produces no output, except a statement about the number of security alarms cleared.

Sample Output

clear security alarms all

```
[3 SECURITY ALARMS] user@router> clear security alarms all

3 security alarms cleared
```

clear security alarms alarm-id <alarm-id>

```
[3 SECURITY ALARMS] user@router> clear security alarms alarm-id 1

1 security alarm cleared
```

clear security alarms alarm-type authentication

```
[3 SECURITY ALARMS] user@router> clear security alarms alarm-type authentication

3 security alarms cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security alarms newer-than <time>

```
[3 SECURITY ALARMS] user@router> clear security alarms newer-than 2010-01-19.13:41:59

1 security alarm cleared
```

show security alarms older-than <time>

```
[3 SECURITY ALARMS] user@router> clear security alarms older-than 2010-01-19.13:41:59

2 security alarms cleared
```

show security alarms process <process>

```
[3 SECURITY ALARMS] user@router> clear security alarms process sshd

3 security alarms cleared
```

show security alarms severity <severity>

```
[3 SECURITY ALARMS] user@router> clear security alarms severity notice

3 security alarms cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security policies hit-count

Syntax

clear security policies hit-count
<from-zone *zone-name*>
<to-zone *zone-name*>

Release Information

Command introduced in Junos Release 12.1.

Description

Clear the hit-count values for security policies.

Options

- **from-zone** ***zone-name***—(Optional) Clear the number of hits for security policies associated with the named source zone.
- **to-zone** ***zone-name***—(Optional) Clear the number of hits for security policies associated with the named destination zone.

Required Privilege Level

clear

Related Documentation

- show security policies hit-count on page 1375
- Monitoring Policy Statistics on page 1162

Output Fields

This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security policies statistics

Syntax clear security policies statistics

Release Information Command introduced in Junos OS Release 8.5. Support for systemwide policies statistics added in Junos OS Release 12.1X46-D10.

Description Clear systemwide policies statistics and security policies statistics configured on the device.

Required Privilege Level clear

Related Documentation
- show security policies on page 802
- Monitoring Policy Statistics on page 1162

List of Sample Output clear security policies statistics on page 1344

Output Fields This command produces no output.

Sample Output

clear security policies statistics

```
user@host> clear security policies statistics
```


Copyright © 2016, Juniper Networks, Inc.

clear system services dns dns-proxy

Syntax clear system services dns dns-proxy

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Clear DNS proxy cache information.

Options

- **cache**—Clear DNS proxy cache information.
- **statistics**—Clear DNS proxy statistics.
- **none**—Clear all DNS proxy cache information.
- **hostname**—(Optional) Clear DNS proxy cache information from the specified host.

Required Privilege Level clear

Related Documentation

- show system services dns dns-proxy on page 1392
- show system services dynamic-dns on page 1395

List of Sample Output clear system services dns-proxy on page 1345

Output Fields When you enter this command no output is produced.

Sample Output

clear system services dns-proxy

```
user@host> clear system services dns dns-proxy cache
user@host> clear system services dns dns-proxy statistics
```

Copyright © 2016, Juniper Networks, Inc.

request security user-identification local-authorization-table add

Syntax request security user-identification local-authorization-table add user *user-name* ip-address *ip-address* roles [*role-name*]

Release Information Command introduced in Junos OS Release 12.1. Command updated in Junos OS Release 12.1X44-D10.

Description This command adds user and role information to the local authentication table. The table is used to retrieve user and role information for traffic from the specified IP address to enforce a user role firewall.

To add an entry, specify the user name, IP address, and up to 40 roles to be associated with this user. Subsequent commands for the same user and IP address aggregates any new roles to the existing list. An authentication entry can contain up to 200 roles.

NOTE: To change the user name of an entry or to remove or change entries in a role list, you must delete the existing entry and create a new one.

An IP address can be associated with only one user. If a second request is made to add a different user using the same IP address, the second authentication entry overwrites the existing entry.

Options user *user-name*—Specify the name of the user to be added to the table.

ip-address *ip-address*—Specify the IP address of the user. Either IPv4 or IPv6 addresses are supported.

roles [*role-name*]—(Optional) Specify the role or list of roles to be associated with the specified user. If the specified user and IP address already exist, any roles specified in the command are added to the existing role list.

Required Privilege Level maintenance

Related Documentation

- request security user-identification local-authentication-table delete on page 1348
- Understanding the User Identification Table on page 1130

List of Sample Output request security user-identification local-authentication-table add on page 1347

Output Fields When you enter this command, either an entry is added to the local authentication table, or the roles of an existing entry are aggregated with additional roles.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

request security user-identification local-authentication-table add

```
user@host> request security user-identification local-authentication-table add user user1
ip-address 192.0.2.1 roles role1
user@host> request security user-identification local-authentication-table add user user2
ip-address 203.0.113.2 roles [role2 role3]
user@host> request security user-identification local-authentication-table add user user2
ip-address 203.0.113.2 roles role1
user@host> show security user-identification local-authentication-table all
Total entries: 2
Source IP       Username      Roles
192.0.2.1       user1      role1
203.0.113.2     user2      role2, role3, role1
```

Copyright © 2016, Juniper Networks, Inc.

request security user-identification local-authentication-table delete

Syntax request security user-identification local-authentication-table delete *ip-address* | *user-name*

Release Information Command introduced in Junos OS Release 12.1.

Description This command removes an entry from the local authentication table. You can identify the entry by IP address or user-name. To change the user name of an entry or to remove or change entries in a role list, you must delete the existing entry and create a new one.

Options ***ip-address***—The IP address of the entry to be deleted.

user-name—The user name of the entry to be deleted. To change the user name of an entry, you must delete the old entry and create a new one.

Required Privilege Level maintenance

Related Documentation

- request security user-identification local-authorization-table add on page 1346
- Understanding the User Identification Table on page 1130

Output Fields The specified show command verifies the table content before and after an entry has been deleted from the local authentication table.

Sample Output

```
user@host> show security user-identification local-authentication-table all
Total entries: 2
    Ip-address: 192.0.2.1
    Username: user1
    Roles: role1

    Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1
```

```
user@host> request security user-identification local-authentication-table delete 203.0.113.2
user@host> show security user-identification local-authentication-table all
Total entries: 1
    Ip-address: 192.0.2.1
    Username: user1
    Roles: role1
```


Copyright © 2016, Juniper Networks, Inc.

show security alarms

Syntax show security alarms
<detail>
<alarm-id *id-number*>
<alarm-type [*types*]>
<newer-than *YYYY-MM-DD.HH:MM:SS*>
<older-than *YYYY-MM-DD.HH:MM:SS*>
<process *process*>
<severity *severity*>

Release Information Command introduced in Junos OS Release 11.2.

Description Display the alarms that are active on the device. Run this command when the CLI prompt indicates that a security alarm has been raised, as shown here:

```
[1 SECURITY ALARM] user@host#
```

Options **none**—Display all active alarms.

detail—(Optional) Display detailed output.

alarm-id *id-number*—(Optional) Display the specified alarm.

alarm-type [*types*]—(Optional) Display the specified alarm type or a set of types. You can specify one or more of the following alarm types:

- **authentication**
- **cryptographic-self-test**
- **decryption-failures**
- **encryption-failures**
- **ike-phase1-failures**
- **ike-phase2-failures**
- **key-generation-self-test**
- **non-cryptographic-self-test**
- **policy**
- **replay-attacks**

newer-than *YYYY-MM-DD.HH:MM:SS*—(Optional) Display active alarms that were raised after the specified date and time.

older-than *YYYY-MM-DD.HH:MM:SS*—(Optional) Display active alarms that were raised before the specified date and time.

process *process*—(Optional) Display active alarms that were raised by the specified system process.

Copyright © 2016, Juniper Networks, Inc.

severity *severity*—(Optional) Display active alarms of the specified severity.

You can specify the following severity levels:

- **alert**
- **crit**
- **debug**
- **emerg**
- **err**
- **info**
- **notice**
- **warning**

Required Privilege Level security—To view this statement in the configuration.

Related Documentation

- clear security alarms on page 1340
- Example: Generating a Security Alarm in Response to Policy Violations on page 1167

List of Sample Output

show security alarms on page 1351
show security alarms detail on page 1351
show security alarms alarm-id on page 1351
show security alarms alarm-type authentication on page 1351
show security alarms newer-than <time> on page 1352
show security alarms older-than <time> on page 1352
show security alarms process <process> on page 1352
show security alarms severity <severity> on page 1352

Output Fields Table 98 lists the output fields for the **show security alarms** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used.

Table 98: show security alarms

Field Name	Field Description	Level of Output
ID	Identification number of the alarm.	All levels
Alarm time	Date and time the alarm was raised..	All levels
Message	Information about the alarm, including the alarm type, username, IP address, and port number.	All levels
Process	System process (For example, login or sshd) and process identification number associated with the alarm.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 98: show security alarms *(continued)*

Field Name	Field Description	Level of Output
Severity	Severity level of the alarm.	**detail**

Sample Output

show security alarms

```
[3 SECURITY ALARMS] user@router> show security alarms

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms detail

```
[3 SECURITY ALARMS] user@router> show security alarms detail

Alarm ID   : 1
Alarm Type : authentication
Time       : 2010-01-19 13:41:36 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
 user 'user' reached from '10.17.0.1'
Process    : sshd (pid 1414)
Severity   : notice

Alarm ID   : 2
Alarm Type : authentication
Time       : 2010-01-19 13:41:52 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
 user 'user' reached from '10.17.0.1'
Process    : sshd (pid 1414)
Severity   : notice

Alarm ID   : 3
Alarm Type : authentication
Time       : 2010-01-19 13:42:13 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
 user 'user' reached from '10.17.0.1'
Process    : sshd (pid 1414)
Severity   : notice
```

show security alarms alarm-id

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-id 1

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms alarm-type authentication

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-type authentication
```

Copyright © 2016, Juniper Networks, Inc.

```
ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms newer-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms newer-than 2010-01-19.13:41:59

3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms older-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms older-than 2010-01-19.13:41:59

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms process <process>

```
[3 SECURITY ALARMS] user@router> show security alarms process sshd

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```

show security alarms severity <severity>

```
[3 SECURITY ALARMS] user@router> show security alarms severity notice

ID    Alarm time               Message
1     2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
2     2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
3     2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
 failures (1) for user 'user' reached from '10.17.0.1'
```


Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users address

Syntax **show security firewall-authentication users address** *ip-address*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display information about the users at the specified IP address that are currently authenticated.

Options

- **address** *ip-address*—IP address of the authentication source.
- none—Display all the firewall authentication information for users at this IP address.
- **node**—(Optional) For chassis cluster configurations, display user firewall authentication entries on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Understanding User Role Firewalls on page 1127

List of Sample Output

show security firewall-authentication users address 192.0.2.9 on page 1354
show security firewall-authentication users address 192.0.2.9 node local on page 1354
show security firewall-authentication users address 192.0.2.9 on page 1355

Output Fields Table 99 lists the output fields for the **show security firewall-authentication users address** command. Output fields are listed in the approximate order in which they appear.

Table 99: show security firewall-authentication users address Output Fields

Field Name	Field Description
Username	User ID.
Source IP	IP address of the authentication source.
Authentication state	Status of authentication (success or failure).
Authentication method	Path chosen for authentication.
Access time remaining	Duration for which the connection exists.

Copyright © 2016, Juniper Networks, Inc.

Table 99: show security firewall-authentication users address Output Fields *(continued)*

Field Name	Field Description
Lsys	The logical system where the traffic was received.
Source zone	User traffic received from the zone.
Destination zone	User traffic destined to the zone.
Policy index	Identification number of the policy.
Policy name	Name of the policy.
Access profile	Name of profile used for authentication.
Interface Name	Name of the interface.
Bytes sent by this user	Number of bytes sent by the user.
Bytes received by this user	Number of bytes received by the user.
Client-groups	Name of the client group.

Sample Output

show security firewall-authentication users address 192.0.2.9

```
user@host>show security firewall-authentication users address 192.0.2.9
Username: hello
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: Pass-through using Telnet
Access time remaining: 0
Source zone: z2
Destination zone: z1
Policy index: 5
Access profile: profile1
Interface Name: ge-0/0/2.0
Bytes sent by this user: 0
Bytes received by this user: 0
Client-groups: Sunnyvale Bangalore
```

Sample Output

show security firewall-authentication users address 192.0.2.9 node local

```
user@host> show security firewall-authentication users address 192.0.2.9 node local
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: Pass-through using Telnet
Age: 2
```


Copyright © 2016, Juniper Networks, Inc.

```
Access time remaining: 4
Source zone: z1
Destination zone: z2
Policy name: POL1
Access profile: p1
Interface Name: reth1.0
Bytes sent by this user: 614
Bytes received by this user: 1880
```

show security firewall-authentication users address 192.0.2.9

```
user@host> show security firewall-authentication users address 192.0.2.9
Username: hello
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: User-firewall
Age: 0
Access time remaining: 10
Lsys: root-logical-system
Source zone: N/A
Destination zone: N/A
Access profile: test
```

show security firewall-authentication users auth-type

Syntax show security firewall-authentication users auth-type [user-firewall | pass-through | web-authentication]

Release Information Command introduced in Junos OS Release 12.1X45-D10

Description Display statistics about the users authenticated by the selected method.

Options
- **user-firewall**—Lists all users authenticated for user firewall.
- **pass-through**—Lists all users authenticated by the pass-through method.
- **web-authentication**—Lists all users authenticated by the user authentication method.

Required Privilege Level view

Related Documentation
- Understanding User Role Firewalls on page 1127

List of Sample Output
show security firewall-authentication users auth-type user-firewall on page 1357
show security firewall-authentication users auth-type pass-through on page 1357
show security firewall-authentication users auth-type pass-through on page 1357

Output Fields Table 100 lists the output fields for the **show security firewall-authentication users auth-type** command. Output fields are listed in the approximate order in which they appear.

Table 100: show security firewall-authentication users auth-type Output Fields

Field Name	Field Description
Total users in table	Total number of users authenticated by this method.
Id	The ID assigned to the entry.
Source IP	The source IP address of the traffic.
Src zone	The source zone of the traffic.
Dst zone	The destination zone of the traffic.
Profile	The profile used to authenticate the used.
Age	The length of time since authentication.
Status	The status of the authentication.
User	The username associated with the traffic.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security firewall-authentication users auth-type user-firewall

```
user@host> show security firewall-authentication users auth-type user-firewall
User-firewall authentication data:
  Total users in table: 2
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
           1 10.208.16.1  N/A      N/A      test         0 Success  jasonliu
           2 10.208.16.6  N/A      N/A      test6        0 Failed   jason
```

show security firewall-authentication users auth-type pass-through

```
user@host> show security firewall-authentication users auth-type pass-through
Pass-through firewall authentication data:
  Total users in table: 1
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
           1 10.208.16.2  zone1    zone2    test2        0 Success  jasonliu2
```

show security firewall-authentication users auth-type pass-through

```
user@host> show security firewall-authentication users auth-type web-authentication
Web firewall authentication data:
  Total users in table: 1
          Id Source Ip   Src zone Dst zone Profile    Age Status   User
           1 10.208.16.3  N/A      N/A      test3        0 Success  jasonliu3
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session application

Syntax

show security flow session application
application-name **[brief | extensive | summary]**

Release Information

Command introduced in Junos OS Release 8.5. Filter and view options added in Junos OS Release 10.2.

Description

Display information about each session of the specified application type.

Options

- ***application-name***—Type of application about which to display sessions information. Possible values are:
  - dns—Domain Name System
  - ftp—File Transfer Protocol
  - ignore—Ignore application type
  - mgcp-ca—Media Gateway Control Protocol with Call Agent
  - mgcp-ua—MGCP with User Agent
  - pptp—Point-to-Point Tunneling Protocol
  - q931—ISDN connection control protocol
  - ras—Remote Access Server
  - realaudio—RealAudio
  - rsh—UNIX remote shell services
  - rtsp—Real-Time Streaming Protocol
  - sccp—Skinny Client Control Protocol
  - sip—Session Initiation Protocol
  - sqlnet-v2—Oracle SQLNET
  - talk—TALK program
  - tftp—Trivial File Transfer Protocol
- brief | extensive | summary—Display the specified level of output.

Required Privilege Level

view

Related Documentation

- clear security flow session application on page 422

List of Sample Output

show security flow session application telnet on page 1360
show security flow session application telnet brief on page 1360
show security flow session application telnet extensive on page 1360
show security flow session application telnet summary on page 1361

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 27 lists the output fields for the **show security flow session application** command. Output fields are listed in the approximate order in which they appear.

Table 101: show security flow session application Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 101: show security flow session application Output Fields *(continued)*

Field Name	Field Description
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session application telnet

```
root> show security flow session application telnet
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes:
610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes:
602
Total sessions: 1
```

show security flow session application telnet brief

```
root> show security flow session application telnet brief
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Policy name: default-policy/2, Timeout: 1796, Valid
  In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp, If: ge-0/0/2.0, Pkts: 10, Bytes:
610
  Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp, If: ge-0/0/1.0, Pkts: 9, Bytes:
602
Total sessions: 1
```

show security flow session application telnet extensive

```
root>  show security flow session application telnet extensive
Flow Sessions on FPC4 PIC1:
Total sessions: 0

Flow Sessions on FPC5 PIC0:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC5 PIC1:

Session ID: 210067547, Status: Normal
Flag: 0x40
Policy name: default-policy/2
Source NAT pool: Null, Application: junos-telnet/10
Maximum timeout: 1800, Current timeout: 1788
Session State: Valid
Start time: 670184, Duration: 33
   In: 203.0.113.2/32781 --> 192.0.2.5/23;tcp,
    Interface: ge-0/0/2.0,
    Session token: 0x180, Flag: 0x0x21
    Route: 0x60010, Gateway: 203.0.113.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 10, Bytes: 610
   Out: 192.0.2.5/23 --> 203.0.113.2/32781;tcp,
    Interface: ge-0/0/1.0,
    Session token: 0x1c0, Flag: 0x0x20
    Route: 0x70010, Gateway: 192.0.2.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 602
Total sessions: 1
```

show security flow session application telnet summary

```
root> show security flow session application telnet summary
Flow Sessions on FPC4 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC5 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC5 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security match-policies

Syntax

show security match-policies
destination-ip *ip-address*
destination-port *port-number*
from-zone *zone-name*
protocol *protocol-name* | *protocol-number*
result-count *number*
source-identity *role-name*
source-ip *ip-address*
source-port *port-number*
to-zone *zone-name*

Release Information

Command introduced in Junos OS Release 10.3. Command updated in Junos OS Release 10.4 . Command updated in Junos OS Release 12.1. Command updated to include optional from-zone and to-zone global match options in Junos OS Release 12.1X47-D10.

Description

The **show security match-policies** command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, then the **show security match-policies** command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The **result-count** option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are "shadowed" by the first and are never encountered by matching traffic.

NOTE: **The show security match-policies command is applicable only to security policies; IDP policies are not supported.**

Options

- **destination-ip** ***destination-ip***—Destination IP address of the traffic.
- **destination-port** ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535
- **from-zone** ***from-zone***—Name or ID of the source zone of the traffic.
- **protocol** ***protocol-name*** **|** ***protocol-number***—Protocol name or numeric value of the traffic.
  - **ah** or **51**
  - **egp** or **8**
  - **esp** or **50**
  - **gre** or **47**
  - **icmp** or **1**
  - **igmp** or **2**

Copyright © 2016, Juniper Networks, Inc.

  - **igp** or **9**
  - **ipip** or **94**
  - **ipv6** or **41**
  - **ospf** or **89**
  - **pgm** or **113**
  - **pim** or **103**
  - **rdp** or **27**
  - **rsvp** or **46**
  - **sctp** or **132**
  - **tcp** or **6**
  - **udp** or **17**
  - **vrrp** or **112**
- **result-count** ***number***—(Optional) The number of policy matches to display. Valid range is from 1 through 16. The default value is 1.
- **source-identity** ***role-name***—Source identity of the traffic determined by the user role.
- **source-ip** ***source-ip***—Source IP address of the traffic.
- **source-port** ***source-port***—Source port number of the traffic. Range is 1 through 65,535.
- **to-zone** ***to-zone***—Name or ID of the destination zone of the traffic.

Required Privilege Level view

Related Documentation
- clear security policies statistics on page 1344
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output
Example 1: show security match-policies on page 1365
Example 2: show security match policies ... result-count on page 1365
Example 3: show security match policies ... source-identity on page 1366

Output Fields Table 102 lists the output fields for the **show security match-policies** command. Output fields are listed in the approximate order in which they appear.

Table 102: show security match-policies Output Fields

Field Name	Field Description
Policy	Name of the applicable policy.

Copyright © 2016, Juniper Networks, Inc.

Table 102: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Action or **Action-type**	The action to be taken for traffic that matches the policy's match criteria. Actions include the following: • **permit** • **firewall-authentication** • **tunnel ipsec-vpn** *vpn-name* • **pair-policy** *pair-policy-name* • **source-nat pool** *pool-name* • **pool-set** *pool-set-name* • **interface** • **destination-nat** *name* • **deny** • **reject**
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	An internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
From zone	Name of the source zone.
To zone	Name of the destination zone.
Source addresses	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
Application	Name of a preconfigured or custom application, or **any** if no application is specified.
IP protocol	Numeric value for the IP protocol used by the application, such as 6 for TCP or 1 for ICMP.
ALG	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
Inactivity timeout	Elapsed time without activity after which the application is terminated.
Source-port range	Range of matching source ports defined in the policy.

Copyright © 2016, Juniper Networks, Inc.

Table 102: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Destination-port range	Range of matching destination ports defined in the policy.
Source identities	One or more user roles defined in the matching policy.

Sample Output

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1 to-zone z2 source-ip 10.10.10.1
destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 198.51.100.0/24
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.0/24
    d3: 192.0.2.1/32
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [21-21]
```

Example 2: show security match policies ... result-count

```
user@host> show security match-policies source-ip 10.10.10.1 destination-ip 192.0.2.5 source_port
1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/32
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]
```

Copyright © 2016, Juniper Networks, Inc.

Example 3: show security match policies ... source-identity

```
user@host> show security match-policies from-zone untrust to-zone trust source-ip 10.10.10.1
destination-ip 192.0.2.1 destination_port 21 protocol 6 source-port 1234 source-identity role1
Policy: p1, action-type: permit, State: enabled, Index: 40
      Policy Type: Configured
      Sequence number: 1
      From zone: untrust, To zone: trust
      Source addresses:
        a1: 10.0.0.0/8
      Destination addresses:
        d1: 192.0.2.0/24
      Application: junos-ftp
        IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
          Source port range: [0-0]
          Destination port range: [21-21]
      Source identities: role1
      Per policy TCP Options: SYN check: No, SEQ check: No
```


Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 1370
show security policies policy-name p1 detail on page 1371
show security policies (services-offload) on page 1372
show security policies detail on page 1372
show security policies detail (TCP Options) on page 1373
show security policies policy-name p1 (Negated Address) on page 1373
show security policies policy-name p1 detail (Negated Address) on page 1374
show security policies global on page 1374

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 103: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 103: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 103: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```


Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                  545 bps
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                  272 bps
    Output bytes       :               18144                  545 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
      Initial direction:                9072                  272 bps
      Reply direction  :                9072                  272 bps
    Input  packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Output packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Session rate       :                 108                    3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

show security policies hit-count

Syntax

show security policies hit-count
<ascending | descending>
<from-zone *zone-name*>
<greater-than *count*>
<less-than *count*>
<to-zone *zone-name*>

Release Information

Command introduced in Junos OS Release 12.1.

Description

Display the utility rate of security policies according to the number of hits they receive. The number of hits can be listed without an order or sorted in either ascending or descending order, and they can be restricted to the number of hits that fall above or below a specific count or within a range. Data is shown for all zones associated with the policies or named zones.

Use this command without options to display the number of hits in random order for all security policies and for all zones.

Options

- **ascending** | **descending**—(Optional) Display the number of hits for security policies in ascending or descending order.
- **from-zone** ***zone-name***—(Optional) Display the number of hits for security policies associated with the named source zone.
- **greater-than** ***count***—(Optional) Display security policies for which the number of hits is greater than the specified number.

  Range: 0 through 4,294,967,295
- **less-than** ***count***—(Optional) Display security policies for which the number of hits is less than the specified number.

  Range: 0 through 4,294,967,295
- **to-zone** ***zone-name***—(Optional) Display the number of hits for security policies associated with the named destination zone.

Required Privilege Level

view

Related Documentation

- clear security policies hit-count on page 1343
- Security Policies Overview on page 1085

List of Sample Output

show security policies hit-count on page 1376
show security policies hit-count ascending on page 1376
show security policies hit-count descending greater-than 70 less-than 100 on page 1376
show security policies hit-count from-zone untrust to-zone trust on page 1376

Output Fields

Table 104 lists the output fields for the **show security policies hit-count** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 104: show security policies hit-count Output Fields

Field Name	Field Description
from-zone	Name of the source zone.
to-zone	Name of the destination zone.
policy	Name of the security policy.
hit-count	Number of hits for each security policy.
Number of policy	Number of security policies for which hit counts are displayed.

Sample Output

show security policies hit-count

```
user@host> show security policies hit-count
from-zone   to-zone  policy  hit-count
untrust     vrtrust  u2t1      40
untrust     trust    u2t2      20
untrust     trust    u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count ascending

```
user@host> show security policies hit-count ascending
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      20
untrust     vrtrust  u2t1      40
untrust     trust    u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count descending greater-than 70 less-than 100

```
user@host> show security policies hit-count descending greater-than 70 less-than 100
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      100
untrust     vrtrust  u2t1      90
untrust     vrtrust  u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count from-zone untrust to-zone trust

```
user@host> show security policies hit-count from-zone untrust to-zone trust
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      20
```


Copyright © 2016, Juniper Networks, Inc.

```
untrust     trust    u2t3      80

Number of policy: 2
```

Copyright © 2016, Juniper Networks, Inc.

show security policies unknown-source-identity

Syntax show security policies unknown-source-identity

Release Information Command introduced in Junos OS Release 12.1X45-D10.

Description Display a list of any user or role that is referenced in a policy as a source-identity, but is not yet included in the role provisioning table.

The role provisioning table is created from the local authentication table, UAC authentication tables, and firewall authentication tables. The UAC and firewall authentication tables are dynamic and contain only those users currently authenticated. Because of this, a role can be listed as unknown because no user associated with the role has authenticated yet. There is no consequence if a role remains unknown.

Required Privilege Level view

Related Documentation
- Security Policies Overview on page 1085

List of Sample Output show security policies unknown-source-identity on page 1378

Output Fields Table 105 lists the output fields for the **show security policies unknown-source-identity** command. Output fields are listed in the approximate order in which they appear.

Table 105: show security policies unknown-source-identity Output Fields

Field Name	Field Description
From zone	Part of the zone pair that identifies the source of the traffic to which a policy applies. Affected policies are grouped by their zone pair.
To zone	Part of the zone pair that identifies the destination of the traffic to which a policy applies. Affected policies are grouped by their zone pair.
Policy	The name of the policy that contains the unknown source identity.
Unknown source identities	A list of user names and roles specified in the source-identity field of the named policy that are unknown.

Sample Output

show security policies unknown-source-identity

In the following sample output, policy p1 which controls traffic from the untrust zone to the trust zone specifies two roles, r1 and r3, that are not yet provisioned. Similarly, policy p2 affecting traffic from the trust zone to the trust zone also contains two roles that are not provisioned, role1 and abc.

```
user@host> show security policies unknown-source-identity
```


Copyright © 2016, Juniper Networks, Inc.

```
From zone: untrust, To zone: trust
  Policy: p1
  Unknown source identities: r1, r3
From zone: trust, To zone: trust
  Policy: p2
  Unknown source identities: role1, abc
```

Copyright © 2016, Juniper Networks, Inc.

show security shadow-policies logical-system

Syntax show security shadow-policies logical-system *lsys-name* [from-zone *from-zone-name* to-zone *to-zone-name* policy *policy-name* reverse | global policy *policy-name* reverse]

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display the shadowing and shadowed policies in a policy list.

Required Privilege Level view

Related Documentation
- show security policies on page 802

List of Sample Output
show security shadow-policies from-zone zone-a to-zone zone-b on page 1380
show security shadow-policies from-zone zone-a to-zone zone-b policy P1 on page 1380
show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse on page 1380

Output Fields Table 106 lists the output fields for the **show security shadow-policies logical-system** command. Output fields are listed in the approximate order in which they appear.

Table 106: show security shadow-policies logical-system Output Fields

Field Name	Field Description
Policies	The policies shadowing one or more policies in the policy list.
Shadowed policies	The policies shadowed by one or more policies in the policy list.

Sample Output

show security shadow-policies from-zone zone-a to-zone zone-b

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b
    Policies              Shadowed policies
    P1                    P3
    P1                    P4
    P2                    P5
```

show security shadow-policies from-zone zone-a to-zone zone-b policy P1

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P1
   Policies              Shadowed policies
    P1                    P3
    P1                    P4
```

show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse

```
root@host> show security shadow-policies from-zone zone-a to-zone zone-b policy P4 reverse
   Policies              Shadowed policies
    P1                    P4
```


Copyright © 2016, Juniper Networks, Inc.

show security user-identification local-authentication-table

Syntax show security user-identification local-authentication-table [(all [brief | extensive])
|ip-address *ip-address* | role *role-name* | start *value* | count *value* | user *user-name*]

Release Information Command introduced in release 12.1 of Junos OS.

Description This command displays the content of the local authentication table by IP address.

all—(Optional) All entries displayed from the beginning of the table or from the specified starting entry.

brief—(Default) Uses a tabular format and truncates longer entries:
username—displays up to 13 characters, roles—displays up to 32 characters.

extensive—(Optional) Displays the full names and all items.

count *value*—(Optional) The total number of entries to display.

ip-address *ip-address*—(Optional) The IP address of the entry to display.

role *role-name*—(Optional) The role name of the entries to display.

start *value*—(Optional) The first entry to display.

user *user-name*—(Optional) The username of the entry to display.

Required Privilege Level view

Related Documentation
- request security user-identification local-authorization-table add on page 1346
- Understanding the User Identification Table on page 1130

List of Sample Output
show security user-identification local-authentication-table all on page 1382
show security user-identification local-authentication-table ip-address on page 1382
show security user-identification local-authentication-table start on page 1382
show security user-identification local-authentication-table role on page 1382

Output Fields Table 107 lists the output fields for the **show security user-identification local-authentication-table** command. Output fields are listed in the approximate order in which they appear.

Table 107: show security user-identification local-authentication-table Output Fields

Field Name	Field Description
Total entries	The number of entries in the table.
IP address	IP address of the associated user. NOTE: Only one user can be associated with an IP address.
Username	User associated with the specified IP address.

Copyright © 2016, Juniper Networks, Inc.

Table 107: show security user-identification local-authentication-table Output Fields *(continued)*

Field Name	Field Description
Roles	A comma-separated list of all roles associated with this IP address and user.

Sample Output

show security user-identification local-authentication-table all

```
user@host> show security user-identification local-authentication-table all
Total entries: 3
Source IP       Username     Roles
192.0.2.1       user1        role1
203.0.113.2     user1        role2
198.51.100.3    user3        role1, role2
```

show security user-identification local-authentication-table ip-address

```
user@host> show security user-identification local-authentication-table ip-address 203.0.113.2
Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1
```

show security user-identification local-authentication-table start

```
user@host> show security user-identification local-authentication-table start 2 count 2
Total entries: 2
    Ip-address: 203.0.113.2
    Username: user2
    Roles: role2, role3, role1

    Ip-address: 198.51.100.3
    Username: user3
    Roles: role2, role3
```

show security user-identification local-authentication-table role

```
user@host> show security user-identification local-authentication-table role qa3456
Total entries: 3
    Ip-address: 203.0.113.2
    Username: dev-grp-3
    Roles: qa432, qa3456, qa84, qa794

    Ip-address: 198.51.100.3
    Username: dev-qa
    Roles: qa3456, qa3985, qa23

    Ip-address: 203.0.113.2
    Username: brandall
    Roles: qa3456
```


Copyright © 2016, Juniper Networks, Inc.

show security user-identification role-provision all

Syntax show security user-identification role-provision all

Release Information Command introduced in Junos OS Release 12.1.

Description Display all the available user roles for policy provisioning. The output combines user roles from all available UITs.

Required Privilege Level view

Related Documentation
- show security user-identification user-provision all on page 1385
- show security user-identification source-identity-provision all on page 1384

List of Sample Output show security user-identification role-provision all on page 1383

Output Fields Table 108 lists the output fields for the **show security user-identification role-provision all** command. Output fields are listed in the approximate order in which they appear.

Table 108: show security user-identification role-provision all Output Fields

Field Name	Field Description
Roles	A comma-separated list of all user roles available for provisioning in user role policies. This list combines user roles from both the local authentication table and any UAC authentication tables that have been configured.

Sample Output

show security user-identification role-provision all

```
user@host> show security user-identification role-provision all
Roles: role1, role2, role3, role4, role_0_1, role_1_1
```

Copyright © 2016, Juniper Networks, Inc.

show security user-identification source-identity-provision all

Syntax show security user-identification source-identity-provision all

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display the available source identities for policy provisioning. The output combines users and user roles from all available UITs.

Required Privilege Level view

Related Documentation
- show security user-identification role-provision all on page 1383
- show security user-identification user-provision all on page 1385

List of Sample Output show security user-identification source-identity-provision all on page 1384

Output Fields Table 109 lists the output fields for the **show security user-identification source-identity-provision all** command. Output fields are listed in the approximate order in which they appear.

Table 109: show security user-identification source-identity-provision all Output Fields

Field Name	Field Description
Source identities	A comma-separated list of all users and user roles available for policy provisioning. This list combines users and user roles from both the local authentication table and any UAC authentication tables that have been configured.

Sample Output

show security user-identification source-identity-provision all

```
user@host> show security user-identification source-identity-provision all
Source identities: ariana, ben, guest5, role1, role2, role3, role4, role_0_1,
role_1_1,u1, user2
```


Copyright © 2016, Juniper Networks, Inc.

show security user-identification user-provision all

Syntax show security user-identification user-provision all

Release Information Command introduced in Release Junos OS 12.1X44-D10.

Description Display the available user names for policy provisioning. The output combines users from all available UITs.

Required Privilege Level view

Related Documentation
- show security user-identification role-provision all on page 1383
- show security user-identification source-identity-provision all on page 1384

List of Sample Output show security user-identification user-provision all on page 1385

Output Fields Table 110 lists the output fields for the **show security user-identification user-provision all** command. Output fields are listed in the approximate order in which they appear.

Table 110: show security user-identification user-provision all Output Fields

Field Name	Field Description
Users	A comma-separated list of all users available for policy provisioning. This list combines users from both the local authentication table and any UAC authentication tables that have been configured.

Sample Output

show security user-identification user-provision all

```
user@host> show security user-identification user-provision all
Users: ariana, ben, guest5, u1 user2, ...
```

Copyright © 2016, Juniper Networks, Inc.

show security zones

Syntax
show security zones
<detail | terse>
< *zone-name* >

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones.

Options
- none—Display information about all zones.
- **detail | terse**—(Optional) Display the specified level of output.
- ***zone-name*** —(Optional) Display information about the specified zone.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones on page 1387
show security zones abc on page 1387
show security zones abc detail on page 1387
show security zones terse on page 1388

Output Fields
Table 32 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 111: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                     Type
my-internal              Security
my-external              Security
dmz                      Security
```


Copyright © 2016, Juniper Networks, Inc.

show security zones type

Syntax
show security zones type
(functional | security)
<detail | terse>

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones of the specified type.

Options
- **functional**—Display functional zones.
- **security**—Display security zones.
- **detail | terse**—(Optional) Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones type functional on page 1390
show security zones type security on page 1390
show security zones type security terse on page 1390
show security zones type security detail on page 1390

Output Fields
Table 33 lists the output fields for the **show security zones type** command. Output fields are listed in the approximate order in which they appear.

Table 112: show security zones type Output Fields

Field Name	Field Description
Security zone	Zone name.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones type functional

```
user@host> show security zones type functional
Functional zone: management
  Description: management zone
  Policy configurable: No
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security

```
user@host> show security zones type security
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security terse

```
user@host> show security zones type security terse
Zone                    Type
trust                   Security
untrust                 Security
junos-host              Security
```

Sample Output

show security zones type security detail

```
user@host> show security zones type security detail
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Copyright © 2016, Juniper Networks, Inc.

show system services dns dns-proxy

Syntax show system services dns dns-proxy

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display domain name system (DNS) proxy information.

Options
- **none**—Display DNS proxy statistics information.
- **cache**—(Optional) Display the DNS proxy cache.
- **statistics**—(Optional) Display the DNS proxy statistics.

Required Privilege Level view

Related Documentation
- clear system services dns dns-proxy on page 1345
- dns-proxy on page 1248

List of Sample Output
show system services dns-proxy statistics on page 1393
show system services dns-proxy cache on page 1393
show system services dns-proxy cache <viewname V1> on page 1393

Output Fields Table 113 lists the output fields for the **show system services dns-proxy** command. Output fields are listed in the approximate order in which they appear.

Table 113: show system services dns-proxy

Field Name	Field Description
DNS proxy statistics	Display information about the DNS proxy. • **Status**—State of the proxy server as Enabled or disabled. • **Queries received**—Number of DNS queries received by the DNS proxy. • **Responses sent**—Number of DNS responses sent by the DNS proxy. • **Queries forwarded**—Number of DNS queries forwarded by the DNS proxy. • **Negative responses**—Number of negative responses the DNS proxy sent to the DNS client. • **Retry requests**—Number of retries the DNS proxy received from the DNS client. • **Pending requests**—Number of pending queries the DNS proxy has yet to send the DNS client a response for. • **Server failures**—Number of DNS proxy server failures.
Hostname	Hostname of the host that has been cached.
IP address	IP address of the host.

Copyright © 2016, Juniper Networks, Inc.

Table 113: show system services dns-proxy *(continued)*

Field Name	Field Description
Time-to-live	Length of time before an entry is purged from the DNS cache.
Type	Type of DNS Resource Record. For example, A records refer to IPv4 host addresses.
Class	Class of DNS. A parameter used to define a DNS Resource Record. For example, IN class is used for Internet domain names.

Sample Output

show system services dns-proxy statistics

```
user@host> show system services dns-proxy statistics
DNS proxy statistics        :
    DNS proxy statistics    :
    Status              : enabled
    IPV4 Queries received   : 30
    IPV6 Queries received   : 0
    Responses sent      : 30
    Queries forwarded   : 13
    Negative responses  : 23
    Positive responses  : 23
    Retry requests      : 0
    Pending requests    : 0
    Server failures     : 0
    Interfaces          : fe-0/0/0.0, fe-1/0/1.0
```

show system services dns-proxy cache

```
user@host>show system services dns-proxy cache
Hostname                Time-to-live      Type        Class   IP address/Hostname
example.com               408               A           IN      192.0.2.50
whitestar.example.com     408               A           IN      172.17.27.50
scarlet.example.com       408               A           IN      172.17.28.11
bng-admin1.example.com    408               A           IN      10.209.194.131
wf-nis1.example.com       408               A           IN      10.10.4.202
asg-ns1.example.com       408               A           IN      10.16.0.11
ruby.example.com          408               A           IN      172.17.28.100
a.l.example.com           408               A           IN      203.0.113.9
b.l.example.com           408               A           IN      198.51.100.9
maps.l.example.com        408               A           IN      198.51.100.104
c.l.example.com           408               A           IN      198.51.100.39
d.l.example.com           408               A           IN      198.51.100.9
e.l.example.com           408               A           IN      198.51.100.19
g.l.example.com           408               A           IN      198.51.100.89
magenta.example.com.      408               A           IN      172.17.27.123
mail.example.com.         408               CNAME       IN
magenta.example.com.
```

show system services dns-proxy cache <viewname V1>

```
user@host>show system services dns-proxy cache <viewname V1>
Hostname                Time-to-live    Type        Class     IP address/Hostname
example.com.              408             A           IN      198.51.100.123
example.com.              408             A           IN      203.0.113.9
```

Copyright © 2016, Juniper Networks, Inc.

```
example.com.              408         A           IN       192.0.2.70
magenta.example.com.      495         A           IN       172.17.27.123
mail.example.com.         495         CNAME       IN
magenta.example.com.
```


Copyright © 2016, Juniper Networks, Inc.

show system services dynamic-dns

Syntax show system services dynamic-dns

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display information about dynamic DNS clients.

Required Privilege Level view

Related Documentation
- dynamic-dns on page 1249

List of Sample Output
show system services dynamic-dns client on page 1395
show system services dynamic-dns client detail on page 1395

Output Fields Table 114 lists the output fields for the **show system services dynamic-dns** command. Output fields are listed in the approximate order in which they appear.

Table 114: show system services dynamic-dns

Field Name	Field Description
Hostname	Hostname of the registered client
Server	DDNS server name
Agent	Name of the DDNS agent
Last response	Status of the last response
Last update	Date and time of the last update
Interface	Name of the interface

Sample Output

show system services dynamic-dns client

```
user@host> show system services dynamic-dns client
Internal hostname      Server                  Last response
example.com             example.net                success
jnr.example.org              example.net                failure
newsam.example.com      members.example.org      nochg
rkhetan.example.com      members.example.org      nochg
```

show system services dynamic-dns client detail

```
user@host>show system services dynamic-dns client detail
Hostname      : example.com
Server        : example.net
Agent         : branch-0.1
```

Copyright © 2016, Juniper Networks, Inc.

```
Last response: success
Last update  : 2006-08-29 04:02:52 PDT
Interface    : fe-0/0/0.0

Hostname     : jnr.example.org
Server       : example.net
Agent        : Branch-0.1
Last response: failure
Last update  : 2006-08-29 04:03:03 PDT
Interface    : fe-0/0/0.0

Hostname     : newsam.example.com
Server       : members.example.org
Agent        : Branch-0.1
Last response: nochg
Last update  : 2006-08-29 04:02:50 PDT
Username     : user1
Interface    : fe-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Class of Service Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 16

Overview

- Introduction to Class of Service on page 1401

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 65

Introduction to Class of Service

- Understanding Class of Service on page 1401
- Benefits of CoS on page 1402
- CoS Across the Network on page 1403
- Junos OS CoS Components on page 1403
- CoS Components Packet Flow on page 1405
- CoS Device Configuration Overview on page 1407
- Understanding CoS Default Settings on page 1407

Understanding Class of Service

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to the rules you configure.

For interfaces that carry IPv4, IPv6, or MPLS traffic, you can configure the Junos OS CoS features to provide multiple classes of service for different applications. On the device, you can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

Traffic shaping is the allocation of the appropriate amount of network bandwidth to every user and application on an interface. The appropriate amount of bandwidth is defined as cost-effective carrying capacity at a guaranteed CoS. You can use a Juniper Networks device to control traffic rate by applying classifiers and shapers.

The CoS features provide a set of mechanisms that you can use to provide differentiated services when best-effort delivery is insufficient.

Using Junos OS CoS features, you can assign service levels with different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows. CoS is especially useful for networks supporting time-sensitive video and audio applications.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Policing, scheduling, and shaping CoS services are not supported for pre-encryption and post-encryption packets going into and coming out of an IPsec VPN tunnel.

Junos OS supports the following RFCs for traffic classification and policing:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2697, *A Single Rate Three Color Marker*
- RFC 2698, A *Two Rate Three Color Marker*

Related Documentation

- Junos OS CoS Components on page 1403
- CoS Components Packet Flow on page 1405
- Understanding CoS Default Settings on page 1407
- CoS Device Configuration Overview on page 1407

Benefits of CoS

IP routers normally forward packets independently, without controlling throughput or delay. This type of packet forwarding, known as *best-effort service*, is as good as your network equipment and links allow. Best-effort service is sufficient for many traditional IP data delivery applications, such as e-mail or Web browsing. However, newer IP applications such as real-time video and audio (or voice) require lower delay, jitter, and packet loss than simple best-effort networks can provide.

CoS features allow a Juniper Networks device to improve its processing of critical packets while maintaining best-effort traffic flows, even during periods of congestion. Network throughput is determined by a combination of available bandwidth and delay. CoS dedicates a guaranteed minimum bandwidth to a particular service class by reducing forwarding queue delays. (The other two elements of overall network delay, serial transmission delays determined by link speeds and propagation delays determined by media type, are not affected by CoS settings.)

Normally, packets are queued for output in their order of arrival, regardless of service class. Queuing delays increase with network congestion and often result in lost packets when queue buffers overflow. CoS packet classification assigns packets to forwarding queues by service class.

Because CoS must be implemented consistently end-to-end through the network, the CoS features on the Juniper Networks device are based on IETF Differentiated Services (DiffServ) standards to interoperate with other vendors' CoS implementations.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Understanding Class of Service on page 1401

CoS Across the Network

CoS works by examining traffic entering at the edge of your network. The edge devices classify traffic into defined service groups, which allow for the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic.

To support CoS, you must configure each device in the network. Generally, each device examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are first transmitted to the next downstream device. In addition, the devices at the edges of your network might be required to alter the CoS settings of the packets transmitting to the neighboring network.

Figure 70 shows an example of CoS operating across an Internet Service Provider (ISP) network.

Figure 70: CoS Across the Network

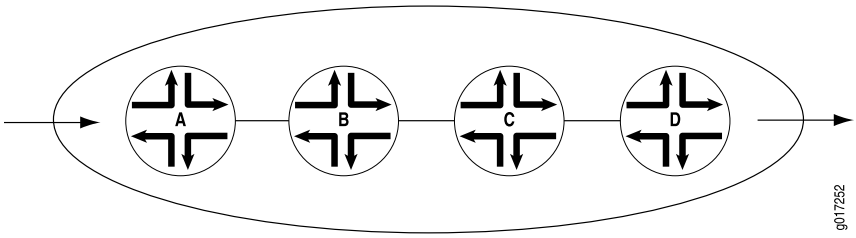

In the ISP network shown in Figure 70, Device A is receiving traffic from your network. As each packet enters, Device A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined by the ISP. This definition allows Device A to prioritize its resources for servicing the traffic streams it is receiving. In addition, Device A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the ISP's traffic groups. When Device B receives the packets, it examines the CoS settings, determines the appropriate traffic group, and processes the packet according to those settings. Device B then transmits the packets to Device C, which performs the same actions. Device D also examines the packets and determines the appropriate group. Because it sits at the far end of the network, the ISP might decide once again to alter the CoS settings of the packets before Device D transmits them to the neighboring network.

Related Documentation
- Understanding Class of Service on page 1401

Junos OS CoS Components

Junos OS supports CoS components on Juniper Networks devices as indicated in Table 115.

Copyright © 2016, Juniper Networks, Inc.

Table 115: Supported Junos OS CoS Components

Junos OS CoS Component	Description	For More Information
Code-point aliases	A code-point alias assigns a name to a pattern of code-point bits. You can use this name, instead of the bit pattern, when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.	"Code-Point Aliases Overview" on page 1549
Classifiers	Packet classification refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level. Two general types of classifiers are supported—behavior aggregate (BA) classifiers and multifield (MF) classifiers. When both BA and MF classifications are performed on a packet, the MF classification has higher precedence.	"Classification Overview" on page 1411
Forwarding classes	Forwarding classes allow you to group packets for transmission. Based on forwarding classes, you assign packets to output queues. The forwarding class plus the loss priority define the per-hop behavior (PHB in DiffServ) of a packet. Juniper Networks routers and services gateways support eight queues (0 through 7).	"Forwarding Classes Overview" on page 1453
Loss priorities	Loss priorities allow you to set the priority of dropping a packet. You can use the loss priority setting to identify packets that have experienced congestion.	"Understanding Packet Loss Priorities" on page 1414
Forwarding policy options	CoS-based forwarding (CBF) enables you to control next-hop selection based on a packet's class of service and, in particular, the value of the IP packet's precedence bits. For example, you can specify a particular interface or next hop to carry high-priority traffic while all best-effort traffic takes some other path.	"Example: Assigning a Forwarding Class to an Interface" on page 1462
Transmission queues	After a packet is sent to the outgoing interface on a device, it is queued for transmission on the physical media. The amount of time a packet is queued on the device is determined by the availability of the outgoing physical media as well as the amount of traffic using the interface. Juniper Networks routers and services gateways support queues 0 through 7.	"Transmission Scheduling Overview" on page 1479
Schedulers	An individual device interface has multiple queues assigned to store packets temporarily before transmission. To determine the order to service the queues, the device uses a round-robin scheduling method based on priority and the queue's weighted round-robin (WRR) credits. Junos OS schedulers allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission.	"Schedulers Overview" on page 1473
Virtual channels	On Juniper Networks routers and services gateways, you can configure virtual channels to limit traffic sent from a corporate headquarters to branch offices. Virtual channels might be required when the headquarters site has an expected aggregate bandwidth higher than that of the individual branch offices. The router at the headquarters site must limit the traffic sent to each branch office router to avoid oversubscribing their links.	"Virtual Channels Overview" on page 1531
Policers for traffic classes	Policers allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with firewall filters that can be associated with input or output interfaces.	"Simple Filters and Policers Overview" on page 1427

Copyright © 2016, Juniper Networks, Inc.

Table 115: Supported Junos OS CoS Components *(continued)*

Junos OS CoS Component	Description	For More Information
Rewrite rules	A rewrite rule modifies the appropriate CoS bits in an outgoing packet. Modification of CoS bits allows the next downstream device to classify the packet into the appropriate service group. Rewriting or marking outbound packets is useful when the device is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.	"Rewrite Rules Overview" on page 1467

Related Documentation

- Understanding Class of Service on page 1401
- CoS Components Packet Flow on page 1405
- Understanding CoS Default Settings on page 1407
- CoS Device Configuration Overview on page 1407

CoS Components Packet Flow

On Juniper Networks devices, you configure CoS functions using different components. These components are configured individually or in a combination to define particular CoS services. Figure 71 displays the relationship of different CoS components to each other and illustrates the sequence in which they interact.

Figure 71: Packet Flow Through Juniper Networks Device

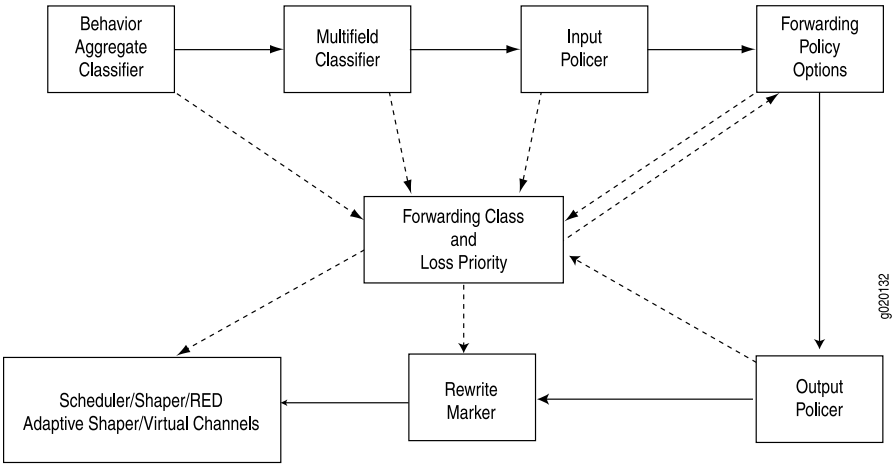

Each box in Figure 71 represents a CoS component. The solid lines show the direction of packet flow in a device. The upper row indicates an incoming packet, and the lower row an outgoing packet. The dotted lines show the inputs and outputs of particular CoS components. For example, the forwarding class and loss priority are outputs of behavior aggregate classifiers and multifield classifiers and inputs for rewrite markers and schedulers.

Typically, only a combination of some components shown in Figure 71 (not all) is used to define a CoS service offering. For example, if a packet's class is determined by a

Copyright © 2016, Juniper Networks, Inc.

behavior aggregate classifier, it is associated with a forwarding class and loss priority and does not need further classification by the multifield classifier.

This section contains the following topics:

- CoS Process on Incoming Packets on page 1406
- CoS Process on Outgoing Packets on page 1406

CoS Process on Incoming Packets

Classifiers and policers perform the following operations on incoming packets:

1. A classifier examines an incoming packet and assigns a forwarding class and loss priority to it.
2. Based on the forwarding class, the packet is assigned to an outbound transmission queue.
3. Input policers meter traffic to see if traffic flow exceeds its service level. Policers might discard, change the forwarding class and loss priority, or set the PLP bit of a packet. A packet for which the PLP bit is set has an increased probability of being dropped during congestion.

CoS Process on Outgoing Packets

The scheduler map and rewrite rules perform the following operations on outgoing packets:

1. Scheduler maps are applied to interfaces and associate the outgoing packets with a scheduler and a forwarding class.
2. The scheduler defines how the packet is treated in the output transmission queue based on the configured transmit rate, buffer size, priority, and drop profile.
   - The buffer size defines the period for which the packet is stored during congestion.
   - The scheduling priority and transmit rate determine the order in which the packet is transmitted.
   - The drop profile defines how aggressively to drop packets that are using a particular scheduler.
3. Output policers meter traffic and might change the forwarding class and loss priority of a packet if a traffic flow exceeds its service level.
4. The rewrite rule writes information to the packet (for example, EXP or DSCP bits) according to the forwarding class and loss priority of the packet.

Related Documentation

- Understanding Class of Service on page 1401
- Junos OS CoS Components on page 1403
- Understanding CoS Default Settings on page 1407
- CoS Device Configuration Overview on page 1407

Copyright © 2016, Juniper Networks, Inc.

CoS Device Configuration Overview

Before you begin configuring a Juniper Networks device for CoS, complete the following tasks:

- Determine whether the device needs to support different traffic streams, such as voice or video. If so, CoS helps to make sure this traffic receives more than basic best-effort packet delivery service.
- Determine whether the device is directly attached to any applications that send CoS-classified packets. If no sources are enabled for CoS, you must configure and apply rewrite rules on the interfaces facing the sources.
- Determine whether the device must support assured forwarding (AF) classes. Assured forwarding usually requires random early detection (RED) drop profiles to be configured and applied.
- Determine whether the device must support expedited forwarding (EF) classes with a policer. Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a high loss priority, so that packets exceeding the policer limits are discarded first.

NOTE: **When the T1/E1 mPIM is oversubscribed, we recommend that you configure its shaping rate for consistent CoS functionality. The shaping rate should be less than the total link speed.**

Related Documentation

- Understanding Class of Service on page 1401
- CoS Components Packet Flow on page 1405
- Understanding CoS Default Settings on page 1407

Understanding CoS Default Settings

The Class of Service menu in J-Web allows you to configure most of the Junos OS CoS components for the IPv4 and MPLS traffic on a Juniper Networks device. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm. After defining the CoS components, you must assign classifiers to the required physical and logical interfaces.

Even when you do not configure any CoS settings on your routing platform, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some default mappings are automatically applied to each logical interface that you configure. Other default mappings, such as explicit default classifiers and rewrite rules, are in operation only if you explicitly associate them with an interface.

Copyright © 2016, Juniper Networks, Inc.

You can display default CoS settings by running the **show class-of-service** operational mode command.

You configure CoS when you need to override the default packet forwarding behavior of a Juniper Networks device—especially in the three areas identified in Table 116.

Table 116: Reasons to Configure Class of Service (CoS)

Default Behavior to Override with CoS	CoS Configuration Area
Packet classification—By default, the Juniper Networks device does not use behavior aggregate (BA) classifiers to classify packets. Packet classification applies to incoming traffic.	Classifiers
Scheduling queues—By default, the Juniper Networks device has only two queues enabled. Scheduling queues apply to outgoing traffic.	Schedulers
Packet headers—By default, the Juniper Networks device does not rewrite CoS bits in packet headers. Rewriting packet headers applies to outgoing traffic.	Rewrite rules

Related Documentation

- Understanding Class of Service on page 1401
- CoS Components Packet Flow on page 1405
- CoS Device Configuration Overview on page 1407

Copyright © 2016, Juniper Networks, Inc.

PART 17

Configuring Class of Service Components

- Assigning Service Levels with Classifiers on page 1411
- Controlling Network Access with Traffic Policing on page 1427
- Controlling Output Queues with Forwarding Classes on page 1453
- Altering Outgoing Packets Headers with Rewrite Rules on page 1467
- Defining Output Queue Properties with Schedulers on page 1473
- Removing Delays with Strict-Priority Queues on page 1505
- Controlling Congestion with Drop Profiles on page 1519
- Controlling Congestion with Adaptive Shapers on page 1527
- Limiting Traffic Using Virtual Channels on page 1531
- Enabling Queuing for Tunnel Interfaces on page 1539
- Naming Components with Code-Point Aliases on page 1549

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 66

Assigning Service Levels with Classifiers

- Classification Overview on page 1411
- Understanding Packet Loss Priorities on page 1414
- Default Behavior Aggregate Classification on page 1414
- Sample Behavior Aggregate Classification on page 1416
- Example: Configuring Behavior Aggregate Classifiers on page 1417

Classification Overview

Packet classification refers to the examination of an incoming packet, which associates the packet with a particular class-of-service (CoS) servicing level. Junos operating system (OS) supports these classifiers:

- Behavior aggregate (BA) classifiers
- Multifield (MF) classifiers
- Default IP precedence classifiers

NOTE: The total number of classifiers supported on a Services Processing Unit (SPU) is 79. Three classifiers are installed on the SPU as default classifiers in the Layer 3 mode, independent of any CoS configuration, which leaves 76 classifiers that can be configured using the CoS CLI commands. The default classifiers number can vary in future releases or in different modes.

Verify the number of default classifiers installed on the SPU to determine how many classifiers can be configured using the CoS CLI commands.

When both BA and MF classifications are performed on a packet, the MF classification has higher precedence.

In Junos OS, classifiers associate incoming packets with a forwarding class (FC) and packet loss priority (PLP), and, based on the associated FC, assign packets to output queues. A packet's FC and PLP specify the behavior of a hop, within the system, to process the packet. The per-hop behavior (PHB) comprises packet forwarding, policing, scheduling, shaping, and marking. For example, a hop can put a packet in one of the priority queues

Copyright © 2016, Juniper Networks, Inc.

according to its FC and then manage the queues by checking the packet's PLP. Junos OS supports up to eight FCs and four PLPs.

This topic includes the following sections:

- Behavior Aggregate Classifiers on page 1412
- Multifield Classifiers on page 1412
- Default IP Precedence Classifier on page 1413

Behavior Aggregate Classifiers

A BA classifier operates on a packet as it enters the device. Using BA classifiers, the device aggregates different types of traffic into a single FC so that all the types of traffic will receive the same forwarding treatment. The CoS value in the packet header is the single field that determines the CoS settings applied to the packet. BA classifiers allow you to set a packet's FC and PLP based on the Differentiated Services (DiffServ) code point (DSCP) value, DSCP IPv4 value, DSCP IPv6 value, IP precedence value, MPLS EXP bits, or IEEE 802.1p value. The default classifier is based on the IP precedence value. For more information, see "Default IP Precedence Classifier" on page 1413.

Junos OS performs BA classification for a packet by examining its Layer 2, Layer 3, and related CoS parameters, as shown in Table 117.

Table 117: BA Classification

Layer	CoS Parameter
Layer 2	IEEE 802.1p value: User Priority
Layer 3	IPv4 precedence IPv4 Differentiated Services code point (DSCP) value IPv6 DSCP value

NOTE: A BA classifier evaluates Layer 2 and Layer 3 parameters independently. The results from Layer 2 parameters override the results from the Layer 3 parameters.

Multifield Classifiers

An MF classifier is a second means of classifying traffic flows. Unlike the BA classifier, an MF classifier can examine multiple fields in the packet—for example, the source and destination address of the packet, or the source and destination port numbers of the packet. With MF classifiers, you set the FC and PLP based on firewall filter rules.

Copyright © 2016, Juniper Networks, Inc.

NOTE: For a specified interface, you can configure both an MF classifier and a BA classifier without conflicts. Because the classifiers are always applied in sequential order (the BA classifier followed by the MF classifier), any BA classification result is overridden by an MF classifier if they conflict.

Junos OS performs MF traffic classification by directly scrutinizing multiple fields of a packet to classify a packet. This avoids having to rely on the output of the previous BA traffic classification. Junos OS can simultaneously check a packet's data for Layers 2, 3, 4, and 7, as shown in Table 118.

Table 118: MF Classification

Layer	CoS Parameter
Layer 2	IEEE 802.1Q: VLAN ID IEEE 802.1p: User priority
Layer 3	IP precedence value DSCP or DSCP IPv6 value Source IP address Destination IP address Protocol ICMP: Code and type
Layer 4	TCP/UDP: Source port TCP/UDP: Destination port TCP: Flags AH/ESP: SPI
Layer 7	Not supported for this release.

Using Junos OS, you configure an MF classifier with a firewall filter and its associated match conditions. This enables you to use any filter match criterion to locate packets that require classification.

Default IP Precedence Classifier

With Junos OS, all logical interface are automatically assigned a default IP precedence classifier when the logical interface is configured. This default traffic classifier maps IP precedence values to an FC and a PLP as shown in Table 119. These mapping results are in effect for an ingress packet until the packet is further processed by another classification method.

Copyright © 2016, Juniper Networks, Inc.

Table 119: Default IP Precedence Classifier

IP Precedence CoS Values	Forwarding Class	Packet Loss Priority
000	best-effort	low
001	best-effort	high
010	best-effort	low
011	best-effort	high
100	best-effort	low
101	best-effort	high
110	network-control	low
111	network-control	high

Related Documentation

- Default Behavior Aggregate Classification on page 1414
- Sample Behavior Aggregate Classification on page 1416
- Example: Configuring Behavior Aggregate Classifiers on page 1417

Understanding Packet Loss Priorities

Packet loss priorities (PLPs) allow you to set the priority for dropping packets. You can use the PLP setting to identify packets that have experienced congestion. Typically, you mark packets exceeding some service level with a high loss priority—that is, a greater likelihood of being dropped. You set PLP by configuring a classifier or a policer. The PLP is used later in the work flow to select one of the drop profiles used by random early detection (RED).

You can configure the PLP bit as part of a congestion control strategy. The PLP bit can be configured on an interface or in a filter. A packet for which the PLP bit is set has an increased probability of being dropped during congestion.

Related Documentation

- Classification Overview on page 1411
- Default Behavior Aggregate Classification on page 1414
- Sample Behavior Aggregate Classification on page 1416
- Example: Configuring Behavior Aggregate Classifiers on page 1417

Default Behavior Aggregate Classification

Table 120 shows the forwarding class (FC) and packet loss priority (PLP) that are assigned by default to each well-known Differentiated Services (DiffServ) code point (DSCP).

Copyright © 2016, Juniper Networks, Inc.

Although several DSCPs map to the expedited-forwarding (ef) and assured-forwarding (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1*x* are mapped to best-effort, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes. Assignment to the best-effort FC implies that the node does not support that class. You can modify the default settings through configuration.

Table 120: Default Behavior Aggregate Classification

DSCP and DSCP IPv6 Alias	Forwarding Class	Packet Loss Priority
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured-forwarding	high
af13	assured-forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low
af43	best-effort	low
be	best-effort	low
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low

Copyright © 2016, Juniper Networks, Inc.

Table 120: Default Behavior Aggregate Classification *(continued)*

DSCP and DSCP IPv6 Alias	Forwarding Class	Packet Loss Priority
nc2/cs7	network-control	low
other	best-effort	low

Related Documentation

- Classification Overview on page 1411
- Sample Behavior Aggregate Classification on page 1416
- Example: Configuring Behavior Aggregate Classifiers on page 1417
- Understanding Packet Loss Priorities on page 1414

Sample Behavior Aggregate Classification

Table 121 shows the device forwarding classes (FCs) associated with each well-known Differentiated Services (DiffServ) code point (DSCP) and the resources assigned to the output queues for a sample DiffServ CoS implementation. This example assigns expedited forwarding to queue 1 and a subset of the assured FCs (af*x*) to queue 2, and distributes resources among all four forwarding classes. Other DiffServ-based implementations are possible.

Table 121: Sample Behavior Aggregate Classification Forwarding Classes and Queues

DSCP and DSCP IPv6 Alias	DSCP and DSCP IPv6 Bits	Forwarding Class	Packet Loss Priority	Queue
ef	101110	expedited-forwarding	low	1
af11	001010	assured-forwarding	low	2
af12	001100	assured-forwarding	high	2
af13	001110	assured-forwarding	high	2
af21	010010	best-effort	low	0
af22	010100	best-effort	low	0
af23	010110	best-effort	low	0
af31	011010	best-effort	low	0
af32	011100	best-effort	low	0
af33	011110	best-effort	low	0
af41	100010	best-effort	low	0

Copyright © 2016, Juniper Networks, Inc.

Table 121: Sample Behavior Aggregate Classification Forwarding Classes and Queues *(continued)*

DSCP and DSCP IPv6 Alias	DSCP and DSCP IPv6 Bits	Forwarding Class	Packet Loss Priority	Queue
af42	100100	best-effort	low	0
af43	100110	best-effort	low	0
be	000000	best-effort	low	0
cs1	0010000	best-effort	low	0
cs2	010000	best-effort	low	0
cs3	011000	best-effort	low	0
cs4	100000	best-effort	low	0
cs5	101000	best-effort	low	0
nc1/cs6	110000=	network-control	low	3
nc2/cs7	111000=	network-control	low	3
other	—	best-effort	low	0

Related Documentation

- Classification Overview on page 1411
- Default Behavior Aggregate Classification on page 1414
- Example: Configuring Behavior Aggregate Classifiers on page 1417
- Understanding Packet Loss Priorities on page 1414

Example: Configuring Behavior Aggregate Classifiers

This example shows how to configure behavior aggregate classifiers for a device to determine forwarding treatment of packets.

- Requirements on page 1417
- Overview on page 1418
- Configuration on page 1419
- Verification on page 1421

Requirements

Before you begin, determine the forwarding class and PLP that are assigned by default to each well-known DSCP that you want to configure for the behavior aggregate classifier. See "Default Behavior Aggregate Classification" on page 1414.

Copyright © 2016, Juniper Networks, Inc.

Overview

You configure behavior aggregate classifiers to classify packets that contain valid DSCPs to appropriate queues. Once configured, you must apply the behavior aggregate classifier to the correct interfaces. You can override the default IP precedence classifier by defining a classifier and applying it to a logical interface. To define new classifiers for all code point types, include the **classifiers** statement at the **[edit class-of-service]** hierarchy level.

In this example, you set the DSCP behavior aggregate classifier to ba-classifier as the default DSCP map. You set a best-effort forwarding class as be-class, an expedited forwarding class as ef-class, an assured forwarding class as af-class, and a network control forwarding class as nc-class. Finally, you apply the behavior aggregate classifier to an interface called ge-0/0/0.

Table 122 shows how the behavior aggregate classifier assigns loss priorities, to incoming packets in the four forwarding classes.

Table 122: Sample ba-classifier Loss Priority Assignments

mf-classifier Forwarding Class	For CoS Traffic Type	ba-classifier Assignments
be-class	Best-effort traffic	High-priority code point: **000001**
ef-class	Expedited forwarding traffic	High-priority code point: **101111**
af-class	Assured forwarding traffic	High-priority code point: **001100**
nc-class	Network control traffic	High-priority code point: **110001**

Topology

Figure 72 shows the sample network.

Figure 72: Behavior Aggregate Classifier Scenario

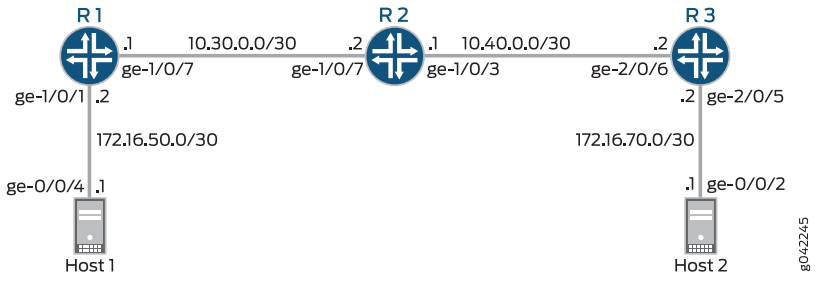

"CLI Quick Configuration" on page 1419 shows the configuration for all of the Juniper Networks devices in Figure 72.

The section "Step-by-Step Procedure" on page 1419 describes the steps on Device R2.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

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

```
set class-of-service classifiers dscp ba-classifier import default
set class-of-service classifiers dscp ba-classifier forwarding-class be-class loss-priority
  high code-points 000001
set class-of-service classifiers dscp ba-classifier forwarding-class ef-class loss-priority
  high code-points 101111
set class-of-service classifiers dscp ba-classifier forwarding-class af-class loss-priority
  high code-points 001100
set class-of-service classifiers dscp ba-classifier forwarding-class nc-class loss-priority
  high code-points 110001
set class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp ba-classifier
```

Step-by-Step Procedure

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

To configure behavior aggregate classifiers for a device:

1. Configure the class of service.

   ```
   [edit]
   user@host# edit class-of-service
   ```

2. Configure behavior aggregate classifiers for DiffServ CoS.

   ```
   [edit class-of-service]
   user@host# edit classifiers dscp ba-classifier
   user@host# set import default
   ```

3. Configure a best-effort forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp ba-classifier]
   user@host# set forwarding-class be-class loss-priority high code-points 000001
   ```

4. Configure an expedited forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp ba-classifier]
   user@host# set forwarding-class ef-class loss-priority high code-points 101111
   ```

5. Configure an assured forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp ba-classifier]
   user@host# set forwarding-class af-class loss-priority high code-points 001100
   ```

6. Configure a network control forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp ba-classifier]
   user@host# set forwarding-class nc-class loss-priority high code-points 110001
   ```

7. Apply the behavior aggregate classifier to an interface.

   ```
   [edit]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp
  ba-classifier
```

NOTE: You can use interface wildcards for **interface-name** and **logical-unit-number**.

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
classifiers {
  dscp ba-classifier {
    import default;
    forwarding-class be-class {
      loss-priority high code-points 000001;
    }
    forwarding-class ef-class {
      loss-priority high code-points 101111;
    }
    forwarding-class af-class {
      loss-priority high code-points 001100;
    }
    forwarding-class nc-class {
      loss-priority high code-points 110001;
    }
  }
  forwarding-classes {
    class BE-data queue-num 0;
    class Premium-data queue-num 1;
    class Voice queue-num 2;
    class NC queue-num 3;
  }
  interfaces {
    ge-0/0/0 {
      unit 0 {
        classifiers {
          dscp ba-classifier;
        }
      }
    }
    ge-1/0/9 {
      unit 0 {
        classifiers {
          dscp v4-ba-classifier;
        }
        ge-1/0/9 {
          unit 0 {
            classifiers {
              dscp v4-ba-classifier;
            }
```


Copyright © 2016, Juniper Networks, Inc.

```
    ge-1/0/9 {
      unit 0 {
        classifiers {
          dscp v4-ba-classifier;
        }
      }
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Code-Point Aliases on page 1421
- Verifying the DSCP Classifier on page 1422
- Verifying the Forwarding Classes and Output Queues on page 1423
- Verifying That the Classifier Is Applied to the Interfaces on page 1424
- Verifying Behavior Aggregate Classifiers on page 1424

Verifying the Code-Point Aliases

Purpose Make sure that the code-point aliases are configured as expected.

Action On Device R2, run the **show class-of-service code-point-aliases dscp** command.

```
user@R2> show class-of-service code-point-aliases dscp
```

```
Code point type: dscp
  Alias              Bit pattern
  af11               001010
  af12               001100
  af13               001110
  af21               010010
  af22               010100
  af23               010110
  af31               011010
  af32               011100
  af33               011110
  af41               100010
  af42               100100
  af43               100110
  be                 000000
  be1                000001
  cs1                001000
  cs2                010000
  cs3                011000
  cs4                100000
  cs5                101000
  cs6                110000
  cs7                111000
  ef                 101110
  ef1                101111
  nc1                110000
```

```
  nc2                 111000
```

Meaning The code-point aliases are configured as expected. Notice that the custom aliases that you configure are added to the default code-point aliases.

Verifying the DSCP Classifier

Purpose Make sure that the DSCP classifier is configured as expected.

Action On Device R2, run the **show class-of-service classifiers name v4-ba-classifier** command.

```
user@R2> show class-of-service classifiers name v4-ba-classifier

Classifier: v4-ba-classifier, Code point type: dscp, Index: 10755
  Code point          Forwarding class                    Loss priority
  000000              BE-data                             high
  000001              BE-data                             low
  000010              BE-data                             low
  000011              BE-data                             low
  000100              BE-data                             low
  000101              BE-data                             low
  000110              BE-data                             low
  000111              BE-data                             low
  001000              BE-data                             low
  001001              BE-data                             low
  001010              Voice                               low
  001011              BE-data                             low
  001100              Voice                               high
  001101              BE-data                             low
  001110              Voice                               high
  001111              BE-data                             low
  010000              BE-data                             low
  010001              BE-data                             low
  010010              BE-data                             low
  010011              BE-data                             low
  010100              BE-data                             low
  010101              BE-data                             low
  010110              BE-data                             low
  010111              BE-data                             low
  011000              BE-data                             low
  011001              BE-data                             low
  011010              BE-data                             low
  011011              BE-data                             low
  011100              BE-data                             low
  011101              BE-data                             low
  011110              BE-data                             low
  011111              BE-data                             low
  100000              BE-data                             low
  100001              BE-data                             low
  100010              BE-data                             low
  100011              BE-data                             low
  100100              BE-data                             low
  100101              BE-data                             low
  100110              BE-data                             low
  100111              BE-data                             low
  101000              BE-data                             low
  101001              BE-data                             low
```


Copyright © 2016, Juniper Networks, Inc.

```
  101010            BE-data                         low
  101011            BE-data                         low
  101100            BE-data                         low
  101101            BE-data                         low
  101110            Premium-data                    high
  101111            Premium-data                    low
  110000            NC                              low
  110001            BE-data                         low
  110010            BE-data                         low
  110011            BE-data                         low
  110100            BE-data                         low
  110101            BE-data                         low
  110110            BE-data                         low
  110111            BE-data                         low
  111000            NC                              low
  111001            BE-data                         low
  111010            BE-data                         low
  111011            BE-data                         low
  111100            BE-data                         low
  111101            BE-data                         low
  111110            BE-data                         low
  111111            BE-data                         low
```

Meaning Notice that the default classifier is incorporated into the customer classifier. If you were to remove the **import default** statement from the custom classifier, the custom classifier would look like this:

```
user@R2> show class-of-service classifier name v4-ba-classifier
Classifier: v4-ba-classifier, Code point type: dscp, Index: 10755
  Code point         Forwarding class                  Loss priority
  000000             BE-data                           high
  000001             BE-data                           low
  101110             Premium-data                      high
  101111             Premium-data                      low
```

Verifying the Forwarding Classes and Output Queues

Purpose Make sure that the forwarding classes are configured as expected.

Action On Device R2, run the **show class-of-service forwarding-class** command.

```
user@R2> show class-of-service forwarding-class

Forwarding class                          ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  BE-data                                  0       0          0             low
            normal            low
  Premium-data                             1       1          1             low
            normal            low
  Voice                                    2       2          2             low
            normal            low
  NC                                       3       3          3             low
            normal            low
```

Meaning The forwarding classes are configured as expected.

Copyright © 2016, Juniper Networks, Inc.

Verifying That the Classifier Is Applied to the Interfaces

Purpose Make sure that the classifier is applied to the correct interfaces.

Action On Device R2, run the **show class-of-service interface** command.

```
user@R2> show class-of-service interface ge-1/0/3

Physical interface: ge-1/0/3, Index: 144
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: ge-1/0/3.0, Index: 333
Object                  Name                    Type                    Index
Classifier              v4-ba-classifier        dscp                    10755

user@R2> show class-of-service interface ge-1/0/9

Physical interface: ge-1/0/9, Index: 150
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: ge-1/0/9.0, Index: 332
Object                  Name                    Type                    Index
Classifier              v4-ba-classifier        dscp                    10755
```

Meaning The interfaces are configured as expected.

Verifying Behavior Aggregate Classifiers

Purpose Verify that the behavior aggregate classifiers were configured properly on the device.

Action From configuration mode, enter the **show class-of-service** command.

When you are using hping to set the DSCP code points in the IPv4 packet header, the type-of-service (ToS) hex value (in this case, BC) is required in the **--tos** option of the **hping** command.

If your binary-to-hex or binary-to-decimal conversion skills are rusty, you can use an online calculator, such as http://www.mathsisfun.com/binary-decimal-hexadecimal-converter.html .

NOTE: When you convert a binary DSCP code point value, be sure to add two extra zeros at the end. So instead of 101111, use 10111100. These 0 values (the 7th and 8th bits) are reserved and ignored, but if you do not include them in the conversion, your hex and decimal values will be incorrect.

Extended Ping Sent from Device R1

```
user@R1> ping 172.16.70.1 tos 188 rapid count 25
```


Copyright © 2016, Juniper Networks, Inc.

```
PING 172.16.70.1 (172.16.70.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!
--- 172.16.70.1 ping statistics ---
25 packets transmitted, 25 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.404/0.483/1.395/0.207 ms
```

hping Sent from Host 1

```
root@host1> hping 172.16.70.1 --tos BC -c 25

HPING 172.16.70.1 (eth1 172.16.70.1): NO FLAGS are set, 40 headers + 0 data bytes
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=0 win=0 rtt=0.3 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=1 win=0 rtt=0.6 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=2 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=3 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=4 win=0 rtt=0.6 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=5 win=0 rtt=0.3 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=6 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=7 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=8 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=9 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=10 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=11 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=12 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=13 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=14 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=15 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=16 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=17 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=18 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=19 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=20 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=21 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=22 win=0 rtt=0.4 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=23 win=0 rtt=0.5 ms
len=46 ip=172.16.70.1 ttl=61 DF id=0 sport=0 flags=RA seq=24 win=0 rtt=0.4 ms
```

On Device R2, Verify that Queue 2 is Incrementing.

Code point 101111 is associated with Premium-data, which uses queue 1.

```
user@R2> show interfaces extensive ge-1/0/3 | find "queue counters"


Queue counters:   Queued packets   Transmitted packets   Dropped packets
  0                           0                     0                 0
  1                          50                    50                 0
  2                           0                     0                 0
  3                          42                    42                 0
  Queue number:         Mapped forwarding classes
    0                   BE-data
    1                   Premium-data
    2                   Voice
    3                   NC
...
```

Meaning The output shows that queue 1 has incremented by 50 packets after sending 50 packets through the router.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *Interfaces Feature Guide for Security Devices*
- Classification Overview on page 1411
- Sample Behavior Aggregate Classification on page 1416
- Understanding Packet Loss Priorities on page 1414

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 67

Controlling Network Access with Traffic Policing

- Simple Filters and Policers Overview on page 1427
- Two-Rate Three-Color Policer Overview on page 1428
- Example: Configuring a Two-Rate Three-Color Policer on page 1429
- Logical Interface (Aggregate) Policer Overview on page 1434
- Two-Color Policer Configuration Overview on page 1434
- Example: Configuring a Two-Color Logical Interface (Aggregate) Policer on page 1436
- Guidelines for Configuring Simple Filters on page 1442
- Example: Configuring and Applying a Firewall Filter for a Multifield Classifier on page 1445

Simple Filters and Policers Overview

You can configure simple filters and policers to handle oversubscribed traffic in SRX5600 and SRX5800 devices. In Junos OS, policers can be configured as part of the firewall filter hierarchy.

NOTE: For SRX5600 and SRX5800 devices, the simple filter or policing actions can be applied only to logical interfaces residing in an SRX5000 line Flex IOC (FIOC) because only an SRX5000 line FIOC supports the simple filter and policing features on the SRX5600 and SRX5800 devices.

The simple filter functionality consists of the following:

- Classifying packets according to configured policies
- Taking appropriate actions based on the results of classification

In Junos OS, ingress traffic policers can limit the rate of incoming traffic. Two main reasons to use traffic policing are:

- To enforce traffic rates to conform to the service-level agreement (SLA)
- To protect next hops, such as protecting the central point and the SPU from being overwhelmed by excess traffic like DOS attacks

Copyright © 2016, Juniper Networks, Inc.

Using the results of packet classification and traffic metering, a policer can take one of the following actions for a packet: forward a conforming (green) packet or drop a nonconforming (yellow) packet. Policers always discard a nonconforming red packet. Traffic metering supports the algorithm of the two-rate tricolor marker (TCM). (See RFC 2698, A *Two Rate Three Color Marker*.)

Related Documentation

- Guidelines for Configuring Simple Filters on page 1442
- Example: Configuring a Two-Rate Three-Color Policer on page 1429
- Example: Configuring and Applying a Firewall Filter for a Multifield Classifier on page 1445

Two-Rate Three-Color Policer Overview

A two-rate three-color policer defines two bandwidth limits (one for guaranteed traffic and one for peak traffic) and two burst sizes (one for each of the bandwidth limits). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

Two-rate three-color policing meters a traffic stream based on the following configured traffic criteria:

- Committed information rate (CIR)—Bandwidth limit for guaranteed traffic.
- Committed burst size (CBS)—Maximum packet size permitted for bursts of data that exceed the CIR.
- Peak information rate (PIR)—Bandwidth limit for peak traffic.
- Peak burst size (PBS)—Maximum packet size permitted for bursts of data that exceed the PIR.

Two-rate tricolor marking (two-rate TCM) classifies traffic as belonging to one of three color categories and performs congestion-control actions on the packets based on the color marking:

- Green—Traffic that conforms to the bandwidth limit and burst size for guaranteed traffic (CIR and CBS). For a green traffic flow, two-rate TCM marks the packets with an implicit loss priority of **low** and transmits the packets.
- Yellow—Traffic that exceeds the bandwidth limit or burst size for guaranteed traffic (CIR or CBS) but not the bandwidth limit and burst size for peak traffic (PIR and PBS). For a yellow traffic flow, two-rate TCM marks packets with an implicit loss priority of **medium-high** and transmits the packets.
- Red—Traffic that exceeds the bandwidth limit and burst size for peak traffic (PIR and PBS). For a red traffic flow, two-rate TCM marks packets with an implicit loss priority of **high** and, optionally, discards the packets.

If congestion occurs downstream, the packets with higher loss priority are more likely to be discarded.

Copyright © 2016, Juniper Networks, Inc.

NOTE: For both single-rate and two-rate three-color policers, the only *configurable* action is to discard packets in a red traffic flow.

For a tricolor marking policer referenced by a firewall filter term, the **discard** policing action is supported on the following routing platforms:

- EX Series switches
- M7i and M10i routers with the Enhanced CFEB (CFEB-E)
- M120 and M320 routers with Enhanced-III FPCs
- MX Series routers with Trio MPCs

To apply a tricolor marking policer on these routing platforms, it is not necessary to include the **logical-interface-policer** statement.

Related Documentation

- Example: Configuring a Two-Rate Three-Color Policer on page 1429

Example: Configuring a Two-Rate Three-Color Policer

This example shows how to configure a two-rate three-color policer.

- Requirements on page 1429
- Overview on page 1429
- Configuration on page 1430
- Verification on page 1433

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

A two-rate three-color policer meters a traffic flow against a bandwidth limit and burst-size limit for guaranteed traffic, plus a bandwidth limit and burst-size limit for peak traffic. Traffic that conforms to the limits for guaranteed traffic is categorized as green, and nonconforming traffic falls into one of two categories:

- Nonconforming traffic that does not exceed peak traffic limits is categorized as yellow.
- Nonconforming traffic that exceeds peak traffic limits is categorized as red.

Each category is associated with an action. For green traffic, packets are implicitly set with a loss-priority value of **low** and then transmitted. For yellow traffic, packets are implicitly set with a loss-priority value of **medium-high** and then transmitted. For red traffic, packets are implicitly set with a loss-priority value of **high** and then transmitted. If the policer configuration includes the optional **action** statement (**action loss-priority high then discard**), then packets in a red flow are discarded instead.

Copyright © 2016, Juniper Networks, Inc.

You can apply a three-color policer to Layer 3 traffic as a firewall filter policer only. You reference the policer from a stateless firewall filter term, and then you apply the filter to the input or output of a logical interface at the protocol level.

Topology

In this example, you apply a color-aware, two-rate three-color policer to the input IPv4 traffic at logical interface **ge-2/0/5.0**. The IPv4 firewall filter term that references the policer does not apply any packet-filtering. The filter is used only to apply the three-color policer to the interface.

You configure the policer to rate-limit traffic to a bandwidth limit of 40 Mbps and a burst-size limit of 100 KB for green traffic, and you configure the policer to also allow a peak bandwidth limit of 60 Mbps and a peak burst-size limit of 200 KB for yellow traffic. Only nonconforming traffic that exceeds the peak traffic limits is categorized as red. In this example, you configure the three-color policer action **loss-priority high then discard**, which overrides the implicit marking of red traffic to a **high** loss priority.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure this example, perform the following tasks:

- Configuring a Two-Rate Three-Color Policer on page 1430
- Configuring an IPv4 Stateless Firewall Filter That References the Policer on page 1432
- Applying the Filter to a Logical Interface at the Protocol Family Level on page 1432

CLI Quick Configuration

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

```
set firewall three-color-policer trTCM1-ca two-rate color-aware
set firewall three-color-policer trTCM1-ca two-rate committed-information-rate 40m
set firewall three-color-policer trTCM1-ca two-rate committed-burst-size 100k
set firewall three-color-policer trTCM1-ca two-rate peak-information-rate 60m
set firewall three-color-policer trTCM1-ca two-rate peak-burst-size 200k
set firewall three-color-policer trTCM1-ca action loss-priority high then discard
set firewall family inet filter filter-trtcm1ca-all term 1 then three-color-policer two-rate
    trTCM1-ca
set interfaces ge-2/0/5 unit 0 family inet address 10.10.10.1/30
set interfaces ge-2/0/5 unit 0 family inet filter input filter-trtcm1ca-all
set class-of-service interfaces ge-2/0/5 forwarding-class af
```

Configuring a Two-Rate Three-Color Policer

Step-by-Step Procedure

To configure a two-rate three-color policer:

1. Enable configuration of a three-color policer.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set firewall three-color-policer trTCM1-ca
```

2. Configure the color mode of the two-rate three-color policer.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate color-aware
```

3. Configure the two-rate guaranteed traffic limits.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate committed-information-rate 40m
user@host# set two-rate committed-burst-size 100k
```

   Traffic that does not exceed both of these limits is categorized as green. Packets in a green flow are implicitly set to **low** loss priority and then transmitted.

4. Configure the two-rate peak traffic limits.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate peak-information-rate 60m
user@host# set two-rate peak-burst-size 200k
```

   Nonconforming traffic that does not exceed both of these limits is categorized as yellow. Packets in a yellow flow are implicitly set to **medium-high** loss priority and then transmitted. Nonconforming traffic that exceeds both of these limits is categorized as red. Packets in a red flow are implicitly set to **high** loss priority.

5. (Optional) Configure the policer action for red traffic.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set action loss-priority high then discard
```

   For three-color policers, the only configurable action is to discard red packets. Red packets are packets that have been assigned high loss priority because they exceeded the peak information rate (PIR) and the peak burst size (PBS).

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
three-color-policer trTCM1-ca {
    action {
        loss-priority high then discard;
    }
    two-rate {
        color-aware;
        committed-information-rate 40m;
        committed-burst-size 100k;
        peak-information-rate 60m;
        peak-burst-size 200k;
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

Configuring an IPv4 Stateless Firewall Filter That References the Policer

Step-by-Step Procedure

To configure an IPv4 stateless firewall filter that references the policer:

1. Enable configuration of an IPv4 standard stateless firewall filter.

   ```
   [edit]
   user@host# set firewall family inet filter filter-trtcm1ca-all
   ```

2. Specify the filter term that references the policer.

   ```
   [edit firewall family inet filter filter-trtcm1ca-all]
   user@host# set term 1 then three-color-policer two-rate trTCM1-ca
   ```

   Note that the term does not specify any match conditions. The firewall filter passes all packets to the policer.

Results

Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
   filter filter-trtcm1ca-all {
      term 1 {
         then {
            three-color-policer {
               two-rate trTCM1-ca;
            }
         }
      }
   }
}
three-color-policer trTCM1-ca {
   action {
      loss-priority high then discard;
   }
   two-rate {
      color-aware;
      committed-information-rate 40m;
      committed-burst-size 100k;
      peak-information-rate 60m;
      peak-burst-size 200k;
   }
}
```

Applying the Filter to a Logical Interface at the Protocol Family Level

Step-by-Step Procedure

To apply the filter to the logical interface at the protocol family level:

1. Enable configuration of an IPv4 firewall filter.

   ```
   [edit]
   user@host# edit interfaces ge-2/0/5 unit 0 family inet
   ```


Copyright © 2016, Juniper Networks, Inc.

2. Apply the policer to the logical interface at the protocol family level.

```
[edit interfaces ge-2/0/5 unit 0 family inet]
user@host# set address 10.10.10.1/30
user@host# set filter input filter-trtcm1ca-all
```

3. (MX Series routers and EX Series switches only) (Optional) For input policers, you can configure a fixed classifier. A fixed classifier reclassifies all incoming packets, regardless of any preexisting classification.

```
[edit]
user@host# set class-of-service interfaces ge-2/0/5 forwarding-class af
```

The classifier name can be a configured classifier or one of the default classifiers.

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-2/0/5 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
      filter {
        input filter-trtcm1ca-all;
      }
    }
  }
}
```

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

Verification

Displaying the Firewall Filters Applied to the Logical Interface

Purpose Verify that the firewall filter is applied to IPv4 input traffic at the logical interface.

Action Use the **show interfaces** operational mode command for the logical interface **ge-2/0/5.0**, and specify **detail** mode. The **Protocol inet** section of the command output displays IPv4 information for the logical interface. Within that section, the **Input Filters** field displays the name of IPv4 firewall filters associated with the logical interface.

```
user@host> show interfaces ge-2/0/5.0 detail
  Logical interface ge-2/0/5.0 (Index 105) (SNMP ifIndex 556) (Generation 170)
    Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

Copyright © 2016, Juniper Networks, Inc.

```
 Output packets:                    0
Transit statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Protocol inet, MTU: 1500, Generation: 242, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Input Filters: filter-trtcm1ca-all
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.20.130/24, Local: 10.20.130.1, Broadcast: 10.20.130.255,

    Generation: 171
Protocol multiservice, MTU: Unlimited, Generation: 243, Route table: 0
  Policer: Input: __default_arp_policer__
```

Related Documentation

- Two-Rate Three-Color Policer Overview on page 1428

Logical Interface (Aggregate) Policer Overview

A *logical interface policer*—also called an *aggregate policer*—is a two-color or three-color policer that defines traffic rate limiting. You can apply a policer to input or output traffic for multiple protocol families on the same logical interface without needing to create multiple instances of the policer.

To configure a single-rate two-color logical interface policer, include the **logical-interface-policer** statement at the **[edit firewall policer policer-name]** hierarchy level.

You apply a logical interface policer to Layer 3 traffic directly to the interface configuration at the protocol family level (to rate-limit traffic of a specific protocol family). You cannot reference a logical interface policer from a stateless firewall filter term and then apply the filter to a logical interface.

You can apply a logical interface policer to unicast traffic only. .

To display a logical interface policer on a particular interface, issue the **show interfaces policers** operational mode command.

Related Documentation

- Two-Color Policer Configuration Overview on page 1434
- Example: Configuring a Two-Color Logical Interface (Aggregate) Policer on page 1436
- **logical-interface-policer on page 1639**

Two-Color Policer Configuration Overview

Table 123 describes the hierarchy levels at which you can configure and apply single-rate two-color policers to Layer 3 traffic. For information about applying single-rate two-color policers to Layer 2 traffic, see *Two-Color Policing at Layer 2 Overview*.

Copyright © 2016, Juniper Networks, Inc.

Table 123: Two-Color Policer Configuration and Application Overview

Policer Configuration	Layer 3 Application	Key Points
Single-Rate Two-Color Policer *Defines traffic rate limiting that you can apply to Layer 3 protocol-specific traffic at a logical interface. Can be applied as an interface policer or as a firewall filter policer.*		

Basic policer configuration:
 	
[edit firewall]
policer *policer-name* {
  if-exceeding {
  bandwidth-limit *bps*;
  burst-size-limit *bytes*;
  }
  then {
  discard;
  forwarding-class *class-name*;
  loss-priority *supported-value*;
  }
}
	Method A—Apply as an interface policer at the protocol family level:
	
	[edit interfaces]
	interface-name {
	  unit *unit-number* {
	  family *family-name* {
	  policer {
	  input *policer-name*;
	  output *policer-name*;
	  }
	  }
	  }
	}
	
	Method B—Apply as a firewall filter policer at the protocol family level:
	
	[edit firewall]
	family *family-name* {
	  filter *filter-name* {
	  interface-specific; # (*)
	  from {
	  ... *match-conditions* ...
	  }
	  then {
	  policer *policer-name*;
	  }
	  }
	}
	
	[edit interfaces]
	interface-name {
	  unit *unit-number* {
	  family *family-name* {
	  filter {
	  input *filter-name*;
	  output *filter-name*;
	  }
	  ... *protocol-configuration* ...
	  }
	  }
	}
		Policer configuration:
		
		• Use **bandwidth-limit *bps*** to specify an absolute value.
		
		Firewall filter configuration (*)
		
		• If applying to multiple interfaces, include the **interface-specific** statement to create unique policers and counters for each interface.
		
		Interface policer verification:
		
		• Use the **show interfaces (detail | extensive)** operational mode command.
		• Use the **show policer** operational mode command.
		
		Firewall filter policer verification:
		
		• Use the **show interfaces (detail | extensive)** operational mode command.
		• Use the **show firewall filter *filter-name*** operational mode command.

Copyright © 2016, Juniper Networks, Inc.

Table 123: Two-Color Policer Configuration and Application Overview *(continued)*

Policer Configuration	Layer 3 Application	Key Points
Logical Interface (Aggregate) Policer *Defines traffic rate limiting that you can apply to multiple protocol families on the same logical interface without creating multiple instances of the policer. Can be applied directly to a logical interface configuration only.*		

Logical interface policer configuration:
[edit firewall]
policer *policer-name* {
logical-interface-policer;
if-exceeding {
bandwidth-limit *bps*;
burst-size-limit *bytes*;
}
then {
discard;
forwarding-class *class-name*;
loss-priority *supported-value*;
}
}
	Method A—Apply as an interface policer only:
	[edit interfaces]
	interface-name {
	unit *unit-number* {
	policer { # All protocols
	input *policer-name*;
	output *policer-name*;
	}
	family *family-name* {
	policer { # One protocol
	input *policer-name*;
	output *policer-name*;
	}
	}
	}
	}
	Method B—Apply as a firewall filter policer at the protocol family level:
	[edit firewall]
	family *family-name* {
	filter *filter-name* {
	interface-specific;
	term *term-name*{
	from {
	... *match-conditions* ...
	}
	}
	then {
	policer *policer-name*;
	}
	}
	}
		Policer configuration:
		• Include the **logical-interface-policer** statement.
		Two options for interface policer application:
		• To rate-limit all traffic types, regardless of the protocol family, apply the logical interface policer at the logical unit level.
		• To rate-limit traffic of a specific protocol family, apply the logical interface policer at the protocol family level.
		Interface policer verification:
		• Use the **show interfaces (detail | extensive)** operational mode command.
		• Use the **show policer** operational mode command.

Related Documentation

- logical-interface-policer on page 1639
- Example: Configuring a Two-Color Logical Interface (Aggregate) Policer on page 1436

Example: Configuring a Two-Color Logical Interface (Aggregate) Policer

This example shows how to configure a single-rate two-color policer as a logical interface policer and apply it to incoming IPv4 traffic on a logical interface.

- Requirements on page 1437
- Overview on page 1437

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1437
- Verification on page 1441

Requirements

Before you begin, make sure that the logical interface to which you apply the two-color logical interface policer is hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**).

Overview

In this example, you configure the single-rate two-color policer **policer_IFL** as a logical interface policer and apply it to incoming IPv4 traffic at logical interface **ge-1/3/1.0**.

Topology

If the input IPv4 traffic on the physical interface **ge-1/3/1** exceeds the bandwidth limit equal to 90 percent of the media rate with a 300 KB burst-size limit, then the logical interface policer **policer_IFL** rate-limits the input IPv4 traffic on the logical interface **ge-1/3/1.0**. Configure the policer to mark nonconforming traffic by setting packet loss priority (PLP) levels to **high** and classifying packets as **best-effort**.

As the incoming IPv4 traffic rate on the physical interface slows and conforms to the configured limits, Junos OS stops marking the incoming IPv4 packets at the logical interface.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure this example, perform the following tasks:

- Configuring the Logical Interfaces on page 1438
- Configuring the Single-Rate Two-Color Policer as a Logical Interface Policer on page 1439
- Applying the Logical Interface Policer to Input IPv4 Traffic at a Logical Interface on page 1440

CLI Quick Configuration

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

```
set interfaces ge-1/3/1 vlan-tagging
set interfaces ge-1/3/1 unit 0 vlan-id 100
set interfaces ge-1/3/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/3/1 unit 1 vlan-id 101
set interfaces ge-1/3/1 unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
    00:00:11:22:33:44
set firewall policer policer_IFL logical-interface-policer
set firewall policer policer_IFL if-exceeding bandwidth-percent 90
set firewall policer policer_IFL if-exceeding burst-size-limit 300k
set firewall policer policer_IFL then loss-priority high
```

Copyright © 2016, Juniper Networks, Inc.

```
set firewall policer policer_IFL then forwarding-class best-effort
set interfaces ge-1/3/1 unit 0 family inet policer input policer_IFL
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/3/1
   ```

2. Configure single tagging.

   ```
   [edit interfaces ge-1/3/1]
   user@host# set vlan-tagging
   ```

3. Configure logical interface **ge-1/3/1.0**.

   ```
   [edit interfaces ge-1/3/1]
   user@host# set unit 0 vlan-id 100
   user@host# set unit 0 family inet address 10.10.10.1/30
   ```

4. Configure logical interface **ge-1/3/1.0**.

   ```
   [edit interfaces ge-1/3/1]
   user@host# set unit 1 vlan-id 101
   user@host# set unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
     00:00:11:22:33:44
   ```

Results

Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
   vlan-tagging;
   unit 0 {
      vlan-id 100;
      family inet {
         address 10.10.10.1/30;
      }
   }
   unit 1 {
      vlan-id 101;
      family inet {
         address 20.20.20.1/30 {
            arp 20.20.20.2 mac 00:00:11:22:33:44;
         }
      }
   }
}
```


Copyright © 2016, Juniper Networks, Inc.

Configuring the Single-Rate Two-Color Policer as a Logical Interface Policer

Step-by-Step Procedure

To configure a single-rate two-color policer as a logical interface policer:

1. Enable configuration of a single-rate two-color policer.

   ```
   [edit]
   user@host# edit firewall policer policer_IFL
   ```

2. Specify that the policer is a logical interface (aggregate) policer.

   ```
   [edit firewall policer policer_IFL]
   user@host# set logical-interface-policer
   ```

   A logical interface policer rate-limits traffic based on a percentage of the media rate of the physical interface underlying the logical interface to which the policer is applied. The policer is applied directly to the interface rather than referenced by a firewall filter.

3. Specify the policer traffic limits.

   a. Specify the bandwidth limit.

      - To specify the bandwidth limit as an absolute rate, from 8,000 bits per second through 50,000,000,000 bits per second, include the **bandwidth-limit** ***bps*** statement.
      - To specify the bandwidth limit as a percentage of the physical port speed on the interface, include the **bandwidth-percent** ***percent*** statement.

      In this example, the CLI commands and output are based on a bandwidth limit specified as a percentage rather than as an absolute rate.

      ```
      [edit firewall policer policer_IFL]
      user@host# set if-exceeding bandwidth-percent 90
      ```

   b. Specify the burst-size limit, from 1,500 bytes through 100,000,000,000 bytes, which is the maximum packet size to be permitted for bursts of data that exceed the specified bandwidth limit.

      ```
      [edit firewall policer policer_IFL]
      user@host# set if-exceeding burst-size-limit 300k
      ```

4. Specify the policer actions to be taken on traffic that exceeds the configured rate limits.

   - To discard the packet, include the **discard** statement.
   - To set the loss-priority value of the packet, include the **loss-priority (low | medium-low | medium-high | high)** statement.
   - To classify the packet to a forwarding class, include the **forwarding-class (*forwarding-class* | assured-forwarding | best-effort | expedited-forwarding | network-control)** statement.

Copyright © 2016, Juniper Networks, Inc.

In this example, the CLI commands and output are based on both setting the packet loss priority level and classifying the packet.

```
[edit firewall policer policer_IFL]
user@host# set then loss-priority high
user@host# set then forwarding-class best-effort
```

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer policer_IFL {
  logical-interface-policer;
  if-exceeding {
    bandwidth-percent 90;
    burst-size-limit 300k;
  }
  then {
    loss-priority high;
    forwarding-class best-effort;
  }
}
```

Applying the Logical Interface Policer to Input IPv4 Traffic at a Logical Interface

Step-by-Step Procedure To apply the two-color logical interface policer to input IPv4 traffic a logical interface:

1. Enable configuration of the logical interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/3/1 unit 0
   ```

2. Apply the policer to all traffic types or to a specific traffic type on the logical interface.

   - To apply the policer to all traffic types, regardless of the protocol family, include the **policer (input | output)** ***policer-name*** statement at the **[edit interfaces** ***interface-name*** **unit** ***number*****]** hierarchy level.
   - To apply the policer to traffic of a specific protocol family, include the **policer (input | output)** ***policer-name*** statement at the **[edit interfaces** ***interface-name*** **unit** ***unit-number*** **family** ***family-name*****]** hierarchy level.

   To apply the logical interface policer to incoming packets, use the **policer input** ***policer-name*** statement. To apply the logical interface policer to outgoing packets, use the **policer output** ***policer-name*** statement.

   In this example, the CLI commands and output are based on rate-limiting the IPv4 input traffic at logical interface **ge-1/3/1.0**.

   ```
   [edit interfaces ge-1/3/1 unit 0]
   user@host# set family inet policer input policer_IFL
   ```


Copyright © 2016, Juniper Networks, Inc.

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
   vlan-tagging;
   unit 0 {
      vlan-id 100;
      family inet {
         policer input policer_IFL;
         address 10.10.10.1/30;
      }
   }
   unit 1 {
      vlan-id 101;
      family inet {
         address 20.20.20.1/30 {
            arp 20.20.20.2 mac 00:00:11:22:33:44;
         }
      }
   }
}
```

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

Verification

Confirm that the configuration is working properly.

- Displaying Traffic Statistics and Policers for the Logical Interface on page 1441
- Displaying Statistics for the Policer on page 1442

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluating packets received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interface **ge-1/3/1.0**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface. The **Protocol inet** subsection contains a **Policer** field that would list the policer **policer_IFL** as an input or output logical interface policer as follows:

- **Input: policer_IFL-ge-1/3/1.0-log_int-i**
- **Output: policer_IFL-ge-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

Copyright © 2016, Juniper Networks, Inc.

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the **show policer** operational mode command and, optionally, specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **policer_IFL**, the input and output policer names are displayed as follows:

- **policer_IFL-ge-1/3/1.0-log_int-i**
- **policer_IFL-ge-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

Related Documentation

- logical-interface-policer on page 1639
- Two-Color Policer Configuration Overview on page 1434

Guidelines for Configuring Simple Filters

This topic covers the following information:

- Statement Hierarchy for Configuring Simple Filters on page 1442
- Simple Filter Protocol Families on page 1443
- Simple Filter Names on page 1443
- Simple Filter Terms on page 1443
- Simple Filter Match Conditions on page 1443
- Simple Filter Terminating Actions on page 1445
- Simple Filter Nonterminating Actions on page 1445

Statement Hierarchy for Configuring Simple Filters

To configure a simple filter, include the **simple-filter** ***simple-filter-name*** statement at the **[edit firewall family inet]** hierarchy level.

```
[edit]
firewall {
   family inet {
      simple-filter  simple-filter-name {
         term term-name {
            from {
               match-conditions;
            }
            then {
               actions;
            }
         }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
```

Individual statements supported under the **simple-filter *simple-filter-name*** statement are described separately in this topic and are illustrated in the example of configuring and applying a simple filter.

Simple Filter Protocol Families

You can configure simple filters to filter IPv4 traffic (**family inet**) only. No other protocol family is supported for simple filters.

NOTE: You can apply simple filters to the family inet only, and only in the input direction. Because of hardware limitations on the SRX5600 and SRX5800 devices, a maximum of 400 logical input interfaces and 2000 terms (in one Broadcom packet processor) can be applied with simple filters.

Simple Filter Names

Under the **family inet** statement, you can include **simple-filter *simple-filter-name*** statements to create and name simple filters. The filter name can contain letters, numbers, and hyphens (-) and be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ ”).

Simple Filter Terms

Under the **simple-filter *simple-filter-name*** statement, you can include **term *term-name*** statements to create and name filter terms.

- You must configure at least one term in a firewall filter.
- You must specify a unique name for each term within a firewall filter. The term name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ ”).
- The order in which you specify terms within a firewall filter configuration is important. Firewall filter terms are evaluated in the order in which they are configured. By default, new terms are always added to the end of the existing filter. You can use the **insert** configuration mode command to reorder the terms of a firewall filter.

Simple filters do *not* support the **next term** action.

NOTE: In one Broadcom packet processor, a maximum of 2000 terms can be applied with simple filters on the SRX5600 and SRX5800 devices.

Simple Filter Match Conditions

Simple filter terms support only a subset of the IPv4 match conditions that are supported for standard stateless firewall filters.

Copyright © 2016, Juniper Networks, Inc.

Unlike standard stateless firewall filters, the following restrictions apply to simple filters:

- On MX Series routers with the Enhanced Queuing DPC, simple filters do *not* support the **forwarding- class** match condition.
- Simple filters support only one **source-address** and one **destination-address** prefix for each filter term. If you configure multiple prefixes, only the last one is used.
- Simple filters do *not* support multiple source addresses and destination addresses in a single term. If you configure multiple addresses, only the last one is used.
- Simple filters do *not* support negated match conditions, such as the **protocol-except** match condition or the **exception** keyword.
- Simple filters support a range of values for **source-port** and **destination-port** match conditions only. For example, you can configure **source-port 400-500** or **destination-port 600-700**.
- Simple filters do *not* support noncontiguous mask values.

Table 124 lists the simple filter match conditions.

Table 124: Simple Filter Match Conditions

Match Condition	Description
destination-address *destination-address*	Match IP destination address.
destination-port *number*	TCP or UDP destination port field. If you configure this match condition, we recommend that you also configure the **protocol** match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text aliases (the port numbers are also listed): **afs** (1483), **bgp** (179), **biff** (512), **bootpc** (68), **bootps** (67), **cmd** (514), **cvspserver** (2401), **dhcp** (67), **domain** (53), **eklogin** (2105), **ekshell** (2106), **exec** (512), **finger** (79), **ftp** (21), **ftp-data** (20), **http** (80), **https** (443), **ident** (113), **imap** (143), **kerberos-sec** (88), **klogin** (543), **kpasswd** (761), **krb-prop** (754), **krbupdate** (760), **kshell** (544), **ldap** (389), **login** (513), **mobileip-agent** (434), **mobilip-mn** (435), **msdp** (639), **netbios-dgm** (138), **netbios-ns** (137), **netbios-ssn** (139), **nfsd** (2049), **nntp** (119), **ntalk** (518), **ntp** (123), **pop3** (110), **pptp** (1723), **printer** (515), **radacct** (1813), **radius** (1812), **rip** (520), **rkinit** (2108), **smtp** (25), **snmp** (161), **snmptrap** (162), **snpp** (444), **socks** (1080), **ssh** (22), **sunrpc** (111), **syslog** (514), **tacacs-ds** (65), **talk** (517), **telnet** (23), **tftp** (69), **timed** (525), **who** (513), or **xdmcp** (177).
forwarding-class *class*	Match the forwarding class of the packet. Specify **assured-forwarding**, **best-effort**, **expedited-forwarding**, or **network-control**.
protocol *number*	IP protocol field. In place of the numeric value, you can specify one of the following text aliases (the field values are also listed): **ah** (51), **dstopts** (60), **egp** (8), **esp** (50), **fragment** (44), **gre** (47), **hop-by-hop** (0), **icmp** (1), **icmp6** (58), **icmpv6** (58), **igmp** (2), **ipip** (4), **ipv6** (41), **ospf** (89), **pim** (103), **rsvp** (46), **sctp** (132), **tcp** (6), **udp** (17), or **vrrp** (112).
source-address *ip-source-address*	Match the IP source address.

Copyright © 2016, Juniper Networks, Inc.

Table 124: Simple Filter Match Conditions *(continued)*

Match Condition	Description
source-port *number*	Match the UDP or TCP source port field. If you configure this match condition, we recommend that you also configure the **protocol** match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text aliases listed for **destination-port**.

Simple Filter Terminating Actions

Simple filters do *not* support explicitly configurable terminating actions, such as **accept**, **reject**, and **discard**. Terms configured in a simple filter always accept packets.

Simple filters do *not* support the **next** action.

Simple Filter Nonterminating Actions

Simple filters support only the following nonterminating actions:

- **forwarding-class (*forwarding-class* | assured-forwarding |best-effort | expedited-forwarding | network-control)**

  NOTE: On the MX Series routers with the Enhanced Queuing DPC, the forwarding class is not supported as a **from** match condition.

- **loss-priority (high | low | medium-high | medium-low)**

Simple filters do not support actions that perform other functions on a packet (such as incrementing a counter, logging information about the packet header, sampling the packet data, or sending information to a remote host using the system log functionality).

Related Documentation

- Simple Filters and Policers Overview on page 1427

Example: Configuring and Applying a Firewall Filter for a Multifield Classifier

This example shows how to configure a firewall filter to classify traffic using a multifield classifier. The classifier detects packets of interest to CoS as they arrive on an interface.

- Requirements on page 1445
- Overview on page 1446
- Configuration on page 1447
- Verification on page 1449

Requirements

To verify this procedure, this example uses a traffic generator. The traffic generator can be hardware-based or it can be software running on a server or host machine.

Copyright © 2016, Juniper Networks, Inc.

The functionality in this procedure is widely supported on devices that run Junos OS. The example shown here was tested and verified on MX Series routers running Junos OS Release 10.4.

Overview

A classifier is a software operation that inspects a packet as it enters the router or switch. The packet header contents are examined, and this examination determines how the packet is treated when the network becomes too busy to handle all of the packets and you want your devices to drop packets intelligently, instead of dropping packets indiscriminately. One common way to detect packets of interest is by source port number. The TCP port numbers 80 and 12345 are used in this example, but many other matching criteria for packet detection are available to multifield classifiers, using firewall filter match conditions. The configuration in this example specifies that TCP packets with source port 80 are classified into the BE-data forwarding class and queue number 0. TCP packets with source port 12345 are classified into the Premium-data forwarding class and queue number 1.

Multifield classifiers are typically used at the network edge as packets enter an autonomous system (AS).

In this example, you configure the firewall filter mf-classifier and specify some custom forwarding classes on Device R1. In specifying the custom forwarding classes, you also associate each class with a queue.

The classifier operation is shown in Figure 73.

Figure 73: Multifield Classifier Based on TCP Source Ports

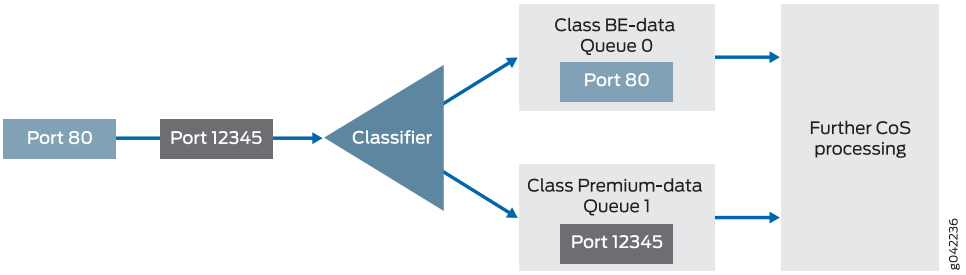

You apply the multifield classifier's firewall filter as an input filter on each customer-facing or host-facing interface that needs the filter. The incoming interface is ge-1/0/0 on Device R1. The classification and queue assignment is verified on the outgoing interface. The outgoing interface is Device R1's ge-1/0/2 interface.

Topology

Figure 74 shows the sample network.

Copyright © 2016, Juniper Networks, Inc.

Figure 74: Multifield Classifier Scenario

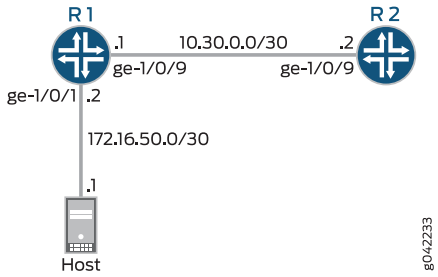

"CLI Quick Configuration" on page 1447 shows the configuration for all of the Juniper Networks devices in Figure 74.

The section "Step-by-Step Procedure" on page 1447 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```
set interfaces ge-1/0/0 description to-host
set interfaces ge-1/0/0 unit 0 family inet filter input mf-classifier
set interfaces ge-1/0/0 unit 0 family inet address 172.16.50.2/30
set interfaces ge-1/0/2 description to-R2
set interfaces ge-1/0/2 unit 0 family inet address 10.30.0.1/30
set class-of-service forwarding-classes class BE-data queue-num 0
set class-of-service forwarding-classes class Premium-data queue-num 1
set class-of-service forwarding-classes class Voice queue-num 2
set class-of-service forwarding-classes class NC queue-num 3
set firewall family inet filter mf-classifier term BE-data from protocol tcp
set firewall family inet filter mf-classifier term BE-data from port 80
set firewall family inet filter mf-classifier term BE-data then forwarding-class BE-data
set firewall family inet filter mf-classifier term Premium-data from protocol tcp
set firewall family inet filter mf-classifier term Premium-data from port 12345
set firewall family inet filter mf-classifier term Premium-data then forwarding-class
    Premium-data
set firewall family inet filter mf-classifier term accept-all-else then accept
```

Device R2

```
set interfaces ge-1/0/2 description to-R1
set interfaces ge-1/0/2 unit 0 family inet address 10.30.0.2/30
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the device interfaces.

```
[edit interfaces]
user@R1# set ge-1/0/0 description to-host
```

Copyright © 2016, Juniper Networks, Inc.

```
user@R1# set ge-1/0/0 unit 0 family inet address 172.16.50.2/30

user@R1# set ge-1/0/2 description to-R2
user@R1# set ge-1/0/2 unit 0 family inet address 10.30.0.1/30
```

2. Configure the custom forwarding classes and associated queue numbers.

```
[edit class-of-service forwarding-classes]
user@R1# set BE-data queue-num 0
user@R1# set Premium-data queue-num 1
user@R1# set Voice queue-num 2
user@R1# set NC queue-num 3
```

3. Configure the firewall filter term that places TCP traffic with a source port of 80 (HTTP traffic) into the BE-data forwarding class, associated with queue 0.

```
[edit firewall family inet filter mf-classifier]
user@R1# set term BE-data from protocol tcp
user@R1# set term BE-data from port 80
user@R1# set term BE-data then forwarding-class BE-data
```

4. Configure the firewall filter term that places TCP traffic with a source port of 12345 into the Premium-data forwarding class, associated with queue 1.

```
[edit firewall family inet filter mf-classifier]
user@R1# set term Premium-data from protocol tcp
user@R1# set term Premium-data from port 12345
user@R1# set term Premium-data then forwarding-class Premium-data
```

5. At the end of your firewall filter, configure a default term that accepts all other traffic.

   Otherwise, all traffic that arrives on the interface and is not explicitly accepted by the firewall filter is discarded.

```
[edit firewall family inet filter mf-classifier]
user@R1# set term accept-all-else then accept
```

6. Apply the firewall filter to the ge-1/0/0 interface as an input filter.

```
[edit interfaces]
user@R1# set ge-1/0/0 unit 0 family inet filter input mf-classifier
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show class-of-service**, **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-1/0/0 {
    description to-host;
    unit 0 {
        family inet {
            filter {
                input mf-classifier;
            }
            address 172.16.50.2/30;
        }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
ge-1/0/2 {
  description to-R2;
  unit 0 {
    family inet {
      address 10.30.0.1/30;
    }
  }
}
```

```
user@R1# show class-of-service
forwarding-classes {
  class BE-data queue-num 0;
  class Premium-data queue-num 1;
  class Voice queue-num 2;
  class NC queue-num 3;
}
```

```
user@R1# show firewall
family inet {
  filter mf-classifier {
    term BE-data {
      from {
        protocol tcp;
        port 80;
      }
      then forwarding-class BE-data;
    }
    term Premium-data {
      from {
        protocol tcp;
        port 12345;
      }
      then forwarding-class Premium-data;
    }
    term accept-all-else {
      then accept;
    }
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- Checking the CoS Settings on page 1449
- Sending TCP Traffic into the Network and Monitoring the Queue Placement on page 1450

Checking the CoS Settings

Purpose Confirm that the forwarding classes are configured correctly.

Copyright © 2016, Juniper Networks, Inc.

Action From Device R1, run the **show class-of-service forwardng-classes** command.

```
user@R1> show class-of-service forwarding-class

Forwarding class                          ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  BE-data                                 0       0          0              low
            normal            low
  Premium-data                            1       1          1              low
            normal            low
  Voice                                   2       2          2              low
            normal            low
  NC                                      3       3          3              low
            normal            low
```

Meaning The output shows the configured custom classifier settings.

Sending TCP Traffic into the Network and Monitoring the Queue Placement

Purpose Make sure that the traffic of interest is sent out the expected queue.

Action

1. Clear the interface statistics on Device R1's outgoing interface.

   ```
   user@R1> clear interfaces statistics ge-1/0/2
   ```

2. Use a traffic generator to send 50 TCP port 80 packets to Device R2 or to some other downstream device.

3. On Device R1, check the queue counters.

   Notice that you check the queue counters on the downstream output interface, not on the incoming interface.

   ```
   user@R1> show interfaces extensive ge-1/0/2 | find "Queue counters"

     Queue counters:      Queued packets  Transmitted packets     Dropped packets

       0                                 50                    50
   0
       1                                  0                    57
   0
       2                                  0                     0
   0
       3                                  0                     0
   0
   ```

4. Use a traffic generator to send 50 TCP port 12345 packets to Device R2 or to some other downstream device.

   ```
   [root@host]#  hping 172.16.60.1 -c 50 -s 12345 -k
   ```

5. On Device R1, check the queue counters.

   ```
   user@R1> show interfaces extensive ge-1/0/2 | find "Queue counters"

     Queue counters:      Queued packets  Transmitted packets     Dropped packets

       0                                 50                    50
   0
       1                                 50                    57
   0
   ```


Copyright © 2016, Juniper Networks, Inc.

```
    2                  0                  0
0
    3                  0                  0
0
```

Meaning The output shows that the packets are classified correctly. When port 80 is used in the TCP packets, queue 0 is incremented. When port 12345 is used, queue 1 is incremented.

Related Documentation

- Example: Configuring a Two-Rate Three-Color Policer on page 1429

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 68

Controlling Output Queues with Forwarding Classes

- Forwarding Classes Overview on page 1453
- Example: Configuring Forwarding Classes on page 1455
- Example: Assigning Forwarding Classes to Output Queues on page 1460
- Example: Assigning a Forwarding Class to an Interface on page 1462
- Understanding the SPC High-Priority Queue on page 1463
- Example: Configuring the SPC High-Priority Queue on page 1463

Forwarding Classes Overview

Forwarding classes (FCs) allow you to group packets for transmission and to assign packets to output queues. The forwarding class and the loss priority define the per-hop behavior (PHB in DiffServ) of a packet.

Juniper Networks devices support eight queues (0 through 7). For a classifier to assign an output queue (default queues 0 through 3) to each packet, it must associate the packet with one of the following forwarding classes:

- Expedited forwarding (EF)—Provides a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service.
- Assured forwarding (AF)—Provides a group of values you can define and includes four subclasses—AF1, AF2, AF3, and AF4—each with three drop probabilities (low, medium, and high).
- Best effort (BE)—Provides no service profile. For the BE forwarding class, loss priority is typically not carried in a class-of-service (CoS) value, and random early detection (RED) drop profiles are more aggressive.
- Network Control (NC)—This class is typically high priority because it supports protocol control.

In addition to behavior aggregate (BA) and mulitfield (MF) classification, the forwarding class (FC) of a packet can be directly determined by the logical interface that receives the packet. The packet FC can be configured using CLI commands, and if configured, this

Copyright © 2016, Juniper Networks, Inc.

FC overrides the FC from any BA classification that was previously configured on the logical interface.

The following CLI command can assign an FC directly to packets received at a logical interface:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
forwarding-class class-name;
```

This section contains the following topics:

- Forwarding Class Queue Assignments on page 1454
- Forwarding Policy Options on page 1455

Forwarding Class Queue Assignments

Juniper Networks devices have eight queues built into the hardware. By default, four queues are assigned to four FCs. Table 125 shows the four default FCs and queues that Juniper Networks classifiers assign to packets, based on the class-of-service (CoS) values in the arriving packet headers.

NOTE: **Queues 4 through 7 have no default assignments to FCs and are not mapped. To use queues 4 through 7, you must create custom FC names and map them to the queues.**

By default, all incoming packets, except the IP control packets, are assigned to the FC associated with queue 0. All IP control packets are assigned to the FC associated with queue 3.

Table 125: Default Forwarding Class Queue Assignments

Forwarding Queue	Forwarding Class	Forwarding Class Description
Queue 0	best-effort (BE)	The Juniper Networks device does not apply any special CoS handling to packets with 000000 in the DiffServ field, a backward compatibility feature. These packets are usually dropped under congested network conditions.
Queue 1	expedited-forwarding (EF)	The Juniper Networks device delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class. Devices accept excess traffic in this class, but in contrast to assured forwarding, out-of-profile expedited-forwarding packets can be forwarded out of sequence or dropped.

Copyright © 2016, Juniper Networks, Inc.

Table 125: Default Forwarding Class Queue Assignments *(continued)*

Forwarding Queue	Forwarding Class	Forwarding Class Description
Queue 2	assured-forwarding (AF)	The Juniper Networks device offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The device accepts excess traffic, but applies a random early detection (RED) drop profile to determine whether the excess packets are dropped and not forwarded. Three drop probabilities (low, medium, and high) are defined for this service class.
Queue 3	network-control (NC)	The Juniper Networks device delivers packets in this service class with a low priority. (These packets are not delay sensitive.) Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard.

Forwarding Policy Options

CoS-based forwarding (CBF) enables you to control next-hop selection based on a packet's CoS and, in particular, the value of the IP packet's precedence bits. For example, you can specify a particular interface or next hop to carry high-priority traffic while all best-effort traffic takes some other path. CBF allows path selection based on FC. When a routing protocol discovers equal-cost paths, it can either pick a path at random or load-balance the packets across the paths, through either hash selection or round-robin selection.

A forwarding policy also allows you to create CoS classification overrides. You can override the incoming CoS classification and assign the packets to an FC based on the packets' input interfaces, input precedence bits, or destination addresses. When you override the classification of incoming packets, any mappings you configured for associated precedence bits or incoming interfaces to output transmission queues are ignored.

Related Documentation

- Example: Assigning Forwarding Classes to Output Queues on page 1460
- Example: Assigning a Forwarding Class to an Interface on page 1462

Example: Configuring Forwarding Classes

By default on all platforms, four output queues are mapped to four FCs as shown in "Forwarding Classes Overview" on page 1453. On Juniper Networks devices, you can configure up to eight FCs and eight queues.

To configure up to eight FCs, include the **queue** statement at the **[edit class-of-service forwarding-classes]** hierarchy level:

Copyright © 2016, Juniper Networks, Inc.

```
[edit class-of-service forwarding-classes]
queue queue-number class-name;
```

The output queue number can be from 0 through 7, and you must map the forwarding classes one-to-one with the output queues. The default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

For example, to configure a one-to-one mapping between eight FCs and eight queues, you would use the following configuration:

```
[edit class-of-service]
forwarding-classes {
  queue 0 be;
  queue 1 ef;
  queue 2 af;
  queue 3 nc;
  queue 4 ef1;
  queue 5 ef2;
  queue 6 af1;
  queue 7 nc1;
}
```

Defining Eight Classifiers

```
[edit class-of-service]
classifiers {
  dscp dscp-table {
    forwarding-class ef {
      loss-priority low code-points [101000, 101001];
      loss-priority high code-points [101010, 101011];
    }
    forwarding-class af {
      loss-priority low code-points [010000, 010001];
      loss-priority high code-points [010010, 010011];
    }
    forwarding-class be {
      loss-priority low code-points [000000];
    }
    forwarding-class nc {
      loss-priority low code-points [111000];
    }
    forwarding-class ef1 {
      loss-priority low code-points [101100, 101101];
      loss-priority high code-points [101110];
    }
    forwarding-class af1 {
      loss-priority high code-points [101110];
    }
    forwarding-class ef2 {
      loss-priority low code-points [101111];
    }
    forwarding-class nc1 {
      loss-priority low code-points [111001];
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Adding Eight Schedulers to a Scheduler Map

Configure a custom scheduler map that applies globally to all interfaces, except those that are restricted to four queues:

```
[edit class-of-service]
scheduler-maps {
  sched {
    forwarding-class be scheduler Q0;
    forwarding-class ef scheduler Q1;
    forwarding-class af scheduler Q2;
    forwarding-class nc scheduler Q3;
    forwarding-class ef1 scheduler Q4;
    forwarding-class ef2 scheduler Q5;
    forwarding-class af1 scheduler Q6;
    forwarding-class nc1 scheduler Q7;
  }
}
schedulers {
  Q0 {
    transmit-rate percent 25;
    buffer-size percent 25;
    priority low;
    drop-profile-map loss-priority any protocol both drop-default;
  }
  Q1 {
    buffer-size temporal 2000;
    priority strict-high;
    drop-profile-map loss-priority any protocol both drop-ef;
  }
  Q2 {
    transmit-rate percent 35;
    buffer-size percent 35;
    priority low;
    drop-profile-map loss-priority any protocol both drop-default;
  }
  Q3 {
    transmit-rate percent 5;
    buffer-size percent 5;
    drop-profile-map loss-priority any protocol both drop-default;
  }
  Q4 {
    transmit-rate percent 5;
    priority high;
    drop-profile-map loss-priority any protocol both drop-ef;
  }
  Q5 {
    transmit-rate percent 10;
    priority high;
    drop-profile-map loss-priority any protocol both drop-ef;
  }
  Q6 {
    transmit-rate remainder;
    priority low;
    drop-profile-map loss-priority any protocol both drop-default;
  }
  Q7 {
    transmit-rate percent 5;
```

Copyright © 2016, Juniper Networks, Inc.

```
      priority high;
      drop-profile-map loss-priority any protocol both drop-default;
    }
  }
```

Configuring an IP Precedence Classifier and Rewrite Tables

```
[edit class-of-service]
classifiers {
  inet-precedence inet-classifier {
    forwarding-class be {
      loss-priority low code-points 000;
    }
    forwarding-class af11 {
      loss-priority high code-points 001;
    }
    forwarding-class ef {
      loss-priority low code-points 010;
    }
    forwarding-class nc1 {
      loss-priority high code-points 011;
    }
    forwarding-class {
      loss-priority low code-points 100;
    }
    forwarding-class af12 {
      loss-priority high code-points 101;
    }
    forwarding-class ef1 {
      loss-priority low code-points 110;
    }
    forwarding-class nc2 {
      loss-priority high code-points 111;
    }
  }
}
exp exp-rw-table {
  forwarding-class be {
    loss-priority low code-point 000;
  }
  forwarding-class af11 {
    loss-priority high code-point 001;
  }
  forwarding-class ef {
    loss-priority low code-point 010;
  }
  forwarding-class nc1 {
    loss-priority high code-point 111;
  }
  forwarding-class be1 {
    loss-priority low code-point 100;
  }
  forwarding-class af12 {
    loss-priority high code-point 101;
  }
  forwarding-class ef1 {
    loss-priority low code-point 110;
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    forwarding-class nc2 {
      loss-priority low code-point 111;
    }
  }
  inet-precedence inet-rw-table {
    forwarding-class be {
      loss-priority low code-point 000;
    }
    forwarding-class af11 {
      loss-priority high code-point 001;
    }
    forwarding-class ef1 {
      loss-priority low code-point 010;
    }
    forwarding-class nc1 {
      loss-priority low code-point 111;
    }
    forwarding-class be1 {
      loss-priority low code-point 100;
    }
    forwarding-class af12 {
      loss-priority high code-point 101;
    }
    forwarding-class ef1 {
      loss-priority low code-point 111;
    }
    forwarding-class nc2 {
      loss-priority low code-point 110;
    }
  }
```

Configuring an IDP Policy with a Forwarding Class

Configure an IDP policy with a forwarding class as an action to rewrite DSCP values of IP packets:

```
[edit class-of-service]
security idp idp-policy policy_name rulebase-ips rule rule_name {
  then {
    action {
      class-of-service {
        forwarding-class forwarding-class-name;
        dscp-code-point value;
        }
      }
    }
  }
```

Related Documentation

- Forwarding Classes Overview on page 1453
- Example: Assigning Forwarding Classes to Output Queues on page 1460
- Example: Assigning a Forwarding Class to an Interface on page 1462

Copyright © 2016, Juniper Networks, Inc.

Example: Assigning Forwarding Classes to Output Queues

This example shows how to assign forwarding classes to output queues.

- Requirements on page 1460
- Overview on page 1460
- Configuration on page 1460
- Verification on page 1461

Requirements

Before you begin, determine the MF classifier. See "Example: Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 1445.

Overview

In this example, you configure class of service and assign best-effort traffic to queue 0 as be-class, expedited forwarding traffic to queue 1 as ef-class, assured forwarding traffic to queue 2 as af-class, and network control traffic to queue 3 as nc-class.

You must assign the forwarding classes established by the MF classifier to output queues. Table 126 shows how this example assigns output queues.

Table 126: Sample Output Queue Assignments for mf-classifier Forwarding Queues

mf-classifier Forwarding Class	For Traffic Type	Output Queue
be-class	Best-effort traffic	Queue 0
ef-class	Expedited forwarding traffic	Queue 1
af-class	Assured forwarding traffic	Queue 2
nc-class	Network control traffic	Queue 3

Configuration

CLI Quick Configuration

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

```
set class-of-service forwarding-classes queue 0 be-class
set class-of-service forwarding-classes queue 1 ef-class
set class-of-service forwarding-classes queue 2 af-class
set class-of-service forwarding-classes queue 3 nc-class
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To assign forwarding classes to output queues:

1. Configure class of service.

   ```
   [edit]
   user@host# edit class-of-service forwarding-classes
   ```

2. Assign best-effort traffic to queue 0.

   ```
   [edit class-of-service forwarding-classes]
   user@host# set queue 0 be-class
   ```

3. Assign expedited forwarding traffic to queue 1.

   ```
   [edit class-of-service forwarding-classes]
   user@host# set queue 1 ef-class
   ```

4. Assign assured forwarding traffic to queue 2.

   ```
   [edit class-of-service forwarding-classes]
   user@host# set queue 2 af-class
   ```

5. Assign network control traffic to queue 3.

   ```
   [edit class-of-service forwarding-classes]
   user@host# set queue 3 nc-class
   ```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  forwarding-classes {
  queue 0 be-class;
    queue 1 ef-class;
    queue 2 af-class;
  queue 3 nc-class;
}
```

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

NOTE: You cannot commit a configuration that assigns the same forwarding class to two different queues.

Verification

Verifying Forwarding Classes Are Assigned to Output Queues

Purpose

Verify that the forwarding classes are properly assigned to output queues.

Copyright © 2016, Juniper Networks, Inc.

Action From configuration mode, enter the **show class-of-service** command.

Related Documentation

- Forwarding Classes Overview on page 1453
- Example: Assigning a Forwarding Class to an Interface on page 1462
- Example: Configuring Forwarding Classes on page 1455

Example: Assigning a Forwarding Class to an Interface

This example shows how to assign a forwarding class to an interface.

- Requirements on page 1462
- Overview on page 1462
- Configuration on page 1462
- Verification on page 1463

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

On a device, you can configure fixed classification on a logical interface by specifying a forwarding class to be applied to all packets received by the logical interface, regardless of the packet contents.

In this example, you configure class of service, create interface ge-3/0/0 unit 0 and then set the forwarding class to assured-forwarding.

All packets coming into the device from the ge-3/0/0 unit 0 interface are assigned to the assured-forwarding forwarding class.

Configuration

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

To assign a forwarding class to an interface:

1. Configure class of service and assign the interface.

   ```
   [edit]
   user@host# edit class-of-service interfaces ge-3/0/0 unit 0
   ```

2. Specify the forwarding class.

   ```
   [edit class-of-service interfaces ge-3/0/0 unit 0]
   user@host# set forwarding-class assured-forwarding
   ```

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

Copyright © 2016, Juniper Networks, Inc.

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

Verification

To verify the configuration is working properly, enter the **show class-of-service** command.

Related Documentation

- Forwarding Classes Overview on page 1453
- Example: Assigning Forwarding Classes to Output Queues on page 1460

Understanding the SPC High-Priority Queue

The Services Processing Card (SPC) on SRX5000 line devices provides processing power to run integrated services such as firewall, IPsec, and IDP. All traffic traversing the SRX Series device is passed to an SPC to have service processing applied. Junos OS provides a configuration option to enable packets with specific Differentiated Services (DiffServ) code points (DSCP) precedence bits to enter a high-priority queue on the SPC. The Services Processing Unit (SPU) draws packets from the high-priority queue and only draws packets from a low-priority queue when the high-priority queue is empty. This feature can reduce overall latency for real-time traffic, such as voice traffic.

To designate packets for the high-priority or low-priority queues, use the **spu-priority** configuration statement at the [**edit class-of-service forwarding-classes *class***] hierarchy level. A value of **high** places packets into the high-priority queue, and a value of **low** places packets into the low-priority queue.

Related Documentation

- Example: Configuring the SPC High-Priority Queue on page 1463
- Forwarding Classes Overview on page 1453

Example: Configuring the SPC High-Priority Queue

This example shows how to configure a forwarding class for the high-priority queue on the SPC.

- Requirements on page 1463
- Overview on page 1464
- Configuration on page 1464
- Verification on page 1465

Requirements

This example uses the following hardware and software components:

- SRX5000 line device
- Junos OS Release 11.4R2 or later

Copyright © 2016, Juniper Networks, Inc.

Overview

This example defines the following forwarding classes and assigns a queue number to each class:

Forwarding Class	Queue Number
best-effort	0
assured-forwarding	1
network-control	3
expedited-forwarding	2

The expedited-forwarding class is configured for the high-priority queue on the SPC.

Configuration

CLI Quick Configuration

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

```
set class-of-service forwarding-classes class best-effort queue-num 0
set class-of-service forwarding-classes class assured-forwarding queue-num 1
set class-of-service forwarding-classes class network-control queue-num 3
set class-of-service forwarding-classes class expedited-forwarding queue-num 2
  spu-priority high
```

Step-by-Step Procedure

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

To configure the high-priority queue on the SPC:

1. Define forwarding classes and assign queue numbers.

```
[edit class-of-service forwarding-classes]
user@host# set class best-effort queue-num 0
user@host# set class assured-forwarding queue-num 1
user@host# set class network-control queue-num 3
user@host# set class expedited-forwarding queue-num 2
```

2. Configure the SPC high-priority queue.

```
[edit class-of-service forwarding-classes]
user@host# set class expedited-forwarding spu-priority high
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service forwarding-classes** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show class-of-service forwarding-classes
class best-effort queue-num 0;
class assured-forwarding queue-num 1;
class network-control queue-num 3;
class expedited-forwarding queue-num 2 spu-priority high;
```

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

Verification

Verifying SPU High-Priority Queue Mapping

Purpose Verify that the forwarding class is mapped to the SPU high-priority queue.

Action From operational mode, enter the **show class-of-service forwarding-class** command.

```
user@host> show class-of-service forwarding-class

Forwarding class                          ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  best-effort                              0       0          0             low
            normal             low
  expedited-forwarding                     1       1          1             low
            normal             high
  assured-forwarding                       2       2          2             low
            normal             low
  network-control                          3       3          3             low
            normal             low
```

Related Documentation

- Understanding the SPC High-Priority Queue on page 1463

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 69

Altering Outgoing Packets Headers with Rewrite Rules

- Rewrite Rules Overview on page 1467
- Rewriting Frame Relay Headers on page 1467
- Example: Configuring and Applying Rewrite Rules on page 1469

Rewrite Rules Overview

A rewrite rule modifies the appropriate class-of-service (CoS) bits in an outgoing packet. Modification of CoS bits allows the next downstream device to classify the packet into the appropriate service group. Rewriting or marking outbound packets is useful when the device is at the border of a network and must alter the CoS values to meet the policies of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.

Typically, a device rewrites CoS values in outgoing packets on the outbound interfaces of an edge device, to meet the policies of the targeted peer. After reading the current forwarding class and loss priority information associated with the packet, the transmitting device locates the chosen CoS value from a table, and writes this CoS value into the packet header.

NOTE: You can configure up to 32 IEEE 802.1p rewriters on each SRX5K-MPC on the SRX5600 and SRX5800 devices.

Related Documentation

- Example: Configuring and Applying Rewrite Rules on page 1469

Rewriting Frame Relay Headers

This section contains the following topics:

- Assigning the Default Frame Relay Rewrite Rule to an Interface on page 1468
- Defining a Custom Frame Relay Rewrite Rule on page 1468

Copyright © 2016, Juniper Networks, Inc.

Assigning the Default Frame Relay Rewrite Rule to an Interface

For Juniper Networks device interfaces with Frame Relay encapsulation, you can rewrite the discard eligibility (DE) bit based on the loss priority of Frame Relay traffic. For each outgoing frame with the loss priority set to low, medium-low, medium-high, or high, you can set the DE bit CoS value to **0** or **1**. You can combine a Frame Relay rewrite rule with other rewrite rules on the same interface. For example, you can rewrite both the DE bit and MPLS EXP bit.

The default Frame Relay rewrite rule contains the following settings:

```
loss-priority low code-point 0;
loss-priority medium-low code-point 0;
loss-priority medium-high code-point 1;
loss-priority high code-point 1;
```

This default rule sets the DE CoS value to **0** for each outgoing frame with the loss priority set to low or medium-low. This default rule sets the DE CoS value to **1** for each outgoing frame with the loss priority set to medium-high or high.

To assign the default rule to an interface, include the **frame-relay-de default** statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
frame-relay-de default;
```

Defining a Custom Frame Relay Rewrite Rule

To define a custom Frame Relay rewrite rule, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
rewrite-rules {
  frame-relay-de rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority level code-point (0 | 1);
    }
  }
}
```

A custom rewrite rule sets the DE bit to the **0** or **1** CoS value based on the assigned loss priority of low, medium-low, medium-high, or high for each outgoing frame.

The rule does not take effect until you apply it to a logical interface. To apply a rule to a logical interface, include the **frame-relay-de *map-name*** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
frame-relay-de map-name;
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Rewrite Rules Overview on page 1467
- Example: Configuring and Applying Rewrite Rules on page 1469

Example: Configuring and Applying Rewrite Rules

This example shows how to configure and apply rewrite rules for a device.

- Requirements on page 1469
- Overview on page 1469
- Configuration on page 1470
- Verification on page 1472

Requirements

Before you begin, create and configure the forwarding classes.

Overview

You can configure rewrite rules to replace DSCPs on packets received from the customer or host with the values expected by other devices. You do not have to configure rewrite rules if the received packets already contain valid DSCPs. Rewrite rules apply the forwarding class information and packet loss priority used internally by the device to establish the DSCP on outbound packets. After you configure the rewrite rules, you must apply them to the correct interfaces.

NOTE: You can configure up to 32 IEEE 802.1p rewriters on each SRX5K-MPC on the SRX5600 and SRX5800 devices.

In this example, you configure the rewrite rules for DiffServ CoS as rewrite-dscps. You specify the best-effort forwarding class as be-class, expedited forwarding class as ef-class, an assured forwarding class as af-class, and a network control class as nc-class. Finally, you apply rewrite rules to an interface called ge-0/0/0.

Table 127 shows how the rewrite rules replace the DSCPs on packets in the four forwarding classes.

Table 127: Sample rewrite-dscps Rewrite Rules to Replace DSCPs

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
be-class	Best-effort traffic—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined.	Low-priority code point: 000000 High-priority code point: 000001
ef-class	Expedited forwarding traffic—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped.	Low-priority code point: 101110 High-priority code point: 101111

Copyright © 2016, Juniper Networks, Inc.

Table 127: Sample rewrite-dscps Rewrite Rules to Replace DSCPs *(continued)*

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
af-class	Assured forwarding traffic—Provides high assurance for packets within the specified service profile. Excess packets are dropped.	Low-priority code point: 001010 High-priority code point: 001100
nc-class	Network control traffic—Packets can be delayed but not dropped.	Low-priority code point: 110000 High-priority code point: 110001

NOTE: Forwarding classes can be configured in a DSCP rewriter and also as an action of an IDP policy to rewrite DSCP code points. To ensure that the forwarding class is used as an action in an IDP policy, it is important that you do not configure an IDP policy and interface-based rewrite rules with the same forwarding class.

Configuration

CLI Quick Configuration

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

```
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class be-class
    loss-priority low code-point 000000
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class be-class
    loss-priority high code-point 000001
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class ef-class loss-priority
    low code-point 101110
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class ef-class loss-priority
    high code-point 101111
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class af-class loss-priority
    low code-point 001010
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class af-class loss-priority
    high code-point 001100
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class nc-class loss-priority
    low code-point 110000
set class-of-service rewrite-rules dscp rewrite-dscps forwarding-class nc-class loss-priority
    high code-point 110001
set class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules dscp rewrite-dscps
```

Step-by-Step Procedure

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

To configure and apply rewrite rules for a device:

1. Configure rewrite rules for DiffServ CoS.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# edit class-of-service
user@host# edit rewrite-rules dscp rewrite-dscps
```

2. Configure best-effort forwarding class rewrite rules.

```
[edit class-of-service rewrite-rules dscp rewrite-dscps]
user@host# set forwarding-class be-class loss-priority low code-point 000000
user@host# set forwarding-class be-class loss-priority high code-point 000001
```

3. Configure expedited forwarding class rewrite rules.

```
[edit class-of-service rewrite-rules dscp rewrite-dscps]
user@host# set forwarding-class ef-class loss-priority low code-point 101110
user@host# set forwarding-class ef-class loss-priority high code-point 101111
```

4. Configure an assured forwarding class rewrite rules.

```
[edit class-of-service rewrite-rules dscp rewrite-dscps]
user@host# set forwarding-class af-class loss-priority low code-point 001010
user@host# set forwarding-class af-class loss-priority high code-point 001100
```

5. Configure a network control class rewrite rules.

```
[edit class-of-service rewrite-rules dscp rewrite-dscps]
user@host# set forwarding-class nc-class loss-priority low code-point 110000
user@host# set forwarding-class nc-class loss-priority high code-point 110001
```

6. Apply rewrite rules to an interface.

```
[edit class-of-service]
user@host# set interfaces ge-0/0/0 unit 0 rewrite-rules dscp rewrite-dscps
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  interfaces {
  interfaces {
    unit 0 {
    rewrite-rules {
    dscp rewrite-dscps;
      }
    }
  }
  }
  rewrite-rules {
    dscp rewrite-dscps {
    forwarding-class be-class {
    loss-priority low code-point 000000;
      loss-priority high code-point 000001;
      }
      forwarding-class ef-class {
        loss-priority low code-point 101110;
      loss-priority high code-point 101111;
    }
    forwarding-class af-class {
```

Copyright © 2016, Juniper Networks, Inc.

```
            loss-priority low code-point 001010;
        loss-priority high code-point 001100;
    }
    forwarding-class nc-class {
        loss-priority low code-point 110000;
        loss-priority high code-point 110001;
        }
    }
}
```

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

Verification

Verifying Rewrite Rules Configuration

Purpose Verify that rewrite rules are configured properly.

Action From configuration mode, enter the **show class-of-service** command.

Related Documentation

- Rewrite Rules Overview on page 1467

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 70

Defining Output Queue Properties with Schedulers

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Transmission Scheduling Overview on page 1479
- Excess Bandwidth Sharing and Minimum Logical Interface Shaping on page 1480
- Excess Bandwidth Sharing Proportional Rates on page 1481
- Calculated Weights Mapped to Hardware Weights on page 1482
- Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces on page 1483
- Shared Bandwidth Among Logical Interfaces on page 1484
- Example: Configuring Class-of-Service Schedulers on page 1485
- Scheduler Buffer Size Overview on page 1489
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Configuring Large Delay Buffers in CoS on page 1496
- Example: Configuring and Applying Scheduler Maps on page 1500

Schedulers Overview

You use *schedulers* to define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the random early detection (RED) drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of *scheduler maps*. You can then associate each scheduler map with an interface, thereby configuring the hardware queues, packet schedulers, and RED processes that operate according to this mapping.

An individual device interface has multiple queues assigned to store packets temporarily before transmission. To determine the order to service the queues, the device uses a round-robin scheduling method based on priority and the queue's weighted round-robin (WRR) credits. Junos OS schedulers allow you to define the priority, bandwidth, delay

Copyright © 2016, Juniper Networks, Inc.

buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission.

You can configure *per-unit scheduling* (also called *logical interface scheduling*) to allow multiple output queues on a logical interface and to associate an output scheduler with each queue.

NOTE: For Juniper Network devices, when configuring the "protocol parameter" in the **drop-profile-map** statement, TCP and non-TCP values are not supported; only the value "any" is supported.

This section contains the following topics:

- Transmit Rate on page 1474
- Delay Buffer Size on page 1475
- Scheduling Priority on page 1476
- Shaping Rate on page 1477

Transmit Rate

The transmission rate determines the traffic transmission bandwidth for each forwarding class you configure. The rate is specified in bits per second (bps). Each queue is allocated some portion of the bandwidth of the outgoing interface.

This bandwidth amount can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. You can limit the transmission bandwidth to the exact value you configure, or allow it to exceed the configured rate if additional bandwidth is available from other queues (SRX Series high-end devices do not support an exact value transmit rate). This property helps ensure that each queue receives the amount of bandwidth appropriate to its level of service.

The minimum transmit rate supported on high-speed interfaces is one-ten thousandth of the speed of that interface. For example, on a Gigabit Ethernet interface with a speed of 1,000 Mbps, the minimum transmit rate is 100 Kbps (1,000 Mbps x 1/10,000). You can configure transmit rates in the range 3,200 bps through 160,000,000,000 bps. When the configured rate is less than the minimum transmit rate, the minimum transmit rate is used instead.

NOTE: Interfaces with slower interface speeds, like T1, E1, or channelized T1/E1/ISDN PRI, cannot support minimum transmit rates because the minimum transmit rate supported on a device is 3,200 bps.

Transmit rate assigns the weighted round-robin (WRR) priority values within a given priority level and not between priorities.

Copyright © 2016, Juniper Networks, Inc.

The transmit rate defines the transmission rate of a scheduler. The transmit rate determines the traffic bandwidth from each forwarding class you configure.

By default, queues 0 through 7 have the following percentage of transmission capacity:

- Queue 0—95 percent
- Queue 1—0 percent
- Queue 2—0 percent
- Queue 3—0 percent
- Queue 4—0 percent
- Queue 6—0 percent
- Queue 7—5 percent

To define a transmit rate, select the appropriate option:

- To specify a transmit rate, select **rate** and type an integer from 3200 to 160,000,000,000 bits per second.
- To enforce an exact transmit rate, select **rate**.
- To specify the remaining transmission capacity, select **remainder**.
- To specify a percentage of transmission capacity, select **percent** and type an integer from 1 through 100.

Optionally, you can specify the percentage of the remainder to be used for allocating the transmit rate of the scheduler on a prorated basis. If there are still points left even after allocating the remainder percentage with the transmit rate and there are no queues, then the points are allocated point by point to each queue in a round-robin method. If the remainder percentage is not specified, the remainder value will be shared equally.

Delay Buffer Size

You can configure the delay buffer size to control congestion at the output stage. A delay buffer provides packet buffer space to absorb burst traffic up to a specified duration of delay. When the buffer is full, all packets are dropped.

On Juniper Networks devices, you can configure larger delay buffers on channelized T1/E1 interfaces. Larger delay buffers help these slower interfaces to avoid congestion and packet dropping when they receive large bursts of traffic.

By default, all branch SRX Series device interfaces support a delay buffer time of 100,000 microseconds.

To define a delay buffer size for a scheduler, select the appropriate option:

- To enforce exact buffer size, select **Exact**.
- To specify a buffer size as a temporal value (microseconds), select **Temporal**.

Copyright © 2016, Juniper Networks, Inc.

- To specify buffer size as a percentage of the total buffer, select **Percent** and type an integer from 1 through 100.
- To specify buffer size as the remaining available buffer, select **Remainder**.

Optionally, you can specify the percentage of the remainder to be used for allocating the buffer size of the scheduler on a prorated basis.

By default, sizes of the delay buffer queues 0 through 7 have the following percentage of the total available buffer space:

- Queue 0—95 percent
- Queue 1—0 percent
- Queue 2—0 percent
- Queue 3—0 percent
- Queue 4—0 percent
- Queue 5—0 percent
- Queue 6—0 percent
- Queue 7—5 percent

NOTE: A large buffer size value correlates with a greater possibility of packet delays. This might not be practical for sensitive traffic such as voice or video.

NOTE: For a Juniper Networks device, if the buffer size percentage is set to zero for T1 interfaces, traffic does not pass.

Scheduling Priority

Scheduling priority determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic are provided better access to the outgoing interface.

The queues for an interface are divided into sets based on their priority. Each set contains queues of the same priority. The device examines the sets in descending order of priority. If at least one queue in a set has a packet to transmit, the device selects that set. If multiple queues in the set have packets to transmit, the device selects a queue from the set according to the weighted round-robin (WRR) algorithm that operates within the set.

The packets in a queue are transmitted based on the configured scheduling priority, the transmit rate, and the available bandwidth.

The scheduling priority of the scheduler determines the order in which an output interface transmits traffic from the queues. You can set scheduling priority at different levels in an

Copyright © 2016, Juniper Networks, Inc.

order of increasing priority from low to high. A high-priority queue with a high transmission rate might lock out lower-priority traffic.

To specify a scheduling priority, select one of the following levels:

- **high**—Packets in this queue have high priority.
- **low**—Packets in this queue are transmitted last.
- **medium—low**—Packets in this queue have medium-low priority.
- **medium—high**—Packets in this queue have medium-high priority.
- **strict—high**—Packets in this queue are transmitted first.

Shaping Rate

Shaping rates control the maximum rate of traffic transmitted on an interface. You can configure the shaping rate so that the interface transmits less traffic than it is physically capable of carrying.

You can configure shaping rates on logical interfaces. By default, output scheduling is not enabled on logical interfaces. Logical interface scheduling (also called per-unit scheduling) allows you to enable multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.

By default, the logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. You can specify a peak bandwidth rate in bits per second (bps), either as a complete decimal number or as a decimal number followed by the abbreviation *k* (1,000), *m* (1,000,000), or *g* (1,000,000,000). The range is from 1,000 through 32,000,000,000 bps.

For low-speed interfaces, the queue-limit values might become lower than the interface MTU so that traffic with large packets can no longer pass through some of the queues. If you want larger-sized packets to flow through, set the buffer-size configuration in the scheduler to a larger value. For more accuracy, the 100-ms queue-limit values are calculated based on shaper rates and not on interface rates.

The shaping rate defines the minimum bandwidth allocated to a queue. The default shaping rate is 100 percent, which is the same as no shaping at all. To define a shaping rate, select the appropriate option:

- To specify shaping rate as an absolute number of bits per second, select **rate** and type an integer from 3,200 to 160,000,000,000 bits per second.
- To specify shaping rate as a percentage, select **percent** and type an integer from 0 through 100.

Related Documentation

- Default Scheduler Settings on page 1478
- Example: Configuring Class-of-Service Schedulers on page 1485
- Scheduler Buffer Size Overview on page 1489
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring and Applying Scheduler Maps on page 1500
- Transmission Scheduling Overview on page 1479

Default Scheduler Settings

Each forwarding class has an associated scheduler priority. Only two forwarding classes, best-effort and network-control (queue 0 and queue 3), are used in the Junos OS default scheduler configuration.

By default, the best-effort forwarding class (queue 0) receives 95 percent, and the network-control (queue 3) receives 5 percent of the bandwidth and buffer space for the output link. The default drop profile causes the buffer to fill and then discard all packets until it again has space.

The expedited-forwarding and assured-forwarding classes have no schedulers, because by default no resources are assigned to queue 1 and queue 2. However, you can manually configure resources for the expedited-forwarding and the assured-forwarding classes.

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. If you do not want a queue to use any leftover bandwidth, you must configure it for strict allocation.

The device uses the following default scheduler settings. You can configure these settings.

```
[edit class-of-service]
schedulers {
    network-control {
        transmit-rate percent 5;
        buffer-size percent 5;
        priority low;
        drop-profile-map loss-priority any protocol any drop-profile terminal;
    }
    best-effort {
        transmit-rate percent 95;
        buffer-size percent 95;
        priority low;
        drop-profile-map loss-priority any protocol any drop-profile terminal;
    }
}
drop-profiles {
    terminal {
        fill-level 100 drop-probability 100;
    }
}
```

Related Documentation

- Schedulers Overview on page 1473
- Example: Configuring Class-of-Service Schedulers on page 1485

Copyright © 2016, Juniper Networks, Inc.

- Scheduler Buffer Size Overview on page 1489
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Example: Configuring and Applying Scheduler Maps on page 1500
- Transmission Scheduling Overview on page 1479

Transmission Scheduling Overview

The packets in a queue are transmitted based on their transmission priority, transmit rate, and the available bandwidth.

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. A queue receiving traffic within its bandwidth configuration is considered to have positive bandwidth credit, and a queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit.

NOTE: **The queues in a logical interface do not use the available buffer from other queues for packet transmission. Instead, the packets transmitted to a queue consider only the buffer size available in its own queue.**

A queue with positive credit does not need to use leftover bandwidth, because it can use its own allocation. For such queues, packets are transmitted based on the priority of the queue, with packets from higher-priority queues transmitting first. The transmit rate is not considered during transmission. In contrast, a queue with negative credit needs a share of the available leftover bandwidth.

The leftover bandwidth is allocated to queues with negative credit in proportion to the configured transmit rate of the queues within a given priority set. The queues for an interface are divided into sets based on their priority. If no transmit rate is configured, each queue in the set receives an equal percentage of the leftover bandwidth. However, if a transmit rate is configured, each queue in the set receives the configured percentage of the leftover bandwidth.

Table 128 shows a sample configuration of priority and transmit rate on six queues. The total available bandwidth on the interface is 100 Mbps.

Table 128: Sample Transmission Scheduling

Queue	Scheduling Priority	Transmit Rate	Incoming Traffic
0	Low	10%	20 Mbps
1	High	20%	20 Mbps
2	High	30%	20 Mbps

Copyright © 2016, Juniper Networks, Inc.

Table 128: Sample Transmission Scheduling *(continued)*

Queue	Scheduling Priority	Transmit Rate	Incoming Traffic
3	Low	30%	20 Mbps
4	Medium-high	No transmit rate configured	10 Mbps
5	Medium-high	No transmit rate configured	20 Mbps

In this example, queues are divided into three sets based on their priority:

- High priority set—Consists of queue 1 and queue 2. Packets use 40 Mbps (20+20) of the available bandwidth (100 Mbps) and are transmitted first. Because of positive credit, the configured transmit rate is not considered.
- Medium-high priority set—Consists of queue 4 and queue 5. Packets use 30 Mbps (10+20) of the remaining 60 Mbps bandwidth. Because of positive credit, the transmit rate is not considered. If the queues had negative credit, they would receive an equal share of the leftover bandwidth because no transmit rate is configured.
- Low priority set—Consists of queue 0 and queue 3. Packets share the 20 Mbps of leftover bandwidth based on the configured transmit rate. The distribution of bandwidth is in proportion to the assigned percentages. Because the total assigned percentage is 40 (10 + 30), each queue receives a share of bandwidth accordingly. Thus queue 0 receives 5 Mbps (10/40 x 20), and queue 3 receives 15 Mbps (30/40 x 20).

Related Documentation

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Example: Configuring Class-of-Service Schedulers on page 1485
- Scheduler Buffer Size Overview on page 1489
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Example: Configuring and Applying Scheduler Maps on page 1500

Excess Bandwidth Sharing and Minimum Logical Interface Shaping

The default excess bandwidth sharing proportional rate is 32.65 Mbps (128 Kbps x 255). In order to have better weighed fair queuing (WFQ) accuracy among queues, the shaping rate configured should be larger than the excess bandwidth sharing proportional rate. Some examples are shown in Table 129.

Table 129: Shaping Rates and WFQ Weights

Shaping Rate	Configured Queue Transmit Rate	WFQ Weight	Total Weights
10 Mbps	(30, 40, 25, 5)	(22, 30, 20, 4)	76
33 Mbps	(30, 40, 25, 5)	(76, 104, 64, 13)	257

Copyright © 2016, Juniper Networks, Inc.

Table 129: Shaping Rates and WFQ Weights *(continued)*

Shaping Rate	Configured Queue Transmit Rate	WFQ Weight	Total Weights
40 Mbps	(30, 40, 25, 5)	(76, 104.64, 13)	257

With a 10-Mbps shaping rate, the total weights are 76. This is divided among the four queues according to the configured transmit rate. Note that when the shaping rate is larger than the excess bandwidth sharing proportional rate of 32.65 Mbps, the total weight on the logical interface is 257 and the WFQ accuracy will be the same.

When using the IOC (40x1GE IOC or 4x10GE IOC) on a Juniper Networks device, there are circumstances when you should configure excess bandwidth sharing and minimum logical interface shaping.

Related Documentation

- Schedulers Overview on page 1473
- Excess Bandwidth Sharing Proportional Rates on page 1481
- Calculated Weights Mapped to Hardware Weights on page 1482
- Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces on page 1483
- Shared Bandwidth Among Logical Interfaces on page 1484

Excess Bandwidth Sharing Proportional Rates

To determine a good excess bandwidth-sharing proportional rate to configure, choose the largest CIR (guaranteed rate) among all the logical interfaces (units). If the logical units have PIRs (shaping rates) only, then choose the largest PIR rate. However, this is not ideal if a single logical interface has a large WRR rate. This method can skew the distribution of traffic across the queues of the other logical interfaces. To avoid this issue, set the excess bandwidth-sharing proportional rate to a lower value on the logical interfaces where the WRR rates are concentrated. This improves the bandwidth sharing accuracy among the queues on the same logical interface. However, the excess bandwidth sharing for the logical interface with the larger WRR rate is no longer proportional.

As an example, consider five logical interfaces on the same physical port, each with four queues, all with only PIRs configured and no CIRs. The WRR rate is the same as the PIR for the logical interface. The excess bandwidth is shared proportionally with a rate of 40 Mbps. The traffic control profiles for the logical interfaces are shown in Table 130.

Table 130: Example Shaping Rates and WFQ Weights

Shaping Rate	Configured Queue Transmit Rate	WFQ Weight	Total Weights
(Unit 0) 10 Mbps	(95, 0, 0, 5)	(60, 0, 0, 3)	63
(Unit 1) 20 Mbps	(25, 25, 25, 25)	(32, 32, 32, 32)	128
(Unit 2) 40 Mbps	(40, 30, 20, 10)	(102, 77, 51, 26)	255

Copyright © 2016, Juniper Networks, Inc.

Table 130: Example Shaping Rates and WFQ Weights *(continued)*

Shaping Rate	Configured Queue Transmit Rate	WFQ Weight	Total Weights
(Unit 3) 200 Mbps	(70, 10, 10, 10)	(179, 26, 26, 26)	255
(Unit 4) 2 Mbps	(25, 25, 25, 25)	(5, 5, 5, 5)	20

Even though the maximum transmit rate for the queue on logical interface unit 3 is 200 Mbps, the excess bandwidth-sharing proportional rate is kept at a much lower value. Within a logical interface, this method provides a more accurate distribution of weights across queues. However, the excess bandwidth is now shared equally between unit 2 and unit 3 (total weights = 255).

Related Documentation

- Schedulers Overview on page 1473
- Excess Bandwidth Sharing and Minimum Logical Interface Shaping on page 1480
- Calculated Weights Mapped to Hardware Weights on page 1482
- Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces on page 1483
- Shared Bandwidth Among Logical Interfaces on page 1484

Calculated Weights Mapped to Hardware Weights

The calculated weight in a traffic control profile is mapped to hardware weight, but the hardware only supports a limited WFQ profile. The weights are rounded to the nearest hardware weight according to the values in Table 131.

Table 131: Rounding Configured Weights to Hardware Weights

Traffic Control Profile Number	Number of Traffic Control Profiles	Weights	Maximum Error
1–16	16	1–16 (interval of 1)	50.00%
17–29	13	18–42 (interval of 2)	6.25%
30–35	6	45–60 (interval of 3)	1.35%
36–43	8	64–92 (interval of 4)	2.25%
44–49	6	98–128 (interval of 6)	3.06%
50–56	7	136–184 (interval of 8)	3.13%
57–62	6	194–244 (interval of 10)	2.71%
63–63	1	255–255 (interval of 11)	2.05%

Copyright © 2016, Juniper Networks, Inc.

As shown in Table 131, the calculated weight of 18.9 is mapped to a hardware weight of 18, because 18 is closer to 18.9 than 20 (an interval of 2 applies in the range of 18 to 42).

Related Documentation

- Schedulers Overview on page 1473
- Excess Bandwidth Sharing and Minimum Logical Interface Shaping on page 1480
- Excess Bandwidth Sharing Proportional Rates on page 1481
- Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces on page 1483
- Shared Bandwidth Among Logical Interfaces on page 1484

Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces

Logical interfaces with only shaping rates (PIRs) or unshaped logical interfaces (units) are given a weight of 10. A logical interface with a small guaranteed rate (CIR) might get an overall weight less than 10. To allocate a higher share of the excess bandwidth to logical interfaces with a small guaranteed rate in comparison to the logical interfaces with only shaping rates configured, a minimum weight of 20 is given to the logical interfaces with guaranteed rates configured.

For example, a logical interface configuration with five units is shown in Table 132.

Table 132: Allocating Weights with PIR and CIR on Logical Interfaces

Logical Interface (Unit)	Traffic Control Profile	WRR Percentages	Weights
Unit 1	PIR 100 Mbps	95, 0, 0, 5	10, 1, 1, 1
Unit 2	CIR 20 Mbps	25, 25, 25, 25	64, 64, 64, 64
Unit 3	PIR 40 Mbps, CIR 20 Mbps	50, 30, 15, 5	128, 76, 38, 13
Unit 4	Unshaped	95, 0, 0, 5	10, 1, 1, 1
Unit 5	CIR 1 Mbps	95, 0, 0, 5	10, 1, 1, 1

The weights for these units are calculated as follows:

- The excess bandwidth-sharing proportional rate is the maximum CIR among all the logical interfaces which is 20 Mbps (unit 2).
- Unit 1 has a PIR and unit 4 is unshaped. The weight for these units is 10.
- The weight for unit 1 queue 0 is 9.5 (10 x 95%), which translates to a hardware weight of 10.
- The weight for unit 1 queue 1 is 0 (0 x 0%) but though the weight is zero, a weight of 1 is assigned to give minimal bandwidth to queues with zero WRR.
- Unit 5 has a very small CIR (1 Mbps), and a weight of 20 is assigned to units with a small CIR.

Copyright © 2016, Juniper Networks, Inc.

- The weight for unit 5 queue 0 is 19 (20 x 95%), which translates to a hardware weight of 18.
- Unit 3 has a CIR of 20 Mbps, which is the same as the excess bandwidth-sharing proportional rate, so it has a total weight of 255.
- The weight of unit 3 queue 0 is 127.5 (255 x 50%), which translates to a hardware weight of 128.

Related Documentation

- Schedulers Overview on page 1473
- Excess Bandwidth Sharing and Minimum Logical Interface Shaping on page 1480
- Excess Bandwidth Sharing Proportional Rates on page 1481
- Calculated Weights Mapped to Hardware Weights on page 1482
- Shared Bandwidth Among Logical Interfaces on page 1484

Shared Bandwidth Among Logical Interfaces

As a simple example showing how bandwidth is shared among the logical interfaces, assume that all traffic is sent on queue 0. Assume also that there is a 40-Mbps load on all of the logical interfaces. Configuration details are shown in Table 133.

Table 133: Example of Shared Bandwidth Among Logical Interfaces

Logical Interface (Unit)	Traffic Control Profile	WRR Percentages	Weights
Unit 1	PIR 100 Mbps	95, 0, 0, 5	10, 1, 1, 1
Unit 2	CIR 20 Mbps	25, 25, 25, 25	64, 64, 64, 64
Unit 3	PIR 40 Mbps, CIR 20 Mbps	50, 30, 15, 5	128, 76, 38, 13
Unit 4	Unshaped	95, 0, 0, 5	10, 1, 1, 1

When the port is shaped at 40 Mbps, because units 2 and 3 have a guaranteed rate (CIR) configured, both units 2 and 3 get 20 Mbps of shared bandwidth.

When the port is shaped at 100 Mbps, because units 2 and 3 have a guaranteed rate (CIR) configured, each of them can transmit 20 Mbps. On units 1, 2, 3, and 4, the 60 Mbps of excess bandwidth is shaped according to the values shown in Table 134.

Table 134: First Example of Bandwidth Sharing

Logical Interface (Unit)	Calculation	Bandwidth
1	10 / (10+64+128+10) x 60 Mbps	2.83 Mbps
2	64 / (10+64+128+10) x 60 Mbps	18.11 Mbps
3	128 / (10+64+128+10) x 60 Mbps	36.22 Mbps

Copyright © 2016, Juniper Networks, Inc.

Table 134: First Example of Bandwidth Sharing *(continued)*

Logical Interface (Unit)	Calculation	Bandwidth
4	10 (10+64+128+10) x 60 Mbps	2.83 Mbps

However, unit 3 only has 20 Mbps extra (PIR and CIR) configured. This means that the leftover bandwidth of 16.22 Mbps (36.22 Mbps – 20 Mbps) is shared among units 1, 2, and 4. This is shown in Table 135.

Table 135: Second Example of Bandwidth Sharing

Logical Interface (Unit)	Calculation	Bandwidth
1	10 / (10+64+128+10) x 16.22 Mbps	1.93 Mbps
2	64 / (10+64+128+10) x 16.22 Mbps	12.36 Mbps
4	10 (10+64+128+10) x 16.22 Mbps	1.93 Mbps

Finally, Table 136 shows the resulting allocation of bandwidth among the logical interfaces when the port is configured with a 100-Mbps shaping rate.

Table 136: Final Example of Bandwidth Sharing

Logical Interface (Unit)	Calculation	Bandwidth
1	2.83 Mbps + 1.93 Mbps	4.76 Mbps
2	20 Mbps + 18.11 Mbps + 12.36 Mbps	50.47 Mbps
3	20 Mbps + 20 Mbps	40 Mbps
4	2.83 Mbps + 1.93 Mbps	4.76 Mbps

Related Documentation

- Schedulers Overview on page 1473
- Excess Bandwidth Sharing and Minimum Logical Interface Shaping on page 1480
- Excess Bandwidth Sharing Proportional Rates on page 1481
- Calculated Weights Mapped to Hardware Weights on page 1482
- Weight Allocation with Only Shaping Rates or Unshaped Logical Interfaces on page 1483

Example: Configuring Class-of-Service Schedulers

This example shows how to configure CoS schedulers on a device.

- Requirements on page 1486
- Overview on page 1486

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1487
- Verification on page 1489

Requirements

Before you begin, determine the buffer size allocation method to use. See "Scheduler Buffer Size Overview" on page 1489.

Overview

An individual device interface has multiple queues assigned to store packets temporarily before transmission. To determine the order in which to service the queues, the device uses a round-robin scheduling method based on priority and the queue's weighted round-robin (WRR) credits. Junos OS schedulers allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission.

You configure schedulers to assign resources, priorities, and drop profiles to output queues. By default, only queues 0 and 3 have resources assigned.

NOTE: **Juniper Network devices support hierarchical schedulers, including per-unit schedulers.**

In this example, you configure a best-effort scheduler called be-scheduler. You set the priority as low and the buffer size to 40. You set the be-scheduler transmit-rate remainder percentage to 40. You configure an expedited forwarding scheduler called ef-scheduler and set the priority as high and the buffer size to 10. You set the ef-scheduler transmit-rate remainder percentage to 50.

Then you configure an assured forwarding scheduler called af-scheduler and set the priority as high and buffer size to 45. You set an assured forwarding scheduler transmit rate to 45. You then configure a drop profile map for assured forwarding as low and high priority. (DiffServ can have a RED drop profile associated with assured forwarding.)

Finally, you configure a network control scheduler called nc-scheduler and set the priority as low and buffer size to 5. You set a network control scheduler transmit rate to 5.

Table 137 shows the schedulers created in this example.

Table 137: Sample Schedulers

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Allocated Portion of Remainder (Transmit Rate)
be-scheduler	Best-effort traffic	Low	40 percent	40 percent
ef-scheduler	Expedited forwarding traffic	High	10 percent	50 percent

Copyright © 2016, Juniper Networks, Inc.

Table 137: Sample Schedulers *(continued)*

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Allocated Portion of Remainder (Transmit Rate)
af-scheduler	Assured forwarding traffic	High	45 percent	—
nc-scheduler	Network control traffic	Low	5 percent	—

Configuration

CLI Quick Configuration

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

```
set class-of-service schedulers be-scheduler priority low buffer-size percent 40
set class-of-service schedulers be-scheduler transmit-rate remainder 40
set class-of-service schedulers ef-scheduler priority high buffer-size percent 10
set class-of-service schedulers ef-scheduler transmit-rate remainder 50
set class-of-service schedulers af-scheduler priority high buffer-size percent 45
set class-of-service schedulers af-scheduler transmit-rate percent 45
set class-of-service schedulers af-scheduler drop-profile-map loss-priority low protocol
    any drop-profile af-normal
set class-of-service schedulers af-scheduler drop-profile-map loss-priority high protocol
    any drop-profile af-with-PLP
set class-of-service schedulers nc-scheduler priority low buffer-size percent 5
set class-of-service schedulers nc-scheduler transmit-rate percent 5
```

Step-by-Step Procedure

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

To configure CoS schedulers:

1. Configure a best-effort scheduler.

   ```
   [edit]
   user@host# edit class-of-service schedulers be-scheduler
   ```

2. Specify a best-effort scheduler priority and buffer size.

   ```
   [edit class-of-service schedulers be-scheduler]
   user@host# set priority low
   user@host# set buffer-size percent 40
   ```

3. Configure a remainder option for a best-effort scheduler transmit rate.

   ```
   [edit class-of-service schedulers be-scheduler]
   user@host# set transmit-rate remainder 40
   ```

4. Configure an expedited forwarding scheduler.

   ```
   [edit]
   user@host# edit class-of-service schedulers ef-scheduler
   ```

Copyright © 2016, Juniper Networks, Inc.

5. Specify an expedited forwarding scheduler priority and buffer size.

```
[edit class-of-service schedulers ef-scheduler]
user@host# set priority high
user@host# set buffer-size percent 10
```

6. Configure a remainder option for an expedited forwarding scheduler transmit rate.

```
[edit class-of-service schedulers ef-scheduler]
user@host# set transmit-rate remainder 50
```

7. Configure an assured forwarding scheduler.

```
[edit]
user@host# edit class-of-service schedulers af-scheduler
```

8. Specify an assured forwarding scheduler priority and buffer size.

```
[edit class-of-service schedulers af-scheduler]
user@host# set priority high
user@host# set buffer-size percent 45
```

9. Configure an assured forwarding scheduler transmit rate.

```
[edit class-of-service schedulers af-scheduler]
user@host# set transmit-rate percent 45
```

10. Configure a drop profile map for assured forwarding low and high priority.

```
[edit class-of-service schedulers af-scheduler]
user@host# set drop-profile-map loss-priority low protocol any drop-profile
  af-normal
user@host# set drop-profile-map loss-priority high protocol any drop-profile
  af-with-PLP
```

11. Configure a network control scheduler.

```
[edit]
user@host# edit class-of-service schedulers nc-scheduler
```

12. Specify a network control scheduler priority and buffer size.

```
[edit class-of-service schedulers nc-scheduler]
user@host# set priority low
user@host# set buffer-size percent 5
```

13. Configure a network control scheduler transmit rate.

```
[edit class-of-service schedulers nc-scheduler]
user@host# set transmit-rate percent 5
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  schedulers {
  be-scheduler {
  transmit-rate remainder 40;
    buffer-size percent 40;
```


Copyright © 2016, Juniper Networks, Inc.

```
    priority low;
    }
    ef-scheduler {
      transmit-rate remainder 50;
      buffer-size percent 10;
    priority high;
    }
    af-scheduler {
      transmit-rate percent 45;
      buffer-size percent 45;
      priority high;
        drop-profile-map loss-priority low protocol any drop-profile af-normal;
      drop-profile-map loss-priority high protocol any drop-profile af-with-PLP;
      }
      nc-scheduler {
        transmit-rate percent 5;
        buffer-size percent 5;
      priority low;
    }
  }
```

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

Verification

Verifying Schedulers Configuration

Purpose Verify that the schedulers are configured properly.

Action From operational mode, enter the **show class-of-service** command.

Related Documentation

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Example: Configuring and Applying Scheduler Maps on page 1500
- Transmission Scheduling Overview on page 1479

Scheduler Buffer Size Overview

Large bursts of traffic from faster interfaces can cause congestion and dropped packets on slower interfaces that have small delay buffers. For example, a Juniper Networks device operating at the edge of the network can drop a portion of the burst traffic it receives on a channelized T1/E1 interface from a Fast Ethernet or Gigabit Ethernet interface on a router at the network core. On Juniper Networks devices, large delay buffers can be configured for both channelized T1/E1 and nonchannelized T1/E1 interfaces.

To ensure that traffic is queued and transmitted properly on slower interfaces, you can configure a buffer size larger than the default maximum.

Copyright © 2016, Juniper Networks, Inc.

This section contains the following topics:

- Maximum Delay Buffer Sizes Available to Channelized T1/E1 Interfaces on page 1490
- Maximum Delay Buffer Size for vSRX Interfaces on page 1491
- Delay Buffer Size Allocation Methods on page 1491
- Delay Buffer Sizes for Queues on page 1492

Maximum Delay Buffer Sizes Available to Channelized T1/E1 Interfaces

When you enable the large delay buffer feature on interfaces, a larger buffer is available for allocation to scheduler queues. The maximum delay buffer size that is available for an interface depends on the maximum available delay buffer time and the speed of the interface as shown in Table 138.

The default values are as follows:

- Clear-channel interface—The default delay buffer time is 500,000 microseconds (0.5 s).
- *N*xDS0 interface—The default delay buffer time is 1,200,000 microseconds (1.2 s).

Table 138: Maximum Available Delay Buffer Time by Channelized Interface and Rate

Effective Line Rate	Maximum Available Delay Buffer Time
< 4xDS0	4,000,000 microseconds (4 s)
< 8xDS0	2,000,000 microseconds (2 s)
< 16xDS0	1,000,000 microseconds (1 s)
<= 32xDS0	500,000 microseconds (0.5 s)
<= 10 mbps	400,000 microseconds (0.4 s)
<= 20 mbps	300,000 microseconds (0.3 s)
<= 30 mbps	200,000 microseconds (0.2 s)
<= 40 mbps	150,000 microseconds (0.15 s)

You can calculate the maximum delay buffer size available for an interface, with the following formula:

interface speed x *maximum delay buffer time* = *maximum available delay buffer size*

For example, the following maximum delay buffer sizes are available to 1xDS0 and 2xDS0 interfaces:

1xDS0—64 Kbps x 4 s = 256 Kb (32 KB)

2xDS0—128 Kbps x 4 s = 512 Kb (64 KB)

Copyright © 2016, Juniper Networks, Inc.

If you configure a delay buffer size larger than the maximum, the system allows you to commit the configuration but displays a system log warning message and uses the default buffer size setting instead of the configured maximum setting.

Maximum Delay Buffer Size for vSRX Interfaces

For a vSRX virtual machine, 1 Gbps interfaces have a default delay buffer time of 1 second, a maximum buffer time of 32 seconds, and a maximum buffer size of 128 MB. Use the following CLI command to set the maximum delay buffer time for a scheduler:

set class-of-service schedulers be-scheduler buffer-size temporal 32m

On a logical vSRX interface, the delay buffer size for a queue that does not have a specific shaping rate acts as a guaranteed minimum buffer size, and the queue is allowed to grow without any packet drops if the queue size is less then the guaranteed buffer size.

The sum of the guaranteed delay buffer sizes for all the queues acts as a pool that can be shared among the queues that do not have a specific shaping rate.

NOTE: **The delay buffers are used to control the size of the queues, but do not represent actual memory. The packet buffer pool contains the actual memory used to store packets.**

Packets are tail-dropped (100% probability) from the queue if:

- The total buffer limit would be exceeded.
- The queue size would exceed the total free buffer size.
- The packet buffer pool is less then 25% free and the queue exceeds the guaranteed minimum buffer size.
- The packet buffer pool is only 5% free (or less).

Packets also can be dropped by a RED profile (RED-dropped) if the queue size exceeds the guaranteed buffer size. The queue size will be restricted to be less then or equal to the free shared buffers available.

Delay Buffer Size Allocation Methods

You can specify delay buffer sizes for each queue using schedulers. The queue buffer can be specified as a period of time (microseconds) or as a percentage of the total buffer or as the remaining buffer. Table 139 shows different methods that you can specify for buffer allocation in queues.

Table 139: Delay Buffer Size Allocation Methods

Buffer Size Allocation Method	Description
Percentage	A percentage of the total buffer.

Copyright © 2016, Juniper Networks, Inc.

Table 139: Delay Buffer Size Allocation Methods *(continued)*

Buffer Size Allocation Method	Description
Temporal	A period of time, value in microseconds. When you configure a temporal buffer, you must also configure a transmit rate. The system calculates the queue buffer size by multiplying the available bandwidth of the interface times the configured temporal value and transmit rate. When you specify a temporal method, the drop profile is assigned a static buffer and the system starts dropping packets once the queue buffer size is full. By default, the other buffer types are assigned dynamic buffers that use surplus transmission bandwidth to absorb bursts of traffic.
Remainder	The remaining buffer available. The remainder is the percentage buffer that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 3 to keep the default allotment of 5 percent, and assign the remainder to queue 7, then queue 7 uses approximately 55 percent of the delay buffer. Optionally, you can specify the percentage of the remainder to be used for allocating the buffer size of the scheduler on a prorated basis. If the remainder percentage is not specified, the remainder value will be shared equally.

Delay Buffer Sizes for Queues

You specify delay buffer sizes for queues using schedulers. The system calculates the buffer size of a queue based on the buffer allocation method you specify for it in the scheduler. See Table 139 for different buffer allocation methods and Table 140 for buffer size calculations.

Table 140: Delay Buffer Allocation Method and Queue Buffer

Buffer Size Allocation Method	Queue Buffer Calculation	Example
Percentage	*available interface bandwidth* x *configured buffer size percentage* x *maximum delay buffer time* = *queue buffer*	Suppose you configure a queue on a 1xDS0 interface to use 30 percent of the available delay buffer size. The system uses the maximum available delay buffer time (4 seconds) and allocates the queue 9600 bytes of delay buffer: 64 Kbps x 0.3 x 4 s=76,800 bits=9,600 bytes

Copyright © 2016, Juniper Networks, Inc.

Table 140: Delay Buffer Allocation Method and Queue Buffer *(continued)*

Buffer Size Allocation Method	Queue Buffer Calculation	Example
Temporal	*available interface bandwidth* x *configured transmit rate percentage* x *configured temporal buffer size* = *queue buffer*	Suppose you configure a queue on a 1xDS0 interface to use 300,000,000 microseconds (3 seconds) of delay buffer, and you configure the transmission rate to be 20 percent. The queue receives 4800 bytes of delay buffer: 64 Kbps x 0.2 x 3 s=38,400 bits=4,800 bytes If you configure a temporal value that exceeds the maximum available delay buffer time, the queue is allocated the buffer remaining after buffers are allocated for the other queues. Suppose you configure a temporal value of 6,000,000 microseconds on a 1xDS0 interface. Because this value exceeds the maximum allowed value of 4,000,000 microseconds, the queue is allocated the remaining delay buffer.

When you specify the buffer size as a percentage, the system ignores the transmit rate and calculates the buffer size based only on the buffer size percentage.

Related Documentation

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Example: Configuring Class-of-Service Schedulers on page 1485
- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Example: Configuring and Applying Scheduler Maps on page 1500
- Transmission Scheduling Overview on page 1479

Example: Configuring a Large Delay Buffer on a Channelized T1 Interface

This example shows how to configure a large delay buffer on a channelized T1 interface to help slower interfaces avoid congestion and packet dropping when they receive large bursts of traffic.

- Requirements on page 1493
- Overview on page 1494
- Configuration on page 1494
- Verification on page 1495

Requirements

Before you begin, enable the large buffer feature on the channelized T1/E1 PIM and then configure a buffer size for each queue in the CoS scheduler. See "Scheduler Buffer Size Overview" on page 1489.

Copyright © 2016, Juniper Networks, Inc.

Overview

On devices, you can configure large delay buffers on channelized T1/E1 interfaces. Each channelized T1/E1 interface can be configured as a single clear channel, or for channelized (NxDS0) operation, where N denotes channels 1 to 24 for a T1 interface and channels 1 to 32 for an E1 interface.

In this example, you specify a queue buffer of 30 percent in scheduler **be-scheduler** and associate the scheduler to a defined forwarding class **be-class** using scheduler map **large-buf-sched-map**. Finally, you apply the scheduler map to channelized T1 interface **t1-3/0/0**.

Configuration

CLI Quick Configuration

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

```
set chassis fpc 3 pic 0 q-pic-large-buffer
set class-of-service schedulers be-scheduler buffer-size percent 30
set class-of-service scheduler-maps large-buf-sched-map forwarding-class be-class
  scheduler be-scheduler
set class-of-service interfaces t1-3/0/0 unit 0 scheduler-map large-buf-sched-map
```

Step-by-Step Procedure

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

To configure a large delay buffer on a channelized T1 interface:

1. Enable the large buffer size feature on the channelized T1 interface.

   ```
   [edit]
   user@host# edit chassis
   user@host# set fpc 3 pic 0 q-pic-large-buffer
   ```

2. Create best-effort traffic and specify a buffer size.

   ```
   [edit]
   user@host# edit class-of-service
   user@host# set schedulers be-scheduler buffer-size percent 30
   ```

3. Configure the scheduler map to associate schedulers with defined forwarding classes.

   ```
   [edit class-of-service]
   user@host# set scheduler-maps large-buf-sched-map forwarding-class be-class
     scheduler be-scheduler
   ```

4. Apply the scheduler map to the channelized T1 interface.

   ```
   [edit class-of-service]
   user@host# set interfaces t1-3/0/0 unit 0 scheduler-map large-buf-sched-map
   ```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show class-of-service** and **show chassis** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  interfaces {
  t1-3/0/0 {
  unit 0 {
    scheduler-map large-buf-sched-map;
    }
  }
  }
  scheduler-maps {
    large-buf-sched-map {
    forwarding-class be-class scheduler be-scheduler;
  }
  }
  schedulers {
    be-scheduler {
    buffer-size percent 30;
  }
  }
  [edit]
    user@host# show chassis
    fpc 3 {
    pic 0 {
    q-pic-large-buffer;
  }
}
```

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

Verification

Verifying Large Delay Buffers Configuration

Purpose Verify that the large delay buffers are configured properly.

Action From configuration mode, enter the **show class-of-service** and **show chassis** commands.

Related Documentation

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Example: Configuring Class-of-Service Schedulers on page 1485
- Example: Configuring and Applying Scheduler Maps on page 1500
- Transmission Scheduling Overview on page 1479

Copyright © 2016, Juniper Networks, Inc.

Configuring Large Delay Buffers in CoS

You can configure very large delay buffers using the **buffer-size-temporal** command combined with the **q-pic-large-buffer** command. The **buffer-size temporal** option in combination with **q-pic-large-buffer** can create extra-large delay buffer allocations for one or several queues on an interface.

NOTE: If the configured buffer size is too low, the buffer size for the forwarding class defaults to 9192 and the following log message is displayed: "fwdd_cos_set_delay_bandwidth:queue:16 delay buffer size (1414) too low, setting to default 9192."

Configuring Large Delay Buffers

The following configuration applies to the examples that follow:

1. Configure two VLANs (one ingress, one egress) on one interface. No interface shaping rate is initially defined for this configuration.

   ```
   [edit]
   set interfaces ge-0/0/3 per-unit-scheduler
   set interfaces ge-0/0/3 vlan-tagging
   set interfaces ge-0/0/3 unit 102 vlan-id 102
   set interfaces ge-0/0/3 unit 102 family inet address 203.0.113.2/24
   set interfaces ge-0/0/3 unit 201 vlan-id 201
   set interfaces ge-0/0/3 unit 201 family inet address 198.51.100.2/24
   set routing-options static route 192.02.1/32 next-hop 198.51.100.3
   ```

2. Enable the **q-pic-large-buffer** option on the same PIC, in addition to the **buffer-size temporal** option on the queue, to create a large buffer on the queue:

   ```
   [edit]
   set chassis fpc 0 pic 0 q-pic-large-buffer
   ```

   NOTE: The CLI does not provide a warning when you use **buffer-size temporal** without **q-pic-large-buffer**. When you use **buffer-size temporal**, verify that the configuration also includes the **q-pic-large buffer** command.

3. Define four forwarding-classes (queue names) for the four queues:

   ```
   [edit]
   set class-of-service forwarding-classes queue 0 be-Queue0
   set class-of-service forwarding-classes queue 1 video-Queue1
   set class-of-service forwarding-classes queue 2 voice-Queue2
   set class-of-service forwarding-classes queue 3 nc-Queue3
   ```

4. Configure the forwarding classes (queue names) included in a scheduler map, applied to the egress VLAN:

   ```
   [edit]
   set class-of-service interfaces ge-0/0/3 unit 201 scheduler-map schedMapM
   ```


Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service scheduler-maps schedMapM forwarding-class be-Queue0
  scheduler be-Scheduler0
set class-of-service scheduler-maps schedMapM forwarding-class video-Queue1
  scheduler video-Scheduler1
set class-of-service scheduler-maps schedMapM forwarding-class voice-Queue2
  scheduler voice-Scheduler2
set class-of-service scheduler-maps schedMapM forwarding-class nc-Queue3
  scheduler nc-Scheduler3
```

5. Set the queue priorities. Only queue priorities are initially defined, not transmit rates or buffer sizes.

```
[edit]
set class-of-service schedulers be-Scheduler0 priority low
set class-of-service schedulers video-Scheduler1 priority medium-low
set class-of-service schedulers voice-Scheduler2 priority medium-high
set class-of-service schedulers nc-Scheduler3 priority high
```

Example: Simple Configuration Using Four Queues

This configuration allocates 12,500,000 bytes of buffer for each of the four queues. To avoid exceeding the limits of the delay buffer calculation, this initial example has no interface shaping rate, scheduler transmit rate, or scheduler buffer size percent configuration.

1. Specify the maximum 4-second delay buffer on each of the four queues:

```
[edit]
set class-of-service schedulers be-Scheduler0 buffer-size temporal 4m
set class-of-service schedulers video-Scheduler1 buffer-size temporal 4m
set class-of-service schedulers voice-Scheduler2 buffer-size temporal 4m
set class-of-service schedulers nc-Scheduler3 buffer-size temporal 4m
```

   Specifying **buffer-size temporal** on some or all queues uses implicit (default) or explicit transmit rate percentages as the buffer-size percentages of the temporal values for those queues. Because there are no explicitly specified transmit rate percentages, divide 100 percent by the number of configured queues (queues with schedulers configured in the scheduler map) to get the implicit (default) per-queue transmit rate percentages. Each queue gets an implicit (default) transmit rate of 100% / 4 = 25%.

   In this example, specifying the maximum 4-second delay on each queue, with no shaping rate on the interface and implicit (default) per-queue transmit rates of 25 percent, the total buffer for all temporal 4m queues on an interface = 4 seconds * 100,000,000 maximum interface bps / 8 bits/byte = 4 seconds * 12,500,000 bytes = 50,000,000 bytes. Each queue specifying temporal 4m gets 25% * 50,000,000 = 12,500,000 bytes.

2. Add a shaping rate of 4 Mbps to the interface:

```
[edit]
set class-of-service interfaces ge-0/0/3 unit 201 shaping-rate 4m
```

   The total buffer for all temporal 4m queues on an interface = 4 sec * 4,000,000 bps shaping-rate / 8 bits/byte = 4 sec * 500,000 bytes = 2,000,000 bytes. Therefore, each queue specifying temporal 4m receives 25% * 2,000,000 = 500,000 bytes.

Copyright © 2016, Juniper Networks, Inc.

When using **buffer-size temporal** on any interface queues, if you also use the **transmit-rate percent** command, or the **buffer-size percent** command, or both commands, on any of the interface queues, the buffer size calculations become more complex and the limits of available queue depth might be reached. If the configuration attempts to exceed the available memory, then at commit time two system log messages appear in the **/var/log/messages** file, the interface class-of-service configuration is ignored, and the interface class-of-service configuration reverts to the two-queue defaults:

```
Mar 11 11:02:10.239  elma-n4 elma-n4 COSMAN_FWDD: queue mem underflow for ge-0/0/3
Mar 11 11:02:10.240  elma-n4 elma-n4 cosman_compute_install_sched_params: Failed
 to compute scheduler params for ge-0/0/3.Hence retaining defaults
```

When configuring **buffer-size temporal** along with **transmit-rate percent** or **buffer-size percent**, or both, you must monitor the system log to see whether the available queue depth limit has been reached.

Example: Using buffer-size temporal with Explicit transmit-rate percent Commands

To add explicit transmit rates to all four queues:

```
[edit]
set class-of-service schedulers be-Scheduler0 transmit-rate percent 10
set class-of-service schedulers video-Scheduler1 transmit-rate percent 25
set class-of-service schedulers voice-Scheduler2 transmit-rate percent 25
set class-of-service schedulers nc-Scheduler3 transmit-rate percent 40
```

For example, if an interface is shaped to 4 Mbps, the transmit rate percentage of 10 for a queue means that the bandwidth share for the specific queue is 0.4 Mbps. The queues are allocated portions of the 2,000,000 bytes of total buffer available for temporal queues on this interface, proportionally to their transmit rates. The four queues get 200,000, 500,000, 500,000, and 800,000 bytes of delay buffer, respectively.

To avoid exceeding the queue depth limits and triggering system log messages and default configuration behavior, when configuring queues with **buffer-size temporal** and **transmit rate percent** and other (non-temporal) queues with **buffer-size percent**, the following configuration rule must be followed: When one or more queues on an interface are configured with **buffer-size temporal**, the sum of the temporal queues explicitly configured transmit rate percentages plus other non-temporal queues explicitly configured buffer size percentages must not exceed 100 percent.

If the total of the temporal queues transmit rate percentages and the non-temporal queues buffer-size percentages exceeds 100 percent, the **queue mem underflow** and **Failed to compute scheduler params** system log messages appear in the messages log, the explicitly configured CLI CoS configuration for the interface is ignored, and the interface reverts to a two-queue default CoS configuration.

When **buffer-size temporal** is specified on a queue, if **transmit-rate percent** is also configured on the same queue, the queue depth configured is based on the fractional bandwidth for the queue as obtained by the specified **transmit-rate percent**.

In addition to temporal delay times specified for one or more queues using buffer size temporal, there is another delay time automatically computed for the entire interface.

Copyright © 2016, Juniper Networks, Inc.

This interface delay time is distributed across all non-temporal queues, proportionally to their implicit (default) or explicit transmit-rate percentages. If **q-pic-large-buffer** is not enabled, the interface delay time defaults to 100 ms. As shown in Table 141, when **q-pic-large-buffer** is enabled, interface delay time is calculated according to configured shaping rate for the interface. Because the shaping-rate configured in the example above was 4 Mbps (> 2,048,000 bps), the interface delay time for the configuration is 100 msec.

Table 141: Interface Delay Times Enabled By q-pic-large-buffer

Configured Shaping Rate (bps)	Interface Delay Time (msec) Used for Non-Temporal Queues with q-pic-large-buffer Enabled	Default Delay Time Used (msec) Without q-pic-large-buffer
64,000-255,999	4000	100
256,000 - 511,999	2000	100
512,000 - 102,3999	1000	100
1,024,000 - 2,047,999	500	100
>= 2,048,000	100	100

This example properly computes the delay buffer limits on both temporal and non-temporal queues:

1. Substitute **buffer-size percent** for **buffer-size temporal** on queues 0 and 1:

   ```
   [edit]
   delete class-of-service schedulers be-Scheduler0 buffer-size temporal 4m
   delete class-of-service schedulers video-Scheduler1 buffer-size temporal 4m
   set class-of-service schedulers be-Scheduler0 buffer-size percent 10
   set class-of-service schedulers video-Scheduler1 buffer-size percent 25
   ```

   This deletes the requirement for hard-specified 4 seconds of buffering and replaces it with a proportional limit of 10 percent (or 25 percent) of the total interface delay time for the non-temporal queues. In both cases, the queue depth is calculated based on the share of the interface bandwidth for the specific queues. Total Interface Non-Temporal Queue Memory = shaping-rate * Interface delay time (Table 1) = 4 Mbps * 0.1 seconds = 500,000 bytes per second * 0.1 seconds = 50,000 bytes, therefore queues 0 and 1 get 10% * 50,000 = 5000 bytes and 25% * 50,000 = 12,500 bytes of delay buffer, respectively.

2. Configure **buffer-size temporal** on queues 2 and 3:

   ```
   [edit]
   set class-of-service schedulers voice-Scheduler2 buffer-size temporal 4m
   set class-of-service schedulers voice-Scheduler2 transmit-rate percent 25
   set class-of-service schedulers nc-Scheduler3 buffer-size temporal 4m
   set class-of-service schedulers nc-Scheduler3 transmit-rate percent 40
   ```

   Queues 2 and 3 still get 500,000 and 800,000 bytes of delay buffer, respectively, as previously calculated. This configuration obeys the rule that the sum of the temporal

Copyright © 2016, Juniper Networks, Inc.

queues transmit rate percentages (25% + 40% = 65%), plus the non-temporal queues buffer size percentages (10% + 25% = 35%) do not exceed 100% (65% + 35% <= 100%).

The following example exceeds the delay buffer limit, triggering the system log messages and the default, two-queue class-of-service behavior.

Increase the buffer-size percentage from 25 percent to 26 percent for non-temporal queue 1:

```
[edit]
set class-of-service schedulers video-Scheduler1 buffer-size percent 26
```

This violates the configuration rule that the sum of the non-temporal queues buffer-size percentages (10% + 26% = 36%), plus the temporal queues transmit rate percentages (25% + 40% = 65%) now exceed 100% (36% + 65% = 101%). Therefore, the following two system log messages appear in the **/var/log/messages** file:

```
Mar 23 18:08:23  elma-n4 elma-n4 COSMAN_FWDD: %PFE-3: queue mem underflow for
ge-0/0/3 q_num(3)
Mar 23 18:08:23  elma-n4 elma-n4 cosman_compute_install_sched_params: %PFE-3:
Failed to compute scheduler params for ge-0/0/3.Hence retaining defaults
```

When the delay buffer limits are exceeded, the CLI-configured class-of-service settings are not used and the default class-of-service configuration (the default scheduler-map) is assigned to the interface. This uses two queues: the forwarding-class best-effort (queue 0) has transmit rate percent 95 and buffer-size percent 95 and the forwarding-class network-control (queue 3) has the transmit rate percent 5 and buffer-size percent 5.

```
queue 0: 1,187,500 Bytes
queue 1:     9,192 Bytes
queue 2:     9,192 Bytes
queue 3:    62,500 Bytes
```

Related Documentation

- Example: Configuring and Applying Scheduler Maps on page 1500
- Scheduler Buffer Size Overview on page 1489

Example: Configuring and Applying Scheduler Maps

This example shows how to configure and apply a scheduler map to a device's interface.

- Requirements on page 1500
- Overview on page 1501
- Configuration on page 1501
- Verification on page 1503

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Create and configure the forwarding classes. See *Configuring Forwarding Classes*.
- Create and configure the schedulers. See "Example: Configuring Class-of-Service Schedulers" on page 1485.

Overview

After you define a scheduler, you can include it in a scheduler map, which maps a specified forwarding class to a scheduler configuration. You configure a scheduler map to assign a forwarding class to a scheduler, and then apply the scheduler map to any interface that must enforce DiffServ CoS.

After they are applied to an interface, the scheduler maps affect the hardware queues, packet schedulers, and RED drop profiles.

In this example, you create the scheduler map diffserv-cos-map and apply it to the device's Ethernet interface ge-0/0/0. The map associates the mf-classifier forwarding classes to the schedulers as shown in Table 142.

Table 142: Sample diffserv-cos-map Scheduler Mapping

mf-classifier Forwarding Class	For CoS Traffic Type	diffserv-cos-map Scheduler
be-class	Best-effort traffic	be-scheduler
ef-class	Expedited forwarding traffic	ef-scheduler
af-class	Assured forwarding traffic	af-scheduler
nc-class	Network control traffic	nc-scheduler

Configuration

CLI Quick Configuration

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

```
set class-of-service scheduler-maps diffserv-cos-map forwarding-class be-class scheduler
    be-scheduler
set class-of-service scheduler-maps diffserv-cos-map forwarding-class ef-class scheduler
    ef-scheduler
set class-of-service scheduler-maps diffserv-cos-map forwarding-class af-class scheduler
    af-scheduler
set class-of-service scheduler-maps diffserv-cos-map forwarding-class nc-class scheduler
    nc-scheduler
set class-of-service interfaces ge-0/0/0 unit 0 scheduler-map diffserv-cos-map
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure and apply a scheduler map to a device's interface:

1. Configure a scheduler map for DiffServ CoS.

   ```
   [edit class-of-service]
   user@host# edit scheduler-maps diffserv-cos-map
   ```

2. Configure a best-effort forwarding class and scheduler.

   ```
   [edit class-of-service scheduler-maps diffserv-cos-map]
   user@host# set forwarding-class be-class scheduler be-scheduler
   ```

3. Configure an expedited forwarding class and scheduler.

   ```
   [edit class-of-service scheduler-maps diffserv-cos-map]
   user@host# set forwarding-class ef-class scheduler ef-scheduler
   ```

4. Configure an assured forwarding class and scheduler.

   ```
   [edit class-of-service scheduler-maps diffserv-cos-map]
   user@host# set forwarding-class af-class scheduler af-scheduler
   ```

5. Configure a network control class and scheduler.

   ```
   [edit class-of-service scheduler-maps diffserv-cos-map]
   user@host# set forwarding-class nc-class scheduler nc-scheduler
   ```

6. Apply the scheduler map to an interface.

   ```
   [edit class-of-service]
   user@host# set interfaces ge-0/0/0 unit 0 scheduler-map diffserv-cos-map
   ```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
interfaces {
   ge-0/0/0 {
      unit 0 {
      scheduler-map diffserv-cos-map;
   }
   }
}
scheduler-maps {
   diffserv-cos-map {
      forwarding-class be-class scheduler be-scheduler;
      forwarding-class ef-class scheduler ef-scheduler;
      forwarding-class af-class scheduler af-scheduler;
      forwarding-class nc-class scheduler nc-scheduler;
   }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying the Scheduler Map Configuration

Purpose Verify that scheduler maps are configured properly.

Action From operational mode, enter the **show class-of-service** command.

Related Documentation

- Schedulers Overview on page 1473
- Default Scheduler Settings on page 1478
- Transmission Scheduling Overview on page 1479

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 71

Removing Delays with Strict-Priority Queues

- Strict-Priority Queue Overview on page 1505
- Understanding Strict-Priority Queues on page 1506
- Example: Configuring Priority Scheduling on page 1506
- Example: Configuring Strict-Priority Queuing on page 1508

Strict-Priority Queue Overview

You can configure one queue per interface to have strict-priority, which causes delay-sensitive traffic, such as voice traffic, to be removed and forwarded with minimum delay. Packets that are queued in a strict-priority queue are removed before packets in other queues, including high-priority queues.

The strict-high-priority queuing feature allows you to configure traffic policing that prevents lower priority queues from being starved. The strict-priority queue does not cause starvation of other queues because the configured policer allows the queue to exceed the configured bandwidth only when other queues are not congested. If the interface is congested, the software directs strict-priority queues to the configured bandwidth.

To prevent queue starvation of other queues, you must configure an output (egress) policer that defines a limit for the amount of traffic that the queue can service. The software services all traffic in the strict-priority queue that is under the defined limit. When strict-priority traffic exceeds the limit, the policer marks the traffic in excess of the limit as out-of-profile. If the output port is congested, the software drops out-of-profile traffic.

You can also configure a second policer with an upper limit. When strict-priority traffic exceeds the upper limit, the software drops the traffic in excess of the upper limit, regardless of whether the output port is congested. This upper-limit policer is not a requirement for preventing starvation of the lower priority queues. The policer for the lower limit, which marks the packets as out-of-profile, is sufficient to prevent starvation of other queues.

Related Documentation

- Understanding Strict-Priority Queues on page 1506

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Priority Scheduling on page 1506
- Example: Configuring Strict-Priority Queuing on page 1508

Understanding Strict-Priority Queues

You use strict-priority queuing and policing as follows:

- Identify delay-sensitive traffic by configuring a behavior aggregate (BA) or multifield (MF) classifier.
- Minimize delay by assigning all delay-sensitive packets to the strict-priority queue.
- Prevent starvation on other queues by configuring a policer that checks the data stream entering the strict-priority queue. The policer defines a lower bound, marks the packets that exceed the lower bound as out-of-profile, and drops the out-of-profile packets if the physical interface is congested. If there is no congestion, the software forwards all packets, including the out-of-profile packets.
- Optionally, configure another policer that defines an upper bound and drops the packets that exceed the upper bound, regardless of congestion on the physical interface.

To configure strict-priority queuing and prevent starvation of other queues, include the **priority strict-high** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level and the **if-exceeding** and **then out-of-profile** statements at the **[edit firewall policer *policer-name*]** hierarchy level:

```
[edit class-of-service schedulers scheduler-name]
priority strict-high;

[edit firewall policer policer-name]
if-exceeding {
   bandwidth-limit bps;
   bandwidth-percent number;
   burst-size-limit bytes;
}
then out-of-profile;
```

Related Documentation

- Strict-Priority Queue Overview on page 1505
- Example: Configuring Priority Scheduling on page 1506
- Example: Configuring Strict-Priority Queuing on page 1508

Example: Configuring Priority Scheduling

This example shows how to configure priority scheduling so important traffic receives better access to the outgoing interface.

- Requirements on page 1507
- Overview on page 1507

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 1507
- Verification on page 1508

Requirements

Before you begin, review how to assign forwarding classes. See "Example: Assigning Forwarding Classes to Output Queues" on page 1460.

Overview

In this example, you configure CoS and a scheduler called be-sched with a medium-low priority. Then you configure scheduler map be-map to associate be-sched with the best-effort forwarding class. Finally, you apply be-map to interface ge-0/0/0.

Configuration

CLI Quick Configuration

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

```
set class-of-service schedulers be-sched priority medium-low
set class-of-service scheduler-maps be-map forwarding-class best-effort scheduler
  be-sched
set class-of-service interfaces ge-0/0/0 scheduler-map be-map
```

Step-by-Step Procedure

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

To configure priority scheduling:

1. Configure CoS and a scheduler.

   ```
   [edit]
   user@host# edit class-of-service
   user@host# edit schedulers be-sched
   ```

2. Set a priority.

   ```
   [edit class-of-service schedulers be-sched]
   user@host# set priority medium-low
   ```

3. Configure a scheduler map.

   ```
   [edit]
   user@host# edit class-of-service
   user@host# edit scheduler-maps be-map
   ```

4. Specify the best-effort forwarding class.

   ```
   [edit class-of-service scheduler-maps be-map]
   user@host# set forwarding-class best-effort scheduler be-sched
   ```

5. Apply best-effort map to an interface.

   ```
   [edit]
   user@host# edit class-of-service
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces ge-0/0/0 scheduler-map be-map
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  interfaces {
  ge-0/0/0 {
  scheduler-map be-map;
    }
    }
    scheduler-maps {
      be-map {
      forwarding-class best-effort scheduler be-sched;
    }
    }
    schedulers {
      be-sched {
    priority medium-low;
  }
}
```

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

Verification

Verifying Priority Scheduling

Purpose Verify that the priority scheduling is configured properly on a device.

Action From configuration mode, enter the **show class-of-service** command.

Related Documentation

- Strict-Priority Queue Overview on page 1505
- Understanding Strict-Priority Queues on page 1506
- Example: Configuring Strict-Priority Queuing on page 1508

Example: Configuring Strict-Priority Queuing

This example shows how to configure strict-priority queuing and prevent starvation of other queues.

- Requirements on page 1509
- Overview on page 1509
- Configuration on page 1509
- Verification on page 1516

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, review how to create and configure forwarding classes. See "Forwarding Classes Overview" on page 1453.

Overview

In this example, you create a BA classifier to classify traffic based on the IP precedence of the packet. The classifier defines IP precedence value 101 as voice traffic and 000 as data traffic. You assign forwarding-class priority queue 0 to voice traffic and queue 1 as data traffic. You then configure the scheduler map as corp-map and voice scheduler as voice-sched.

Then you set the priority for the voice traffic scheduler as strict-high and for the data traffic scheduler as strict-low. You apply the BA classifier to input interface ge-0/0/0 and apply the scheduler map to output interface e1-1/0/0. You then configure two policers called voice-drop and voice-excess. You set the burst size limit and bandwidth limit for voice-drop policer and for voice-excess policer. You then create a firewall filter that includes the new policers and add the policer to the term.

Finally, you apply the filter to output interface e1-1/0/1 and set the IP address as 203.0.113.1/24..

Configuration

- Configuring a BA Classifier on page 1509
- Configuring Forwarding Classes on page 1510
- Configuring a Scheduler Map on page 1511
- Configuring a Scheduler on page 1512
- Applying a BA Classifier to an Input Interface on page 1512
- Applying a Scheduler Map to an Output Interface on page 1513
- Configuring Two Policers on page 1514
- Applying a Filter to an Output Interface on page 1515

Configuring a BA Classifier

CLI Quick Configuration

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

```
set class-of-service classifiers inet-precedence corp-traffic forwarding-class voice-class
    loss-priority low code-points 101
set class-of-service classifiers inet-precedence corp-traffic forwarding-class data-class
    loss-priority high code-points 000
```

Step-by-Step Procedure

To configure a BA classifier:

1. Create a BA classifier and set the IP precedence value for voice traffic.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# edit class-of-service classifiers inet-precedence corp-traffic
  forwarding-class voice-class loss-priority low
user@host# set code-points 101
```

2. Create a BA classifier and set the IP precedence value for data traffic.

```
[edit]
user@host# edit class-of-service classifiers inet-precedence corp-traffic
  forwarding-class data-class loss-priority high
user@host# set code-points 000
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  classifiers {
  inet-precedence corp-traffic {
  forwarding-class voice-class {
    loss-priority low code-points 101;
    }
    forwarding-class data-class {
      loss-priority high code-points 000;
    }
  }
}
```

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

Configuring Forwarding Classes

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

```
set class-of-service forwarding-classes queue 0 voice-class
set class-of-service forwarding-classes queue 1 data-class
```

Step-by-Step Procedure To configure forwarding classes:

1. Assign priority queuing to voice traffic.

```
[edit]
user@host# set class-of-service forwarding-classes queue 0 voice-class
```

2. Assign priority queuing to data traffic.

```
[edit]
user@host# set class-of-service forwarding-classes queue 1 data-class
```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
forwarding-classes {
   queue 0 voice-class;
   queue 1 data-class;
}
```

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

Configuring a Scheduler Map

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

```
set class-of-service scheduler-maps corp-map forwarding-class voice-class scheduler
   voice-sched
set class-of-service scheduler-maps corp-map forwarding-class data-class scheduler
   data-sched
```

Step-by-Step Procedure To configure a scheduler map:

1. Configure a scheduler map and voice scheduler.

   ```
   [edit]
   user@host# edit class-of-service scheduler-maps corp-map forwarding-class
      voice-class
   user@host# set scheduler voice-sched
   ```

2. Configure a scheduler map and data scheduler.

   ```
   [edit]
   user@host# edit class-of-service scheduler-maps corp-map forwarding-class
      data-class
   user@host# set scheduler data-sched
   ```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
scheduler-maps {
   corp-map {
   forwarding-class voice-class scheduler voice-sched;
   forwarding-class data-class scheduler data-sched;
   }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Configuring a Scheduler

CLI Quick Configuration

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

```
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers data-sched priority lowset xxx
```

Step-by-Step Procedure

To configure schedulers:

1. Configure a voice traffic scheduler and set the priority.

   ```
   [edit]
   user@host# edit class-of-service schedulers voice-sched
   user@host# set priority strict-high
   ```

2. Configure a data traffic scheduler and set the priority.

   ```
   [edit]
   user@host# edit class-of-service schedulers data-sched
   user@host# set priority low
   ```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
schedulers {
   voice-sched {
   priority strict-high;
   }
   data-sched {
      priority low;
   }
}
```

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

Applying a BA Classifier to an Input Interface

CLI Quick Configuration

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

```
set class-of-service interfaces ge-0/0/0 unit 0 classifiers inet-precedence corp-traffic
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To apply a BA classifier to an input interface:

1. Configure an interface.

```
[edit]
user@host# edit class-of-service interfaces ge-0/0/0 unit 0
```

2. Apply a BA classifier to an input interface.

```
[edit class-of-service interfaces ge-0/0/0 unit 0]
user@host# set classifiers inet-precedence corp-traffic
```

Results

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

```
[edit]
  user@host# show class-of-service interfaces
  ge-0/0/0 {
  unit 0 {
  classifiers {
    inet-precedence corp-traffic;
    }
  }
}
```

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

Applying a Scheduler Map to an Output Interface

CLI Quick Configuration

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

```
set class-of-service interfaces e1-1/0/0 unit 0 scheduler-map corp-map
```

Step-by-Step Procedure

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

To apply the scheduler map to an output interface:

1. Configure an interface.

```
[edit]
user@host# edit class-of-service interfaces e1-1/0/0 unit 0
```

2. Apply a scheduler map to an output interface.

```
[edit class-of-service interfaces e1-1/0/0 unit 0]
user@host# set scheduler-map corp-map
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show class-of-service
  interfaces {
  e1-1/0/0 {
  unit 0 {
    scheduler-map corp-map;
    }
  }
}
```

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

Configuring Two Policers

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

```
set firewall policer voice-drop if-exceeding burst-size-limit 200000 bandwidth-limit
  2000000
set firewall policer voice-drop then discard
set firewall policer voice-excess if-exceeding burst-size-limit 200000 bandwidth-limit
  1000000
set firewall policer voice-excess then out-of-profile
set firewall filter voice-term term 01 from forwarding-class voice-class
set firewall filter voice-term term 01 then policer voice-drop next term
set firewall filter voice-term term 02 from forwarding-class voice-class
set firewall filter voice-term term 02 then policer voice-excess accept
```

Step-by-Step Procedure To configure two policers:

1. Configure a policer voice drop.

   ```
   [edit]
   user@host# edit firewall policer voice-drop
   user@host# set if-exceeding burst-size-limit 200000 bandwidth-limit 2000000
   user@host# set then discard
   ```

2. Configure a policer voice excess.

   ```
   [edit]
   user@host# edit firewall policer voice-excess
   user@host# set if-exceeding burst-size-limit 200000 bandwidth-limit 1000000
   user@host# set then out-of-profile
   ```

3. Create a firewall filter that includes the new policers.

   ```
   [edit]
   user@host# edit firewall filter voice-term term 01
   user@host# set from forwarding-class voice-class
   user@host# set then policer voice-drop next term
   ```


Copyright © 2016, Juniper Networks, Inc.

4. Add the policer to the term.

```
[edit]
user@host# edit firewall filter voice-term term 02
user@host# set from forwarding-class voice-class
user@host# set then policer voice-excess accept
```

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

```
[edit]
  user@host# show firewall
  policer voice-drop {
  if-exceeding {
  bandwidth-limit 2m;
    burst-size-limit 200k;
    }
    then discard;
    }
    policer voice-excess {
      if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 200k;
      }
        then out-of-profile;
      }
      filter voice-term {
        term 01 {
      from {
      forwarding-class voice-class;
      }
      then {
    policer voice-drop;
    next term;
      }
    }
    term 02 {
      from {
      forwarding-class voice-class;
      }
      then {
        policer voice-excess;
      accept;
    }
  }
}
```

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

Applying a Filter to an Output Interface

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces e1-1/0/1 unit 0 family inet filter output voice-term
set interfaces e1-1/0/1 unit 0 family inet address 203.0.113.1/24
```

Step-by-Step Procedure

To apply a filter to an output interface:

1. Apply a filter to an interface.

```
[edit]
user@host# edit interfaces e1-1/0/1 unit 0 family inet filter output
user@host# set voice-term
```

2. Set an IP address.

```
[edit]
user@host# set interfaces e1-1/0/1 unit 0 family inet address 203.0.113.1/24
```

Results

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

```
[edit]
  user@host# show interfaces
  e1-1/0/1 {
  unit 0 {
  family inet {
    filter {
    output voice-term;
    }
    address 203.0.113.1/24;
    }
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Scheduler Map on page 1516
- Verifying the Interfaces on page 1517
- Verifying the Interface Queues on page 1517

Verifying the Scheduler Map

Purpose

Verify that the scheduler map is configured properly.

Action

From operational mode, enter the **show class-of-service scheduler-map corp-map** command.

Copyright © 2016, Juniper Networks, Inc.

Verifying the Interfaces

Purpose Verify that the interfaces are configured properly.

Action From configuration mode, enter the **show interfaces** command.

Verifying the Interface Queues

Purpose Verify that the interface queues are configured properly.

Action From configuration mode, enter the **show interfaces queue** command.

Related Documentation

- Strict-Priority Queue Overview on page 1505
- Understanding Strict-Priority Queues on page 1506
- Example: Configuring Priority Scheduling on page 1506

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 72

Controlling Congestion with Drop Profiles

- RED Drop Profiles Overview on page 1519
- RED Drop Profiles and Congestion Control on page 1520
- Configuring RED Drop Profiles on page 1521
- Example: Configuring RED Drop Profiles on page 1523

RED Drop Profiles Overview

A drop profile is a feature of the random early detection (RED) process that allows packets to be dropped before queues are full. Drop profiles are composed of two main values—the queue fullness and the drop probability. The queue fullness represents percentage of memory used to store packets in relation to the total amount that has been allocated for that queue. The drop probability is a percentage value that correlates to the likelihood that an individual packet is dropped from the network. These two variables are combined in a graph-like format.

A random number between 0 and 100 is calculated for each packet. This random number is plotted against the drop profile having the current queue fullness of that particular queue. When the random number falls above the graph line, the packet is transmitted onto the physical media. When the number falls below the graph line, the packet is dropped from the network.

Randomly dropped packets are counted as RED-dropped, while packets dropped for other reasons (100% probability) are counted as tail-dropped.

You specify drop probabilities in the drop profile section of the class-of-service (CoS) configuration hierarchy and reference them in each scheduler configuration. For each scheduler, you can configure multiple separate drop profiles, one for each combination of loss priority (low, medium-low, medium-high, or high) and IP transport protocol (TCP or non-TCP or any).

You can configure a maximum of 32 different drop profiles.

To configure RED drop profiles, include the following statements at the **[edit class-of-service]** hierarchy level of the configuration:

```
[edit class-of-service]
drop-profiles {
  profile-name {
```

Copyright © 2016, Juniper Networks, Inc.

```
            fill-level percentage drop-probability percentage;
            interpolate {
              drop-probability [ values ];
              fill-level [ values ];
            }
          }
        }
```

Default Drop Profiles

By default, if you configure no drop profiles, RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.

Related Documentation

- Example: Configuring RED Drop Profiles on page 1523

RED Drop Profiles and Congestion Control

If the device must support assured forwarding, you can control congestion by configuring random early detection (RED) drop profiles. RED drop profiles use drop probabilities for different levels of buffer fullness to determine which scheduling queue on the device is likely to drop assured forwarding packets under congested conditions. The device can drop packets when the queue buffer becomes filled to the configured percentage.

Assured forwarding traffic with the PLP (packet loss priority) bit set is more likely to be discarded than traffic without the PLP bit set. This example shows how to configure a drop probability and a queue fill level for both PLP and non-PLP assured forwarding traffic. It is only one example of how to use RED drop profiles.

The example shows how to configure the RED drop profiles listed in Table 143.

Table 143: Sample RED Drop Profiles

Drop Profile	Drop Probability	Queue Fill Level
af-normal—For non-PLP (normal) assured forwarding traffic	Between 0 (never dropped) and 100 percent (always dropped)	Between 95 and 100 percent
af-with-plp—For PLP (aggressive packet dropping) assured forwarding traffic	Between 95 and 100 percent (always dropped)	Between 80 and 95 percent

To configure RED drop profiles for assured forwarding congestion control on the device:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 144.
3. If you are finished configuring the device, commit the configuration.
4. Go on to one of the following tasks:

Copyright © 2016, Juniper Networks, Inc.

- To assign resources, priorities, and profiles to output queues, see "Example: Configuring Class-of-Service Schedulers" on page 1485.
- To apply rules to logical interfaces, see "Example: Configuring Virtual Channels" on page 1533.
- To use adaptive shapers to limit bandwidth for Frame Relay, see "Example: Configuring and Applying an Adaptive Shaper" on page 1528.

Table 144: Configuring RED Drop Profiles for Assured Forwarding Congestion Control

Task	CLI Configuration Editor
Navigate to the **Class of service** level in the configuration hierarchy.	From the **[edit]** hierarchy level, enter **edit class-of-service**
Configure the lower drop probability for normal, non-PLP traffic.	Enter **edit drop-profiles af-normal interpolate** **set drop-probability 0** **set drop-probability 100**
Configure a queue fill level for the lower non-PLP drop probability.	Enter **set fill-level 95** **set fill-level 100**
Configure the higher drop probability for PLP traffic.	From the **[edit class of service]** hierarchy level, enter **edit drop-profiles af-with-PLP interpolate** **set drop-probability 95** **set drop-probability 100**
Configure a queue fill level for the higher PLP drop probability.	Enter **set fill-level 80** **set fill-level 95**

Related Documentation
- Example: Configuring RED Drop Profiles on page 1523

Configuring RED Drop Profiles

Create a segmented configuration and an interpolated configuration that correspond to the graphs in Figure 75. The values defined in the configuration are matched to represent the data points in the graph line. In this example, the drop probability is 25 percent when the queue is 50 percent full. The drop probability increases to 50 percent when the queue is 75 percent full.

Copyright © 2016, Juniper Networks, Inc.

Figure 75: Segmented and Interpolated Drop Profiles

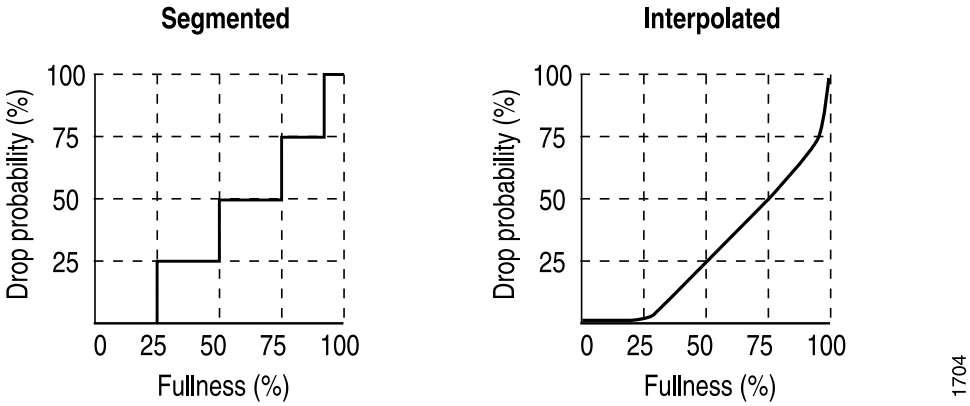

Segmented

```
class-of-service {
  drop-profiles {
    segmented-style-profile {
      fill-level 25 drop-probability 25;
      fill-level 50 drop-probability 50;
      fill-level 75 drop-probability 75;
      fill-level 95 drop-probability 100;
    }
  }
}
```

To create the profile's graph line, the software begins at the bottom-left corner, representing a 0 percent fill level and a 0 percent drop probability. This configuration draws a line directly to the right until it reaches the first defined fill level, 25 percent for this configuration. The software then continues the line vertically until the first drop probability is reached. This process is repeated for all of the defined levels and probabilities until the top-right corner of the graph is reached.

Create a smoother graph line by configuring the profile with the **interpolate** statement. This allows the software to automatically generate 64 data points on the graph beginning at (0, 0) and ending at (100, 100). Along the way, the graph line intersects specific data points, which you define as follows:

Interpolated

```
class-of-service {
  drop-profiles {
    interpolated-style-profile {
      interpolate {
        fill-level [ 50 75 ];
        drop-probability [ 25 50 ];
      }
    }
  }
}
```

Related Documentation

- Example: Configuring RED Drop Profiles on page 1523

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring RED Drop Profiles

This example shows how to configure RED drop profiles.

- Requirements on page 1523
- Overview on page 1523
- Configuration on page 1523
- Verification on page 1525

Requirements

Before you begin, determine which type of profile you want to configure. See *Example: Configuring Segmented and Interpolated Style Profiles*.

Overview

A drop profile is a feature of the RED process that allows packets to be dropped before queues are full. Drop profiles are composed of two main values the queue fullness and the drop probability.

You can control congestion by configuring RED drop profiles, if the device supports assured forwarding. RED drop profiles use drop probabilities for different levels of buffer fullness to determine which scheduling queue on the device is likely to drop assured forwarding packets under congested conditions. The device can drop packets when the queue buffer becomes filled to the configured percentage. Assured forwarding traffic with the PLP bit set is more likely to be discarded than traffic without the PLP bit set.

In this example, you configure a drop probability and a queue fill level for both PLP and non-PLP assured forwarding traffic.

Table 145 shows how to configure the RED drop profiles listed.

Table 145: Sample RED Drop Profiles

Drop Profile	Drop Probability	Queue Fill Level
af-normal—For non-PLP (normal) assured forwarding traffic	Between 0 (never dropped) and 100 percent (always dropped)	Between 95 and 100 percent
af-with-plp—For PLP (aggressive packet dropping) assured forwarding traffic	Between 95 and 100 percent (always dropped)	Between 80 and 95 percent

Configuration

CLI Quick Configuration

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

```
[edit]
set class-of-service drop-profiles af-normal interpolate drop-probability 0
```

Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service drop-profiles af-normal interpolate drop-probability 100
set class-of-service drop-profiles af-normal interpolate fill-level 95
set class-of-service drop-profiles af-normal interpolate fill-level 100
set class-of-service drop-profiles af-with-PLP interpolate drop-probability 95
set class-of-service drop-profiles af-with-PLP interpolate drop-probability 100
set class-of-service drop-profiles af-with-PLP interpolate fill-level 80
set class-of-service drop-profiles af-with-PLP interpolate fill-level 95
```

Step-by-Step Procedure

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

To configure RED drop profiles:

1. Configure the lower drop probability for normal, non-PLP traffic.

```
[edit]
user@host# edit class-of-service
user@host# edit drop-profiles af-normal interpolate
user@host# set drop-probability 0
user@host# set drop-probability 100
```

2. Configure a queue fill level for the lower non-PLP drop probability.

```
[edit class-of-service drop-profiles af-normal interpolate]
user@host# set fill-level 95
user@host# set fill-level 100
```

3. Configure the higher drop probability for PLP traffic.

```
[edit]
user@host# edit class-of-service
user@host# edit drop-profiles af-with-PLP interpolate
user@host# set drop-probability 95
user@host# set drop-probability 100
```

4. Configure a queue fill level for the higher PLP drop probability.

```
[edit class-of-service drop-profiles af-with-PLP interpolate]
user@host# set fill-level 80
user@host# set fill-level 95
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show class-of-service
  drop-profiles {
  af-normal {
    interpolate {
      fill-level [ 95 100 ];
      drop-probability [ 0 100 ];
    }
  }
  af-with-PLP {
    interpolate {
      fill-level [ 80 95 ];
```


Copyright © 2016, Juniper Networks, Inc.

```
            drop-probability [ 95 100 ];
        }
    }
}
```

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

Verification

Verifying RED Drop Profiles Configuration

Purpose Verify that the RED drop profiles are configured properly.

Action From operational mode, enter the **show class-of-service** command.

Related Documentation

- RED Drop Profiles Overview on page 1519
- *Understanding RED Drop Profiles*

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 73

Controlling Congestion with Adaptive Shapers

- Adaptive Shaping Overview on page 1527
- Defining a Custom Frame Relay Loss Priority Map on page 1528
- Example: Configuring and Applying an Adaptive Shaper on page 1528

Adaptive Shaping Overview

You can use adaptive shaping to limit the bandwidth of traffic flowing on a Frame Relay logical interface. If you configure and apply adaptive shaping, the device checks the backward explicit congestion notification (BECN) bit within the last inbound (ingress) packet received on the interface.

NOTE: Adaptive shaping is not available on SRX5600 and SRX5800 devices.

To configure an adaptive shaper, include the **adaptive-shaper** statement at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
adaptive-shaper {
   adaptive-shaper-name {
      trigger type shaping-rate (percent percentage | rate);
   }
}
```

The trigger type can be **BECN** only. If the last ingress packet on the logical interface has its BECN bit set to 1, the output queues on the logical interface are shaped according to the associated shaping rate.

The associated shaping rate can be a percentage of the available interface bandwidth from 0 through 100 percent. Alternatively, you can configure the shaping rate to be an absolute peak rate, in bits per second (bps) from 3200 through 32,000,000,000 bps. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation **K** (1000), **M** (1,000,000), or **G** (1,000,000,000).

Related Documentation

- *Assigning the Default Frame Relay Loss Priority Map to an Interface*

Copyright © 2016, Juniper Networks, Inc.

- Defining a Custom Frame Relay Loss Priority Map on page 1528
- Example: Configuring and Applying an Adaptive Shaper on page 1528

Defining a Custom Frame Relay Loss Priority Map

You can apply a classifier to the same interface on which you configure a Frame Relay loss priority value. The Frame Relay loss priority map is applied first, followed by the classifier. The classifier can change the loss priority to a higher value only (for example, from low to high). If the classifier specifies a loss priority with a lower value than the current loss priority of a particular packet, the classifier does not change the loss priority of that packet.

To define a custom Frame Relay loss priority map, include the following statements at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
loss-priority-maps {
   frame-relay-de map-name {
      loss-priority (low | medium-low | medium-high | high) code-point (0 | 1);
   }
}
```

A custom loss priority map sets the loss priority to low, medium-low, medium-high, or high for each incoming frame with the DE bit containing the specified **0** or **1** CoS value.

The map does not take effect until you apply it to a logical interface. To apply a map to a logical interface, include the **frame-relay-de *map-name*** statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* loss-priority-maps] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number loss-priority-maps]
frame-relay-de map-name;
```

Related Documentation

- *Assigning the Default Frame Relay Loss Priority Map to an Interface*

Example: Configuring and Applying an Adaptive Shaper

This example shows how to configure and apply an adaptive shaper to limit the bandwidth of traffic on a Frame Relay logical interface.

- Requirements on page 1528
- Overview on page 1529
- Configuration on page 1529
- Verification on page 1529

Requirements

Before you begin, review how to create and apply scheduler maps. See "Example: Configuring and Applying Scheduler Maps" on page 1500

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you create adaptive shaper fr-shaper and apply it to T1 interface t1-0/0/2. The adapter shaper limits the transmit bandwidth on the interface to 64 Kbps.

Configuration

Step-by-Step Procedure

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

To configure and apply an adaptive shaper to a logical interface:

1. Specify the name and the maximum transmit rate of the adaptive shaper.

   ```
   [edit]
   user@host# edit class-of-service
   user@host# set adaptive-shapers fr-shaper trigger becn shaping-rate 64k
   ```

2. Apply the adaptive shaper to the logical interface.

   ```
   [edit class-of-service]
   user@host# set interfaces t1-0/0/2 unit 0 adaptive-shaper fr-shaper
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show class-of-service** command.

Related Documentation

- Adaptive Shaping Overview on page 1527
- *Assigning the Default Frame Relay Loss Priority Map to an Interface*
- Defining a Custom Frame Relay Loss Priority Map on page 1528

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 74

Limiting Traffic Using Virtual Channels

- Virtual Channels Overview on page 1531
- Understanding Virtual Channels on page 1532
- Example: Configuring Virtual Channels on page 1533

Virtual Channels Overview

You can configure virtual channels to limit traffic sent from a corporate headquarters to branch its offices. Virtual channels might be required when the headquarters site has an expected aggregate bandwidth higher than that of the individual branch offices. The headquarters router must limit the traffic sent to each branch office router to avoid oversubscribing their links. For instance, if branch 1 has a 1.5 Mbps link and the headquarters router attempts to send 6 Mbps to branch 1, all of the traffic in excess of 1.5 Mbps is dropped in the ISP network.

You configure virtual channels on a logical interface. Each virtual channel has a set of eight queues with a scheduler and an optional shaper. You can use an output firewall filter to direct traffic to a particular virtual channel. For example, a filter can direct all traffic with a destination address for branch office 1 to virtual channel 1, and all traffic with a destination address for branch office 2 to virtual channel 2.

Although a virtual channel group is assigned to a logical interface, a virtual channel is not the same as a logical interface. The only features supported on a virtual channel are queuing, packet scheduling, and accounting. Rewrite rules and routing protocols apply to the entire logical interface.

When you configure virtual channels on an interface, the virtual channel group uses the same scheduler and shaper you configure at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. In this way, virtual channels are an extension of regular scheduling and shaping and not an independent entity.

Related Documentation

- Understanding Virtual Channels on page 1532
- Example: Configuring Virtual Channels on page 1533

Copyright © 2016, Juniper Networks, Inc.

Understanding Virtual Channels

You configure a virtual channel to set up queuing, packet scheduling, and accounting rules to be applied to one or more logical interfaces. You then must apply the virtual channel to a particular logical interface.

You also create a list of virtual channels that you can assign to a virtual channel group. To define a virtual channel group that you can assign to a logical interface, include the **virtual-channel-groups** statement at the [edit class-of-service] hierarchy level.

The *virtual-channel-group-name* can be any name that you want. The *virtual-channel-name* must be one of the names that you define at the **[edit class-of-service virtual-channels]** hierarchy level. You can include multiple virtual channel names in a group.

The scheduler map is required. The *map-name* must be one of the scheduler maps that you configure at the [edit class-of-service scheduler-maps] hierarchy level. For more information, see "Example: Configuring Class-of-Service Schedulers" on page 1485.

The shaping rate is optional. If you configure the shaping rate as a percentage, when the virtual channel is applied to a logical interface, the shaping rate is set to the specified percentage of the interface bandwidth. If you configure a shaper on a virtual channel, the shaper limits the maximum bandwidth transmitted by that virtual channel. Virtual channels without a shaper can use the full logical interface bandwidth. If there are multiple unshaped virtual channels, they share the available logical interface bandwidth equally.

When you apply the virtual channel group to a logical interface, a set of eight queues is created for each of the virtual channels in the group. The **scheduler-map** statement applies a scheduler to these queues. If you include the **shaping-rate** statement, a shaper is applied to the entire virtual channel.

You must configure one of the virtual channels in the group to be the default channel. Therefore, the **default** statement is required in the configuration of one virtual channel per channel group. Any traffic not explicitly directed to a particular channel is transmitted by this default virtual channel.

For the corresponding physical interface, you must also include the **per-unit-scheduler** statement at the [edit interfaces *interface-name*] hierarchy level as follows:

```
[edit interfaces interface-name]
per-unit-scheduler;
```

The **per-unit-scheduler** statement enables one set of output queues for each logical interface configured under the physical interface.

When you apply a virtual channel group to a logical interface, the software creates a set of eight queues for each of the virtual channels in the group.

If you apply a virtual channel group to multiple logical interfaces, the software creates a set of eight queues on each logical interface. The virtual channel names listed in the group are used on all the logical interfaces. We recommend specifying the scheduler and shaping rates in the virtual channel configuration in terms of percentages, rather than

Copyright © 2016, Juniper Networks, Inc.

absolute rates. This allows you to apply the same virtual channel group to logical interfaces that have different bandwidths.

When you apply a virtual channel group to a logical interface, you cannot include the **scheduler-map** and **shaping-rate** statements at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level. In other words, you can configure a scheduler map and a shaping rate on a logical interface, or you can configure virtual channels on the logical interface, but not both.

If you configure multiple logical interfaces on a single physical interface, each logical interface is guaranteed an equal fraction of the physical interface bandwidth as follows:

```
logical-interface-bandwidth =
physical-interface-bandwidth / number-of-logical-interfaces
```

If one or more logical interfaces do not completely use their allocation, the other logical interfaces share the excess bandwidth equally.

If you configure multiple virtual channels on a logical interface, they are each guaranteed an equal fraction of the logical interface bandwidth as follows:

```
virtual-channel-bandwidth =
logical-interface-bandwidth / number-of-virtual-channels
```

If you configure a shaper on a virtual channel, the shaper limits the maximum bandwidth transmitted by that virtual channel. Virtual channels without a shaper can use the full logical interface bandwidth. If there are multiple unshaped virtual channels, they share the available logical interface bandwidth equally.

Related Documentation

- Virtual Channels Overview on page 1531
- Example: Configuring Virtual Channels on page 1533

Example: Configuring Virtual Channels

This example shows how to create virtual channels between a headquarters and its branch office.

- Requirements on page 1533
- Overview on page 1534
- Configuration on page 1534
- Verification on page 1537

Requirements

Before you begin, ensure that your headquarters and branch office have a network connection where the expected aggregate bandwidth is higher for your headquarters than for your branch office. The devices at your headquarters will then be set up to limit the traffic sent to the branch office to avoid oversubscribing the link.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you create the virtual channels as branch1–vc, branch2–vc, branch3–vc, and default-vc. You then define the virtual channel group as wan-vc-group to include the four virtual channels and assign the scheduler map as bestscheduler to each virtual channel. Three of the virtual channels are shaped to 1.5 Mbps. The fourth virtual channel is default-vc, and it is not shaped so it can use the full interface bandwidth.

Then you apply them in the firewall filter as choose-vc to the device's interface t3-1/0/0. The output filter on the interface sends all traffic with a destination address matching 192.168.10.0/24 to branch1-vc, and similar configurations are set for branch2-vc and branch3-vc. Traffic not matching any of the addresses goes to the default, unshaped virtual channel.

Configuration

CLI Quick Configuration

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

```
set class-of-service virtual-channels branch1-vc
set class-of-service virtual-channels branch2-vc
set class-of-service virtual-channels branch3-vc
set class-of-service virtual-channels default-vc
set class-of-service virtual-channel-groups wan-vc-group branch1-vc scheduler-map
  bestscheduler
set class-of-service virtual-channel-groups wan-vc-group branch2-vc scheduler-map
  bestscheduler
set class-of-service virtual-channel-groups wan-vc-group branch3-vc scheduler-map
  bestscheduler
set class-of-service virtual-channel-groups wan-vc-group default-vc scheduler-map
  bestscheduler
set class-of-service virtual-channel-groups wan-vc-group default-vc default
set class-of-service virtual-channel-groups wan-vc-group branch1-vc shaping-rate
  1500000
set class-of-service virtual-channel-groups wan-vc-group branch2-vc shaping-rate
  1500000
set class-of-service virtual-channel-groups wan-vc-group branch3-vc shaping-rate
  1500000
set class-of-service interfaces t3-1/0/0 unit 0 virtual-channel-group wan-vc-group
set firewall family inet filter choose-vc term branch1 from destination-address
  192.168.10.0/24
set firewall family inet filter choose-vc term branch1 then accept
set firewall family inet filter choose-vc term branch1 then virtual-channel branch1-vc
set firewall family inet filter choose-vc term branch1 then virtual-channel branch2-vc
set firewall family inet filter choose-vc term branch1 then virtual-channel branch3-vc
set interfaces t3-1/0/0 unit 0 family inet filter output choose-vc
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

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

To configure virtual channels:

1. Define the virtual channels and the default virtual channel.

```
[edit]
user@host# edit class-of-service
user@host# set virtual-channels branch1-vc
user@host# set virtual-channels branch2-vc
user@host# set virtual-channels branch3-vc
user@host# set virtual-channels default-vc
```

2. Define the virtual channel group and assign each virtual channel a scheduler map.

```
[edit class-of-service]
user@host# set virtual-channel-groups wan-vc-group branch1-vc scheduler-map
  bestscheduler
user@host# set virtual-channel-groups wan-vc-group branch2-vc scheduler-map
  bestscheduler
user@host# set virtual-channel-groups wan-vc-group branch3-vc scheduler-map
  bestscheduler
user@host# set virtual-channel-groups wan-vc-group default-vc scheduler-map
  bestscheduler
user@host# set virtual-channel-groups wan-vc-group default-vc default
```

3. Specify a shaping rate.

```
[edit class-of-service]
user@host# set virtual-channel-groups wan-vc-group branch1-vc shaping-rate 1.5m
user@host# set virtual-channel-groups wan-vc-group branch2-vc shaping-rate
  1.5m
user@host# set virtual-channel-groups wan-vc-group branch3-vc shaping-rate
  1.5m
```

4. Apply the virtual channel group to the logical interface.

```
[edit class-of-service]
user@host# set interfaces t3–1/0/0 unit 0 virtual-channel-group wan-vc-group
```

5. Create the firewall filter to select the traffic.

```
[edit firewall]
user@host# set family inet filter choose-vc term branch1 from destination
  192.168.10.0/24
user@host# set family inet filter choose-vc term branch1 then accept
user@host# set family inet filter choose-vc term branch1 then virtual-channel
  branch1-vc
user@host# set family inet filter choose-vc term branch1 then virtual-channel
  branch2-vc
user@host# set family inet filter choose-vc term branch1 then virtual-channel
  branch3-vc
```

6. Apply the firewall filter to output traffic.

```
[edit interfaces]
user@host# set t3–1/0/0 unit 0 family inet filter output choose-vc
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show class-of-service**, **show firewall**, and **show interfaces t3-1/0/0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show class-of-service
virtual-channels {
   branch1-vc;
   branch2-vc;
   branch3-vc;
   default-vc;
}
virtual-channel-groups {
   wan-vc-group {
      branch1-vc {
      scheduler-map bestscheduler;
      shaping-rate 1500000;
   }
   branch2-vc {
      scheduler-map bestscheduler;
      shaping-rate 1500000;
   }
   branch3-vc {
      scheduler-map bestscheduler;
      shaping-rate 1500000;
   }
   default-vc {
      scheduler-map bestscheduler;
      default;
   }
}
}
interfaces {
   t3-1/0/0 {
      unit 0 {
         virtual-channel-group wan-vc-group;
      }
   }
}
[edit]
user@host# show firewall
family inet {
   filter choose-vc {
      term branch1 {
         from {
            destination-address {
               192.168.10.0/24;
            }
         }
         then {
            virtual-channel branch3-vc;
            accept;
         }
      }
   }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show interfaces t3-1/0/0
unit 0 {
   family inet {
      filter {
         output choose-vc;
      }
   }
}
```

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

Verification

Verifying Virtual Channel Configuration

Purpose Verify that the virtual channels are properly configured.

Action From configuration mode, enter the **show class-of-service**, **show firewall**, and **show interfaces t3-1/0/0** commands.

Related Documentation

- Virtual Channels Overview on page 1531
- Understanding Virtual Channels on page 1532

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 75

Enabling Queuing for Tunnel Interfaces

- CoS Queuing for Tunnels Overview on page 1539
- Understanding the ToS Value of a Tunnel Packet on page 1541
- Example: Configuring CoS Queuing for GRE or IP-IP Tunnels on page 1542
- Copying Outer IP Header DSCP and ECN to Inner IP Header on page 1546

CoS Queuing for Tunnels Overview

On an SRX Series device running Junos OS, a tunnel interface is an internal interface and supports many of the same CoS features as a physical interface. The tunnel interface creates a virtual point-to-point link between two SRX Series devices at remote points over an IP network.

For example, you can configure CoS features for generic routing encapsulation (GRE) and IP-IP tunnel interfaces. Tunneling protocols encapsulate packets inside a transport protocol.

GRE and IP-IP tunnels are used with services like IPsec and NAT to set up point-to-point VPNs. Junos OS allows you to enable CoS queuing, scheduling, and shaping for traffic exiting through these tunnel interfaces. For an example of configuring CoS Queuing for GRE tunnels, see "Example: Configuring CoS Queuing for GRE or IP-IP Tunnels" on page 1542.

This topic includes the following sections:

- Benefits of CoS Queuing for Tunnel Interfaces on page 1539
- How CoS Queuing Works on page 1540
- Limitations on CoS Shapers for Tunnel Interfaces on page 1541

Benefits of CoS Queuing for Tunnel Interfaces

CoS queuing enabled for tunnel interfaces has the following benefits:

- Segregates tunnel traffic.

  Each tunnel can be shaped so that a tunnel with low-priority traffic cannot flood other tunnels that carry high-priority traffic.

  Traffic for one tunnel does not impact traffic on other tunnels.

- Controls tunnel bandwidth.

Copyright © 2016, Juniper Networks, Inc.

Traffic through various tunnels is limited to not exceed a certain bandwidth.

For example, suppose you have three tunnels to three remote sites through a single WAN interface. You can select CoS parameters for each tunnel such that traffic to some sites gets more bandwidth than traffic to other sites.

- Customizes CoS policies.

  You can apply different queuing, scheduling, and shaping policies to different tunnels based on user requirements. Each tunnel can be configured with different scheduler maps, different queue depths, and so on. Customization allows you to configure granular CoS policy providing for better control over tunnel traffic.

- Prioritizes traffic before it enters a tunnel.

  For example, CoS queuing avoids having low-priority packets scheduled ahead of high-priority packets when the interface speed is higher than the tunnel traffic speed. This feature is most useful when combined with IPsec. Typically, IPsec processes packets in a FIFO manner. However, with CoS queuing each tunnel can prioritize high-priority packets over low-priority packets. Also, each tunnel can be shaped so that a tunnel with low-priority traffic does not flood tunnels carrying high-priority traffic.

How CoS Queuing Works

Figure 76 shows CoS-related processing that occurs for traffic entering and exiting a tunnel. For information on flow-based packet processing, see the *Flow-Based and Packet-Based Processing Feature Guide for Security Devices*.

Figure 76: CoS Processing for Tunnel Traffic

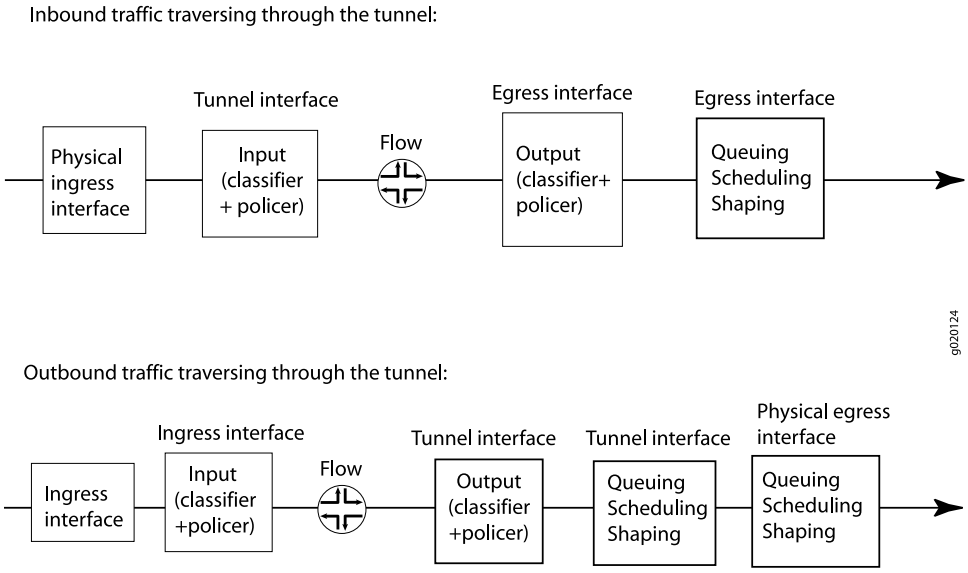

Copyright © 2016, Juniper Networks, Inc.

Limitations on CoS Shapers for Tunnel Interfaces

When defining a CoS shaping rate on a tunnel interface, be aware of the following restrictions:

- The shaping rate on the tunnel interface must be less than that of the physical egress interface.
- The shaping rate only measures the packet size that includes the Layer 3 packet with GRE or IP-IP encapsulation. The Layer 2 encapsulation added by the physical interface is not factored into the shaping rate measurement.
- The CoS behavior works as expected only when the physical interface carries the shaped GRE or IP-IP tunnel traffic alone. If the physical interface carries other traffic, thereby lowering the available bandwidth for tunnel interface traffic, the CoS features do not work as expected.
- You cannot configure a logical interface shaper and a virtual circuit shaper simultaneously on the router. If virtual circuit shaping is desired, do not define a logical interface shaper. Instead, define a shaping rate for all the virtual circuits.

Understanding the ToS Value of a Tunnel Packet

To ensure that the tunneled packet continues to have the same CoS treatment even in the physical interface, you must preserve the type-of-service (ToS) value from the inner IP header to the outer IP header.

For transit traffic, Junos OS preserves the CoS value of the tunnel packet for both GRE and IP-IP tunnel interfaces. The inner IPv4 or IPv6 ToS bits are copied to the outer IPv4 ToS header for both types of tunnel interfaces.

For Routing Engine traffic, however, the router handles GRE tunnel interface traffic differently from IP-IP tunnel interface traffic. Unlike for IP-IP tunnels, the IPv4 ToS bits are not copied to the outer IPv4 header by default. You have a configuration option to copy the ToS value from the packet's inner IPv4 header to the outer IPv4 header.

To copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of a GRE interface.

NOTE: **For IPv6 traffic, the inner ToS value is not copied to the outer IPv4 header for both GRE and IP-IP tunnel interfaces even if the copy-tos-to-outer-ip-header statement is specified.**

This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
```

Copyright © 2016, Juniper Networks, Inc.

```
        family inet;
    }
}
```

Related Documentation

- CoS Queuing for Tunnels Overview on page 1539
- Example: Configuring CoS Queuing for GRE or IP-IP Tunnels on page 1542

Example: Configuring CoS Queuing for GRE or IP-IP Tunnels

This example shows how to configure CoS queuing for GRE or IP-IP tunnels.

- Requirements on page 1542
- Overview on page 1542
- Configuration on page 1543
- Verification on page 1544

Requirements

Before you begin:

- Establish a main office and a branch office connected by a VPN using GRE or IP-IP tunneled interfaces.
- Configure forwarding classes and schedulers. See "Example: Assigning Forwarding Classes to Output Queues" on page 1460 and "Example: Configuring Class-of-Service Schedulers" on page 1485.
- Configure a scheduler map and apply the scheduler map to the tunnel interface. See "Example: Configuring and Applying Scheduler Maps" on page 1500.
- Configure classifiers and apply them to the tunnel interface. See "Example: Configuring Behavior Aggregate Classifiers" on page 1417.
- Create rewrite rules and apply them to the tunnel interface. See "Example: Configuring and Applying Rewrite Rules" on page 1469.

Overview

In this example, you enable tunnel queuing, define the GRE tunnel interface as gr-0/0/0, (Alternatively, you could define the IP-IP tunnel interface as ip-0/0/0.) and set the per unit scheduler. You then set the GRE tunnel's line rate as 100 Mbps by using the shaper definition.

In Figure 77, Router A has a GRE tunnel established with Router B through interface **ge-1/0/0**. Router A also has an IP-IP tunnel established with Router C through interface **ge-1/0/1**. Router A is configured so that tunnel-queuing is enabled. Router B and Router C do not have tunnel-queuing configured.

Copyright © 2016, Juniper Networks, Inc.

Figure 77: Configuring CoS Queuing for GRE Tunnels

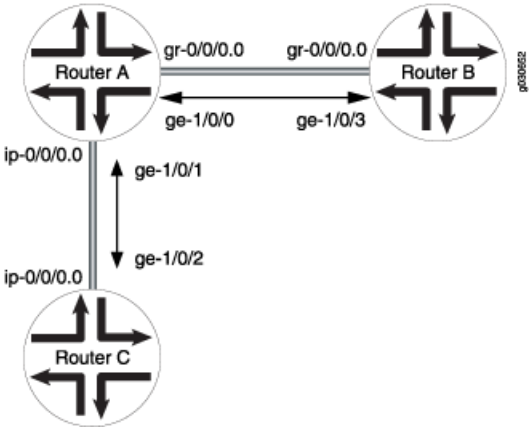

Configuration

CLI Quick Configuration

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

```
set chassis fpc 0 pic 0 tunnel-queuing
set interfaces gr-0/0/0 unit 0
set interfaces gr-0/0/0 per-unit-scheduler
set class-of-services interfaces gr-0/0/0 unit 0 shaping-rate 100m
```

Step-by-Step Procedure

To configure CoS queuing for GRE tunnels:

1. Enable tunnel queuing on the device.

   ```
   [edit]
   user@host# set chassis fpc 0 pic 0 tunnel-queuing
   ```

2. Define the GRE tunnel interface.

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

3. Define the per-unit scheduler for the GRE tunnel interface.

   ```
   [edit]
   user@host# set interfaces gr-0/0/0 per-unit-scheduler
   ```

4. Define the GRE tunnel's line rate by using the shaper definition.

   ```
   [edit]
   user@host# set class-of-services interfaces gr-0/0/0 unit 0 shaping-rate 100m
   ```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service interfaces gr-0/0/0**, **show interfaces gr-0/0/0** , and **show chassis** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show class-of-service interfaces gr-0/0/0
unit 0 {
shaping-rate 100m;
}
[edit]
  user@host# show interfaces gr-0/0/0
  per-unit-scheduler;
unit 0;
[edit]
  user@host# show chassis
  fpc 0 {
  pic 0 {
  tunnel-queuing;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- Verifying a CoS Queuing for GRE Tunnel Configuration on page 1544
- Verifying a CoS Queuing for IP-IP Tunnel Configuration on page 1546

Verifying a CoS Queuing for GRE Tunnel Configuration

Purpose Verify that the device is configured properly for tunnel configuration.

Action From configuration mode, enter the **show interfaces queue gr-0/0/0.0** command.

NOTE: If you enter gr-0/0/0.0 only, queue information for all tunnels is displayed. If you enter gr-0/0/0.0, queue information for the specific tunnel is displayed.

```
user@host> show interfaces queue gr-0/0/0.0
Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 112)
Forwarding classes: 8 supported, 4 in use
Egress queues: 8 supported, 4 in use Burst size: 0
Queue: 0, Forwarding classes: VOICE
  Queued:
    Packets              :              7117734                 7998 pps
    Bytes                :            512476848              4606848 bps
  Transmitted:
    Packets              :              4548146                 3459 pps
    Bytes                :            327466512              1992912 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :              2569421                 4537 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :              2569421                 4537 pps
    RED-dropped bytes    :            184998312              2613640 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
     Low                :                    0                    0 bps
     Medium-low         :                    0                    0 bps
     Medium-high        :                    0                    0 bps
     High               :            184998312              2613640 bps
Queue: 1, Forwarding classes: GOLD
  Queued:
    Packets             :               117600                    0 pps
    Bytes               :              8467200                    0 bps
  Transmitted:
    Packets             :               102435                    0 pps
    Bytes               :              7375320                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                15165                    0 pps
     Low                :                    0                    0 pps
     Medium-low         :                    0                    0 pps
     Medium-high        :                    0                    0 pps
     High               :                15165                    0 pps
    RED-dropped bytes    :              1091880                    0 bps
     Low                :                    0                    0 bps
     Medium-low         :                    0                    0 bps
     Medium-high        :                    0                    0 bps
     High               :              1091880                    0 bps
Queue: 2, Forwarding classes: SILVER
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                :                    0                    0 pps
     Medium-low         :                    0                    0 pps
     Medium-high        :                    0                    0 pps
     High               :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                :                    0                    0 bps
     Medium-low         :                    0                    0 bps
     Medium-high        :                    0                    0 bps
     High               :                    0                    0 bps
Queue: 3, Forwarding classes: BRONZE
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                :                    0                    0 pps
     Medium-low         :                    0                    0 pps
     Medium-high        :                    0                    0 pps
     High               :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                :                    0                    0 bps
     Medium-low         :                    0                    0 bps
     Medium-high        :                    0                    0 bps
     High               :                    0                    0 bps
```

Verifying a CoS Queuing for IP-IP Tunnel Configuration

Purpose Verify that the device is configured properly for tunnel configuration.

Action From configuration mode, enter the **show interfaces queue ip-0/0/0.0** command.

NOTE: If you enter ip-0/0/0.0 only, queue information for all tunnels is displayed. If you enter ip-0/0/0.0, queue information for the specific tunnel is displayed.

Related Documentation

- CoS Queuing for Tunnels Overview on page 1539
- Understanding the ToS Value of a Tunnel Packet on page 1541

Copying Outer IP Header DSCP and ECN to Inner IP Header

This feature enables copying of Differentiated Services Code Point (DSCP) (outer DSCP+ECN) field from the outer IP header encrypted packet to the inner IP header plain text message on the decryption path.

The benefit in enabling this feature is that after IPsec decryption, clear text packets can follow the inner CoS (DSCP+ECN) rules.

This feature supports chassis cluster and also supports IPv6 and IPv4. The following are supported:

- Copying outer IPv4 DSCP and Explicit Congestion Notification (ECN) field to inner IPv4 DSCP and ECN field
- Copying outer IPv6 DSCP and ECN field to inner IPv6 DSCP and ECN field
- Copying outer IPv4 DSCP and ECN field to inner IPv6 DSCP and ECN field
- Copying outer IPv6 DSCP and ECN field to inner IPv4 DSCP and ECN field

By default this feature is disabled. When you enable this feature on a VPN object, the corresponding IPsec security Association (SA) is cleared and reestablished.

- To enable the feature:

  set security ipsec vpn *vpn-name* copy-outer-dscp

- To disable the feature:

  delete security ipsec vpn *vpn-name* copy-outer-dscp

- To verify whether the feature is enabled or not:

  show security ipsec security-associations detail

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *VPN Feature Guide for Security Devices*
- show security ipsec security-associations on page 3681

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 76

Naming Components with Code-Point Aliases

- Code-Point Aliases Overview on page 1549
- Default CoS Values and Aliases on page 1550
- Example: Defining Code-Point Aliases for Bits on page 1553

Code-Point Aliases Overview

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other class-of-service (CoS) components, such as classifiers, drop-profile maps, and rewrite rules.

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure user-defined classifiers in terms of alias names. If the value of an alias changes, it alters the behavior of any classifier that references it.

The following types of code points are supported by Junos operating system (OS):

- DSCP—Defines aliases for DiffServ code point (DSCP) IPv4 values.

  You can refer to these aliases when you configure classes and define classifiers.
- DSCP-IPv6—Defines aliases for DSCP IPv6 values.

  You can refer to these aliases when you configure classes and define classifiers.
- EXP—Defines aliases for MPLS EXP bits.

  You can map MPLS EXP bits to the device forwarding classes.
- inet-precedence—Defines aliases for IPv4 precedence values.

  Precedence values are modified in the IPv4 type-of-service (ToS) field and mapped to values that correspond to levels of service.

Related Documentation

- Default CoS Values and Aliases on page 1550
- Example: Defining Code-Point Aliases for Bits on page 1553

Copyright © 2016, Juniper Networks, Inc.

Default CoS Values and Aliases

Table 146 shows the default mapping between the standard aliases and the bit values.

Copyright © 2016, Juniper Networks, Inc.

Table 146: Standard CoS Aliases and Bit Values

CoS Value Type	Alias	Bit Value
MPLS EXP	be	000
	be1	001
	ef	010
	ef1	011
	af11	100
	af12	101
	nc1/cs6	110
	nc2/cs7	111

Copyright © 2016, Juniper Networks, Inc.

Table 146: Standard CoS Aliases and Bit Values *(continued)*

CoS Value Type	Alias	Bit Value
DSCP and DSCP IPv6	ef	101110
	af11	001010
	af12	001100
	af13	001110
	af21	010010
	af22	010100
	af23	010110
	af31	011010
	af32	011100
	af33	011110
	af41	100010
	af42	100100
	af43	100110
	be	000000
	cs1	001000
	cs2	010000
	cs3	011000
	cs4	100000
	cs5	101000
	nc1/cs6	110000
	nc2/cs7	111000

Copyright © 2016, Juniper Networks, Inc.

Table 146: Standard CoS Aliases and Bit Values *(continued)*

CoS Value Type	Alias	Bit Value
IEEE 802.1	be	000
	be1	001
	ef	010
	ef1	011
	af11	100
	af12	101
	nc1/cs6	110
	nc2/cs7	111
IP precedence	be	000
	be1	001
	ef	010
	ef1	011
	af11	100
	af12	101
	nc1/cs6	110
	nc2/cs7	111

Related Documentation

- Code-Point Aliases Overview on page 1549
- Example: Defining Code-Point Aliases for Bits on page 1553

Example: Defining Code-Point Aliases for Bits

This example shows how to define code-point aliases for bits on a device.

- Requirements on page 1554
- Overview on page 1554
- Configuration on page 1554
- Verification on page 1554

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, determine which default mapping to use. See "Default CoS Values and Aliases" on page 1550.

Overview

In this example, you configure class of service and specify names and values for the CoS code-point aliases that you want to configure. Finally, you specify CoS value using the appropriate formats.

Configuration

Step-by-Step Procedure

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

To define code-point aliases for bits on a device:

1. Configure class of service.

   ```
   [edit]
   user@host# edit class-of-service
   ```

2. Specify CoS values.

   ```
   [edit class-of-service]
   user@host# set code-point-aliases dscp my1 110001
   user@host# set code-point-aliases dscp my2 101110
   user@host# set code-point-aliases dscp be 000001
   user@host# set code-point-aliases dscp cs7 110000
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show class-of-service code-point-aliases dscp** command.

Related Documentation

- Code-Point Aliases Overview on page 1549

Copyright © 2016, Juniper Networks, Inc.

PART 18

Configuring Class of Service Scheduler Hierarchy

- Controlling Traffic by Configuring Scheduler Hierarchy on page 1557

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 77

Controlling Traffic by Configuring Scheduler Hierarchy

- Understanding Hierarchical Schedulers on page 1557
- Understanding Internal Scheduler Nodes on page 1560
- Example: Configuring a Four-Level Scheduler Hierarchy on page 1561
- Example: Controlling Remaining Traffic on page 1573

Understanding Hierarchical Schedulers

Hierarchical schedules consist of nodes and queues. Queues terminate the CLI hierarchy. Nodes can be either root nodes, leaf nodes, or internal (non-leaf) nodes. Internal nodes are nodes that have other nodes as "children" in the hierarchy. For example, if an **interface-set** statement is configured with a logical interface (such as unit 0) and queue, then the **interface-set** is an internal node at level 2 of the hierarchy. However, if there are no traffic control profiles configured on logical interfaces, then the interface set is at level 3 of the hierarchy.

Table 147 shows how the configuration of an interface set or logical interface affects the terminology of hierarchical scheduler nodes.

Table 147: Hierarchical Scheduler Nodes

Root Node (Level 1)	Internal Node (Level 2)	Leaf Node (Level 3)	Queue (Level 4)
Physical interface	Interface set	Logical interfaces	One or more queues
Physical interface	–	Interface set	One or more queues
Physical interface	–	Logical interfaces	One or more queues

When used, the interface set level of the hierarchy falls between the physical interface level (level 1) and the logical interface (level 3). Queues are always level 4 of the hierarchy. The schedulers hold the information about the queues, the last level of the hierarchy. In all cases, the properties for level 4 of the hierarchical schedulers are determined by the scheduler map.

Copyright © 2016, Juniper Networks, Inc.

Hierarchical schedulers add CoS parameters to the new interface set level of the configuration. They use traffic control profiles to set values for parameters such as shaping rate (the peak information rate [PIR]), guaranteed rate (the committed information rate [CIR] on these interfaces), and scheduler maps (the queues and resources assigned to traffic).

The following CoS configuration places the following parameters in traffic control profiles at various levels:

- Traffic control profile at the port level (**tcp-port-level1**):
  - A shaping rate (PIR) of 100 Mbps
  - A delay buffer rate of 100 Mbps
- Traffic control profile at the interface set level (**tcp-interface-level2**):
  - A shaping rate (PIR) of 60 Mbps
  - A guaranteed rate (CIR) of 40 Mbps
- Traffic control profile at the logical interface level (**tcp-unit-level3**):
  - A shaping rate (PIR) of 50 Mbps
  - A guaranteed rate (CIR) of 30 Mbps
  - A scheduler map called smap1 to hold various queue properties (level 4)
  - A delay buffer rate of 40 Mbps

In this case, the traffic control profiles look like this:

```
[edit class-of-service traffic-control-profiles]
tcp-port-level1 { # This is the physical port level
   shaping-rate 100m;
   delay-buffer-rate 100m;
}
tcp-interface-level2 { # This is the interface set level
   shaping-rate 60m;
   guaranteed-rate 40m;
}
tcp-unit-level3 { # This is the logical interface level
   shaping-rate 50m;
   guaranteed-rate 30m;
   scheduler-map smap1;
   delay-buffer-rate 40m;
}
```

Once configured, the traffic control profiles must be applied to the proper places in the CoS interfaces hierarchy.

```
[edit class-of-service interfaces]
interface-set level-2 {
   output-traffic-control-profile tcp-interface-level-2;
}
ge-0/1/0 {
   output-traffic-control-profile tcp-port-level-1;
```


Copyright © 2016, Juniper Networks, Inc.

```
    unit 0 {
      output-traffic-control-profile tcp-unit-level-3;
    }
  }
```

Interface sets can be defined as a list of logical interfaces, for example, unit 100, unit 200, and so on. Service providers can use these statements to group interfaces to apply scheduling parameters such as guaranteed rate and shaping rate to the traffic in the groups. Interface sets are currently only used by CoS, but they are applied at the **[edit interfaces]** hierarchy level so that they might be available to other services.

All traffic heading downstream must be gathered into an interface set with the **interface-set** statement at the [edit class-of-service interfaces] hierarchy level.

NOTE: Ranges are not supported; you must list each logical interface separately.

Although the interface set is applied at the [edit interfaces] hierarchy level, the CoS parameters for the interface set are defined at the [edit class-of-service interfaces] hierarchy level, usually with the **output-traffic-control-profile** ***profile-name*** statement.

You cannot specify an interface set mixing the logical interface, S-VLAN, or VLAN outer tag list forms of the **interface-set** statement. A logical interface can only belong to one interface set. If you try to add the same logical interface to different interface sets, the commit will fail.

This example will generate a commit error:

```
[edit interfaces]
interface-set set-one {
  ge-2/0/0 {
    unit 0;
    unit 2;
  }
}
interface-set set-two {
  ge-2/0/0 {
    unit 1;
    unit 3;
    unit 0; # COMMIT ERROR! Unit 0 already belongs to -set-one.
  }
}
```

Members of an interface set cannot span multiple physical interfaces. Only one physical interface is allowed to appear in an interface set.

This configuration is not supported:

```
[edit interfaces]
interface-set set-group {
  ge-0/0/1 {
    unit 0;
    unit 1;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    ge-0/0/2 { # This type of configuration is NOT supported in the same interface set!
      unit 0;
      unit 1;
    }
  }
```

You can configure many logical interfaces under an interface. However, only a subset of them might have a traffic control profile attached. For example, you can configure three logical interfaces (units) over the same service VLAN, but you can apply a traffic control profile specifying best-effort and voice queues to only one of the logical interface units. Traffic from the two remaining logical interfaces is considered remaining traffic.

The scheduler map configured at individual interfaces (Level 3), interface sets (Level 2), or physical ports (Level 1), defines packet scheduling behavior at different levels. You can group logical interfaces in an interface set and configure the interfaces with scheduler maps. Any egress packet arriving at the physical or logical interfaces will be handled by the interface specific scheduler. If the scheduler map is not configured at the interface level, the packet will be handled by the scheduler configured at the interface set level or the port level.

Related Documentation

- Example: Configuring a Four-Level Scheduler Hierarchy on page 1561
- Example: Controlling Remaining Traffic on page 1573
- Understanding Internal Scheduler Nodes on page 1560

Understanding Internal Scheduler Nodes

A node in the hierarchy is considered internal if either of the following conditions apply:

- One of its children nodes has a traffic control profile configured and applied.
- You configure the **internal-node** statement.

There are more resources available at the logical interface (unit) level than at the interface set level. It might be desirable to configure all resources at a single level, rather than spread over several levels. The **internal-node** statement provides this flexibility. This can be a helpful configuration device when interface-set queuing without logical interfaces is used exclusively on the interface.

You can use the **internal-node** statement to raise the interface set without children to the same level as the other configured interface sets with children, allowing them to compete for the same set of resources.

Using the **internal-node** statement allows statements to all be scheduled at the same level with or without children.

The following example makes the interface sets **if-set-1** and **if-set-2** internal:

```
[edit class-of-service interfaces ]
interface-set {
  if-set-1 {
    internal-node;
```


Copyright © 2016, Juniper Networks, Inc.

```
        output-traffic-control-profile tcp-200m-no-smap;
    }
    if-set-2 {
        internal-node;
        output-traffic-control-profile tcp-100m-no-smap;
    }
}
```

If an interface set has logical interfaces configured with a traffic control profile, then the use of the **internal-node** statement has no effect.

Internal nodes can specify a **traffic-control-profile-remaining** statement.

Related Documentation

- Understanding Hierarchical Schedulers on page 1557
- Example: Configuring a Four-Level Scheduler Hierarchy on page 1561
- Example: Controlling Remaining Traffic on page 1573

Example: Configuring a Four-Level Scheduler Hierarchy

This example shows how to configure a 4-level hierarchy of schedulers.

- Requirements on page 1561
- Overview on page 1561
- Configuration on page 1562
- Verification on page 1572

Requirements

Before you begin:

- Review how to configure schedulers. See "Example: Configuring Class-of-Service Schedulers" on page 1485.
- Review RED drop profiles. See *Understanding RED Drop Profiles*.
- Review how to configure and apply scheduler maps. See "Example: Configuring and Applying Scheduler Maps" on page 1500.

Overview

The configuration parameters for this example are shown in Figure 78. The queues are shown at the top of the figure with the other three levels of the hierarchy below.

Copyright © 2016, Juniper Networks, Inc.

Figure 78: Building a Scheduler Hierarchy

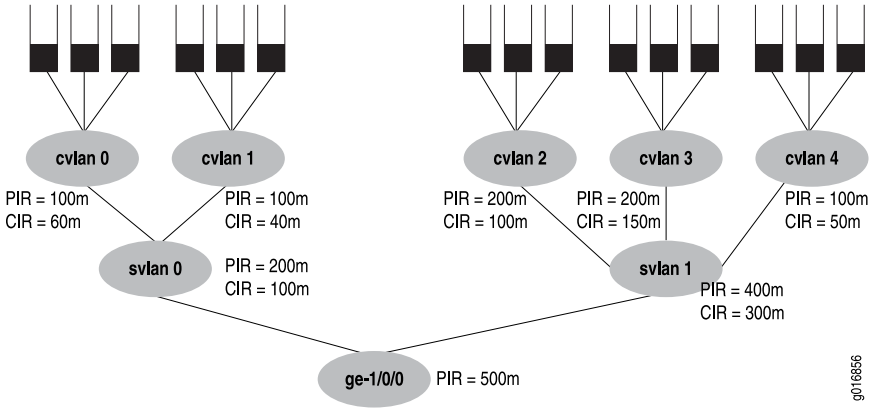

Figure 78's PIR values will be configured as the shaping rates, and the CIRs will be configured as the guaranteed rate on the Ethernet interface **ge-1/0/0**. The PIR can be oversubscribed (that is, the sum of the children PIRs can exceed the parent's, as in **svlan 1**, where 200 + 200 + 100 exceeds the parent rate of 400). However, the sum of the children node level's CIRs must never exceed the parent node's CIR, as shown in all the service VLANs (otherwise, the guaranteed rate could never be provided in all cases).

NOTE: **Although a shaping rate can be applied directly to the physical interface, hierarchical schedulers must use a traffic control profile to hold the shaping rate parameter.**

The keyword to configure hierarchical schedulers is at the physical interface level, as are VLAN tagging and the VLAN IDs. In this example, the interface sets are defined by logical interfaces (units) and not outer VLAN tags. All VLAN tags in this example are customer VLAN tags.

The traffic control profiles in this example are for both the service VLAN level (logical interfaces) and the customer VLAN (VLAN tag) level.

This example shows all details of the CoS configuration for the **ge-1/0/0** interface in Figure 78.

Configuration

This section contains the following topics:

- Configuring the Logical Interfaces on page 1563
- Configuring the Interface Sets on page 1564
- Applying an Interface Set on page 1565
- Configuring the Traffic Control Profiles on page 1565
- Configuring the Schedulers on page 1567
- Configuring the Drop Profiles on page 1568

Copyright © 2016, Juniper Networks, Inc.

- Configuring the Scheduler Maps on page 1569
- Applying Traffic Control Profiles on page 1571

Configuring the Logical Interfaces

CLI Quick Configuration

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

```
edit interface ge-1/0/0
set hierarchical-scheduler vlan-tagging unit 1 vlan-id 101
set hierarchical-scheduler vlan-tagging unit 1 vlan-id 101
set hierarchical-scheduler vlan-tagging unit 2 vlan-id 102
set hierarchical-scheduler vlan-tagging unit 3 vlan-id 103
set hierarchical-scheduler vlan-tagging unit 4 vlan-id 104
```

Step-by-Step Procedure

To configure the logical interfaces:

1. Create the logical interface.

   ```
   [edit]
   user@host# edit interface ge-1/0/0
   ```

2. Create the interface sets by defining the VLAN tagging and the VLAN IDs for each level.

   ```
   [edit interface ge-1/0/0]
   user@host# set hierarchical-scheduler vlan-tagging unit 0 vlan-id 100
   user@host# set hierarchical-scheduler vlan-tagging unit 1 vlan-id 101
   user@host# set hierarchical-scheduler vlan-tagging unit 2 vlan-id 102
   user@host# set hierarchical-scheduler vlan-tagging unit 3 vlan-id 103
   user@host# set hierarchical-scheduler vlan-tagging unit 4 vlan-id 104
   ```

Results

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

```
[edit]
user@host# show interface ge-1/0/0
hierarchical-scheduler;
vlan-tagging;
unit 0 {
   vlan-id 100;
}
unit 1 {
   vlan-id 101;
}
unit 2 {
   vlan-id 102;
}
unit 3 {
   vlan-id 103;
}
unit 4 {
```

Copyright © 2016, Juniper Networks, Inc.

```
    vlan-id 104;
}
```

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

Configuring the Interface Sets

CLI Quick Configuration

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

```
set class-of-service interfaces interface-set svlan-0 interface ge-1/0/0 unit 0
set class-of-service interfaces interface-set svlan-0 interface ge-1/0/0 unit 1
set class-of-service interfaces interface-set svlan-1 interface ge-1/0/0 unit 2
set class-of-service interfaces interface-set svlan-1 interface ge-1/0/0 unit 3
set class-of-service interfaces interface-set svlan-1 interface ge-1/0/0 unit 4
```

Step-by-Step Procedure

To configure the interface sets:

1. Create the first logical interface and its CoS parameters.

```
[edit class-of-service interfaces]
user@host# set interface-set svlan-0 interface ge-1/0/0 unit 0
user@host# set interface-set svlan-0 interface ge-1/0/0 unit 1
```

2. Create the second logical interface and its CoS parameters.

```
[edit class-of-service interfaces]
user@host# set interface-set svlan-1 interface ge-1/0/0 unit 2
user@host# set interface-set svlan-1 interface ge-1/0/0 unit 3
user@host# set interface-set svlan-1 interface ge-1/0/0 unit 4
```

Results

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

```
[edit]
user@host# show interfaces
interface-set svlan-0 {
    interface ge-1/0/0 {
        unit 0;
        unit 1;
    }
}
interface-set svlan-1 {
    interface ge-1/0/0 {
        unit 2;
        unit 3;
        unit 4;
    }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Applying an Interface Set

CLI Quick Configuration

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

```
set class-of-service interfaces interface-set set-ge-0 output-traffic-control-profile tcp-set1
```

Step-by-Step Procedure

To apply an interface set:

1. Create the Ethernet interface set.

   ```
   [edit class-of-service interfaces]
   user@host# set interface-set set-ge-0
   ```

2. Apply a traffic control parameter to the Ethernet interface set.

   ```
   [edit class-of-service interfaces interface-set set-ge-0]
   user@host# set output-traffic-control-profile tcp-set1
   ```

Results

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

```
[edit]
user@host# show interfaces
interface-set set-ge-0 {
  output-traffic-control-profile tcp-set1;
}
```

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

Configuring the Traffic Control Profiles

CLI Quick Configuration

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

```
set class-of-service traffic-control-profiles tcp-500m-shaping-rate shaping-rate 500m
set class-of-service traffic-control-profiles tcp-svlan0 shaping-rate 200m guaranteed-rate
  100m delay-buffer-rate 300m
set class-of-service traffic-control-profiles tcp-svlan1 shaping-rate 400m guaranteed-rate
  30100m delay-buffer-rate 100m
set class-of-service traffic-control-profiles tcp-cvlan0 shaping-rate 100m guaranteed-rate
  60m scheduler-map tcp-map-cvlan0
set class-of-service traffic-control-profiles tcp-cvlan1 shaping-rate 100m guaranteed-rate
  40m scheduler-map tcp-map-cvlan1
set class-of-service traffic-control-profiles tcp-cvlan2 shaping-rate 200m guaranteed-rate
  100m scheduler-map tcp-map-cvlanx
set class-of-service traffic-control-profiles tcp-cvlan3 shaping-rate 200m guaranteed-rate
  150m scheduler-map tcp-map-cvlanx
```

Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service traffic-control-profiles tcp-cvlan4 shaping-rate 100m guaranteed-rate
  50m scheduler-map tcp-map-cvlanx
```

Step-by-Step Procedure

To configure the traffic control profiles:

1. Create the traffic profile parameters.

```
[edit class-of-service traffic-control-profiles]
user@host# tcp-500m-shaping-rate shaping-rate 500m
```

2. Create the traffic control profiles and parameters for the S-VLAN (logical interfaces) level.

```
[edit class-of-service traffic-control-profiles]
user@host# set tcp-svlan0 shaping-rate 200m guaranteed-rate 100m
  delay-buffer-rate 300m
user@host# set tcp-svlan1 shaping-rate 400m guaranteed-rate 30100m
  delay-buffer-rate 100m
```

3. Create the traffic control profiles and parameters for the C-VLAN (VLAN tags) level.

```
[edit class-of-service traffic-control-profiles]
user@host# set tcp-cvlan0 shaping-rate 100m guaranteed-rate 60m scheduler-map
  tcp-map-cvlan0
user@host# set tcp-cvlan1 shaping-rate 100m guaranteed-rate 40m scheduler-map
  tcp-map-cvlan1
user@host# set tcp-cvlan2 shaping-rate 200m guaranteed-rate 100m
  scheduler-map tcp-map-cvlanx
user@host# set tcp-cvlan3 shaping-rate 200m guaranteed-rate 150m scheduler-map
  tcp-map-cvlanx
user@host# set tcp-cvlan4 shaping-rate 100m guaranteed-rate 50m scheduler-map
  tcp-map-cvlanx
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service traffic-control-profiles** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service traffic-control-profiles
tcp-500m-shaping-rate {
  shaping-rate 500m;
}
tcp-svlan0 {
  shaping-rate 200m;
  guaranteed-rate 100m;
  delay-buffer-rate 300m; # This parameter is not shown in the figure
}
tcp-svlan1 {
  shaping-rate 400m;
  guaranteed-rate 300m;
  delay-buffer-rate 100m; # This parameter is not shown in the figure
}
tcp-cvlan0 {
  shaping-rate 100m;
  guaranteed-rate 60m;
  scheduler-map tcp-map-cvlan0; # This example applies scheduler maps to customer
    VLANs
```


Copyright © 2016, Juniper Networks, Inc.

```
}
tcp-cvlan1 {
  shaping-rate 100m;
  guaranteed-rate 40m;
  scheduler-map tcp-map-cvlan1; # This example applies scheduler maps to customer
    VLANs
}
tcp-cvlan2 {
  shaping-rate 200m;
  guaranteed-rate 100m;
  scheduler-map tcp-map-cvlanx; # This example applies scheduler maps to customer
    VLANs
}
tcp-cvlan3 {
  shaping-rate 200m;
  guaranteed-rate 150m;
  scheduler-map tcp-map-cvlanx; # This example applies scheduler maps to customer
    VLANs
}
tcp-cvlan4 {
  shaping-rate 100m;
  guaranteed-rate 50m;
  scheduler-map tcp-map-cvlanx; # This example applies scheduler maps to customer
    VLANs
}
```

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

Configuring the Schedulers

CLI Quick Configuration

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

```
set class-of-service schedulers sched-cvlan0-qx priority low transmit-rate 20m buffer-size
  temporal 100ms drop-profile-map loss-priority low dp-low drop-profile-map loss-priority
  high dp-high
set class-of-service schedulers sched-cvlan1-q0 priority high transmit-rate 20m buffer-size
  percent 40 drop-profile-map loss-priority low dp-low drop-profile-map loss-priority
  high dp-high
set class-of-service schedulers sched-cvlanx-qx transmit-rate percent 30 buffer-size
  percent 30 drop-profile-map loss-priority low dp-low drop-profile-map loss-priority
  high dp-high
set class-of-service schedulers sched-cvlan1-qx transmit-rate 10m buffer-size temporal
  100ms drop-profile-map loss-priority low dp-low drop-profile-map loss-priority high
  dp-high
```

Step-by-Step Procedure

To configure the schedulers:

1. Create the schedulers and their parameters.

   ```
   [edit class-of-service schedulers]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set sched-cvlan0-qx priority low transmit-rate 20m buffer-size temporal
  100ms drop-profile-map loss-priority low dp-low drop-profile-map loss-priority
  high dp-high
user@host# set sched-cvlan1-q0 priority high transmit-rate 20m buffer-size percent
  40 drop-profile-map loss-priority low dp-low drop-profile-map loss-priority high
  dp-high
user@host# set sched-cvlanx-qx transmit-rate percent 30 buffer-size percent 30
  drop-profile-map loss-priority low dp-low drop-profile-map loss-priority high
  dp-high
user@host# set sched-cvlan1-qx transmit-rate 10m buffer-size temporal 100ms
  drop-profile-map loss-priority low dp-low drop-profile-map loss-priority high
  dp-high
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service schedulers** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service schedulers
sched-cvlan0-qx {
   priority low;
   transmit-rate 20m;
   buffer-size temporal 100ms;
   drop-profile-map loss-priority low dp-low;
   drop-profile-map loss-priority high dp-high;
}
sched-cvlan1-q0 {
   priority high;
   transmit-rate 20m;
   buffer-size percent 40;
   drop-profile-map loss-priority low dp-low;
   drop-profile-map loss-priority high dp-high;
}
sched-cvlanx-qx {
   transmit-rate percent 30;
   buffer-size percent 30;
   drop-profile-map loss-priority low dp-low;
   drop-profile-map loss-priority high dp-high;
}
sched-cvlan1-qx {
   transmit-rate 10m;
   buffer-size temporal 100ms;
   drop-profile-map loss-priority low dp-low;
   drop-profile-map loss-priority high dp-high;
}
```

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

Configuring the Drop Profiles

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

Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service drop-profiles dp-low interpolate fill-level 80 drop-probability 80
set class-of-service drop-profiles dp-low interpolate fill-level 100 drop-probability 100
set class-of-service drop-profiles dp-high interpolate fill-level 60 drop-probability 80
set class-of-service drop-profiles dp-high interpolate fill-level 80 drop-probability 100
```

Step-by-Step Procedure

To configure the drop profiles:

1. Create the low drop profile.

```
[edit class-of-service drop-profiles]
user@host# set dp-low interpolate fill-level 80 drop-probability 80
user@host# set dp-low interpolate fill-level 100 drop-probability 100
```

2. Create the high drop profile.

```
[edit class-of-service drop-profiles]
user@host# set dp-high interpolate fill-level 60 drop-probability 80
user@host# set dp-high interpolate fill-level 80 drop-probability 100
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service drop-profiles** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service drop-profiles
dp-low {
   interpolate fill-level 80 drop-probability 80;
   interpolate fill-level 100 drop-probability 100;
}
dp-high {
   interpolate fill-level 60 drop-probability 80;
   interpolate fill-level 80 drop-probability 100;
}
```

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

Configuring the Scheduler Maps

CLI Quick Configuration

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

```
set class-of-service scheduler-maps tcp-map-cvlan0 forwarding-class voice scheduler
   sched-cvlan0-qx
set class-of-service scheduler-maps tcp-map-cvlan0 forwarding-class video scheduler
   sched-cvlan0-qx
set class-of-service scheduler-maps tcp-map-cvlan0 forwarding-class data scheduler
   sched-cvlan0-qx
set class-of-service scheduler-maps tcp-map-cvlan1 forwarding-class voice scheduler
   sched-cvlan1-q0
set class-of-service scheduler-maps tcp-map-cvlan1 forwarding-class video scheduler
   sched-cvlan1-qx
set class-of-service scheduler-maps tcp-map-cvlan1 forwarding-class data scheduler
   sched-cvlan1-qx
```

Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service scheduler-maps tcp-map-cvlanx forwarding-class voice scheduler
  sched-cvlanx-qx
set class-of-service scheduler-maps tcp-map-cvlanx forwarding-class video scheduler
  sched-cvlanx-qx
set class-of-service scheduler-maps tcp-map-cvlanx forwarding-class data scheduler
  sched-cvlanx-qx
```

Step-by-Step Procedure

To configure three scheduler maps:

1. Create the first scheduler map.

```
[edit class-of-service scheduler-maps]
user@host# set tcp-map-cvlan0 forwarding-class voice scheduler sched-cvlan0-qx
user@host# set tcp-map-cvlan0 forwarding-class video scheduler sched-cvlan0-qx
user@host# set tcp-map-cvlan0 forwarding-class data scheduler sched-cvlan0-qx
```

2. Create the second scheduler map.

```
[edit class-of-service scheduler-maps]
user@host# set tcp-map-cvlan1 forwarding-class voice scheduler sched-cvlan1-q0
user@host# set tcp-map-cvlan1 forwarding-class video scheduler sched-cvlan1-qx
user@host# set tcp-map-cvlan1 forwarding-class data scheduler sched-cvlan1-qx
```

3. Create the third scheduler map.

```
[edit class-of-service scheduler-maps]
user@host# set tcp-map-cvlanx forwarding-class voice scheduler sched-cvlanx-qx
user@host# set tcp-map-cvlanx forwarding-class video scheduler sched-cvlanx-qx
user@host# set tcp-map-cvlanx forwarding-class data scheduler sched-cvlanx-qx
```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service scheduler-maps** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service scheduler-maps
tcp-map-cvlan0 {
   forwarding-class voice scheduler sched-cvlan0-qx;
   forwarding-class video scheduler sched-cvlan0-qx;
   forwarding-class data scheduler sched-cvlan0-qx;
}
tcp-map-cvlan1 {
   forwarding-class voice scheduler sched-cvlan1-q0;
   forwarding-class video scheduler sched-cvlan1-qx;
   forwarding-class data scheduler sched-cvlan1-qx;
}
tcp-map-cvlanx {
   forwarding-class voice scheduler sched-cvlanx-qx;
   forwarding-class video scheduler sched-cvlanx-qx;
   forwarding-class data scheduler sched-cvlanx-qx;
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Applying Traffic Control Profiles

CLI Quick Configuration

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

```
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate unit 0 output-control-traffic-control-profile tcp-cvlan0
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate unit 1 output-control-traffic-control-profile tcp-cvlan1
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate unit 2 output-control-traffic-control-profile tcp-cvlan2
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate unit 3 output-control-traffic-control-profile tcp-cvlan3
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate unit 4 output-control-traffic-control-profile tcp-cvlan4
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate interface-set-svlan0 output-control-traffic-control-profile
  tcp-svlan0
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile
  tcp-500m-shaping-rate interface-set-svlan1 output-control-traffic-control-profile
  tcp-svlan1
```

Step-by-Step Procedure

To apply traffic control profiles:

1. Set the interface.

   ```
   [edit class-of-service]
   user@host# set interfaces ge-1/0/0
   ```

2. Set the traffic control profiles for the C-VLANs.

   ```
   [edit class-of-service interfaces ge-1/0/0 ]
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate unit 0
     output-control-traffic-control-profile tcp-cvlan0
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate unit 1
     output-control-traffic-control-profile tcp-cvlan1
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate unit 2
     output-control-traffic-control-profile tcp-cvlan2
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate unit 3
     output-control-traffic-control-profile tcp-cvlan3
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate unit 4
     output-control-traffic-control-profile tcp-cvlan4
   ```

3. Set the traffic control profiles for the S-VLANs.

   ```
   [edit class-of-service interfaces ge-1/0/0 ]
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate
     interface-set-svlan0 output-control-traffic-control-profile tcp-svlan0
   user@host# set output-traffic-control-profile tcp-500m-shaping-rate
     interface-set-svlan1 output-control-traffic-control-profile tcp-svlan1
   ```

Results

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

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show class-of-service interfaces
ge-1/0/0 {
  output-traffic-control-profile tcp-500m-shaping-rate;
  unit 0 {
    output-traffic-control-profile tcp-cvlan0;
  }
  unit 1 {
    output-traffic-control-profile tcp-cvlan1;
  }
  unit 2 {
    output-traffic-control-profile tcp-cvlan2;
  }
  unit 3 {
    output-traffic-control-profile tcp-cvlan3;
  }
  unit 4 {
    output-traffic-control-profile tcp-cvlan4;
  }
}
interface-set svlan0 {
  output-traffic-control-profile tcp-svlan0;
}
interface-set svlan1 {
  output-traffic-control-profile tcp-svlan1;
}
```

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

Verification

Verifying Scheduler Hierarchy Configuration

Purpose Verify that the scheduler hierarchy is configured properly.

Action From operational mode, enter the following commands:

- **show interface ge-1/0/0**
- **show class-of-service interfaces**
- **show class-of-service traffic-control-profiles**
- **show class-of-service schedulers**
- **show class-of-service drop-profiles**
- **show class-of-service scheduler-maps**

Related Documentation

- Understanding Hierarchical Schedulers on page 1557
- Example: Controlling Remaining Traffic on page 1573
- Understanding Internal Scheduler Nodes on page 1560

Copyright © 2016, Juniper Networks, Inc.

Example: Controlling Remaining Traffic

This example shows how to control remaining traffic from the remaining logical interfaces.

- Requirements on page 1573
- Overview on page 1573
- Configuration on page 1575
- Verification on page 1578

Requirements

Before you begin:

- Review how to configure schedulers. See "Example: Configuring Class-of-Service Schedulers" on page 1485.
- Review how to configure and apply scheduler maps. See "Example: Configuring and Applying Scheduler Maps" on page 1500.

Overview

To configure transmit rate guarantees for the remaining traffic, you configure the **output-traffic-control-profile-remaining** statement specifying a guaranteed rate for the remaining traffic. Without this statement, the remaining traffic gets a default, minimal bandwidth. Similarly, you can specify the **shaping-rate** and **delay-buffer-rate** statements in the traffic control profile referenced with the **output-traffic-control-profile-remaining** statement to shape and provide buffering for remaining traffic.

In the interface shown in Figure 79, customer VLANs 3 and 4 have no explicit traffic control profile. However, the service provider might want to establish a shaping and guaranteed transmit rate for aggregate traffic heading for those C-VLANs. The solution is to configure and apply a traffic control profile for all remaining traffic on the interface.

Figure 79: Example 1 Handling Remaining Traffic with no Explicit Traffic Control Profile

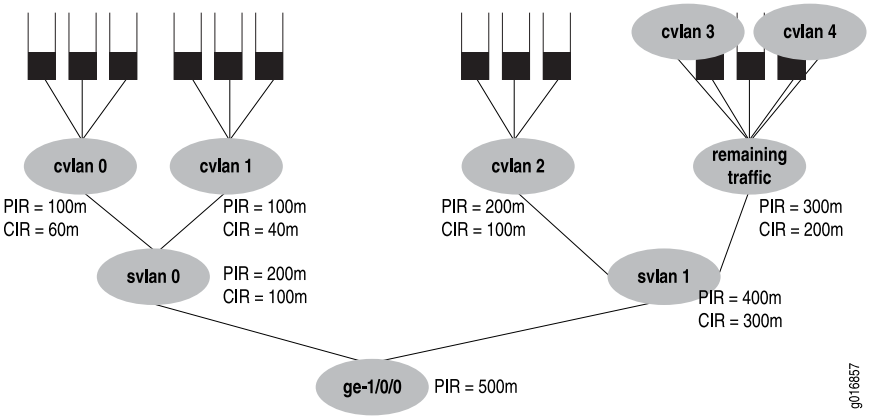

Example 1 considers the case where C-VLANs 3 and 4 have no explicit traffic control profile, yet need to establish a shaping and guaranteed transmit rate for traffic heading for those C-VLANs. The solution is to add a traffic control profile to the **svlan1** interface set. This example builds on the example used in "Example: Configuring a Four-Level Scheduler Hierarchy" on page 1561 and does not repeat all configuration details, only those at the S-VLAN level.

Next, consider Example 2 shown in Figure 80.

Figure 80: Example 2 Handling Remaining Traffic with an Interface Set

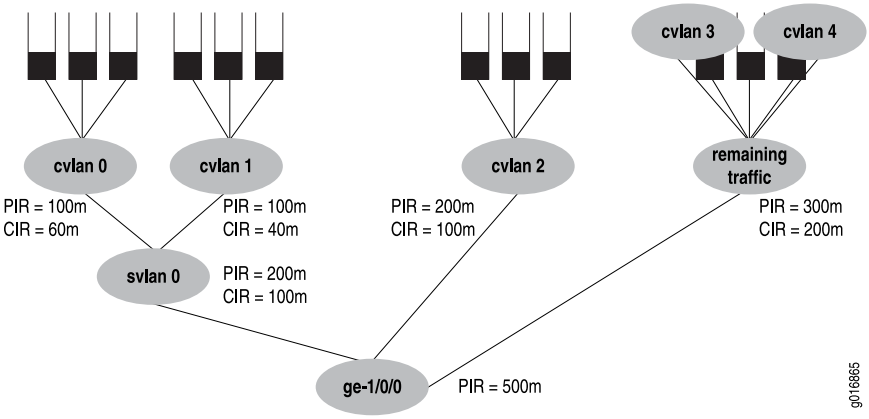

In Example 2, **ge-1/0/0** has five logical interfaces (C-VLAN 0, 1, 2, 3 and 4), and S-VLAN 0, which are covered by the interface set:

- Scheduling for the interface set **svlan0** is specified by referencing an **output-traffic-control-profile** statement, which specifies the **guaranteed-rate**, **shaping-rate**, and **delay-buffer-rate** statement values for the interface set. In this example, the output traffic control profile called **tcp-svlan0** guarantees 100 Mbps and shapes the interface set **svlan0** to 200 Mbps.
- Scheduling and queuing for remaining traffic of **svlan0** is specified by referencing an **output-traffic-control-profile-remaining** statement, which references a **scheduler-map** statement that establishes queues for the remaining traffic. The specified traffic control profile can also configure guaranteed, shaping, and delay-buffer rates for the remaining traffic. In Example 2, **output-traffic-control-profile-remaining tcp-svlan0-rem** references **scheduler-map smap-svlan0-rem**, which calls for a best-effort queue for remaining traffic (that is, traffic on unit 3 and unit 4, which is not classified by the **svlan0** interface set). The example also specifies a **guaranteed-rate** of 200 Mbps and a **shaping-rate** of 300 Mbps for all remaining traffic.
- Scheduling and queuing for logical interface **ge-1/0/0 unit 1** is configured "traditionally" and uses an **output-traffic-control-profile** specified for that unit. In this example, **output-traffic-control-profile tcp-ifl1** specifies scheduling and queuing for **ge-1/0/0 unit 1**.

Copyright © 2016, Juniper Networks, Inc.

Configuration

This section contains the following topics:

- Controlling Remaining Traffic With No Explicit Traffic Control Profile on page 1575
- Controlling Remaining Traffic With An Interface Set on page 1576

Controlling Remaining Traffic With No Explicit Traffic Control Profile

CLI Quick Configuration

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

```
set class-of-service interfaces interface-set svlan0 output-traffic-control-profile
  tcp-svlan0
set class-of-service interfaces interface-set svlan1 output-traffic-control-profile tcp-svlan1
set class-of-service interfaces interface-set svlan1 output-traffic-control-profile-remaining
  tcp-svlan1-remaining
set class-of-service traffic-control-profiles tcp-svlan1 shaping-rate 400m guaranteed-rate
  300m
set class-of-service traffic-control-profiles tcp-svlan1-remaining shaping-rate 300m
  guaranteed-rate 200m scheduler-map smap-remainder
```

Step-by-Step Procedure

To control remaining traffic with no explicit traffic control profile:

1. Set the logical interfaces for the S-VLANs.

   ```
   [edit class-of-service interfaces]
   user@host# set interface-set svlan0 output-traffic-control-profile tcp-svlan0
   user@host# set interface-set svlan1 output-traffic-control-profile tcp-svlan1
   user@host# set interface-set svlan1 output-traffic-control-profile-remaining
     tcp-svlan1-remaining
   ```

2. Set the shaping and guaranteed transmit rates for traffic heading for those C-VLANs.

   ```
   [edit class-of-service traffic-control-profiles]
   user@host# set tcp-svlan1 shaping-rate 400m guaranteed-rate 300m
   user@host# set tcp-svlan1-remaining shaping-rate 300m guaranteed-rate 200m
     scheduler-map smap-remainder
   ```

Results

From configuration mode, confirm your configuration by entering the **show class-of-service interfaces** and **show class-of-service traffic-control-profiles** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service interfaces
interface-set svlan0 {
  output-traffic-control-profile tcp-svlan0;
}
interface-set svlan1 {
  output-traffic-control-profile tcp-svlan1;
  output-traffic-control-profile-remaining tcp-svlan1-remaining; # For all remaining traffic
```

Copyright © 2016, Juniper Networks, Inc.

```
}

[edit]
user@host# show class-of-service traffic-control-profiles
tcp-svlan1 {
   shaping-rate 400m;
   guaranteed-rate 300m;
}
tcp-svlan1-remaining {
   shaping-rate 300m;
   guaranteed-rate 200m;
   scheduler-map smap-remainder; # this smap is not shown in detail
}
```

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

Controlling Remaining Traffic With An Interface Set

CLI Quick Configuration

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

```
set class-of-service interfaces interface-set svlan0 output-traffic-control-profile
   tcp-svlan0
set class-of-service interfaces ge-1/0/0 output-traffic-control-profile-remaining
   tcp-svlan0-rem unit 1output-traffic-control-profile tcp-ifl1
set class-of-service traffic-control-profiles tcp-svlan0 shaping-rate 200m guaranteed-rate
   100m
set class-of-service traffic-control-profiles tcp-svlan0-rem shaping-rate 300m
   guaranteed-rate 200m scheduler-map smap-svlan0-rem
set class-of-service traffic-control-profiles tcp-ifl1 scheduler-map smap-ifl1
set class-of-service scheduler-maps smap-svlan0-rem forwarding-class best-effort
   scheduler-sched-foo
set class-of-service scheduler-maps smap-ifl1 forwarding-class best-effort
   scheduler-sched-bar
set class-of-service scheduler-maps smap-ifl1 forwarding-class assured-forwarding
   scheduler-sched-bar
```

Step-by-Step Procedure

To control remaining traffic with an interface set:

1. Set the interface set for the S-VLAN.

   ```
   [edit class-of-service interfaces]
   user@host# set interface-set svlan0 output-traffic-control-profile tcp-svlan0
   user@host# set ge-1/0/0 output-traffic-control-profile-remaining tcp-svlan0-rem
      unit 1output-traffic-control-profile tcp-ifl1
   ```

2. Set the traffic control profiles.

   ```
   [edit class-of-service traffic-control-profiles]
   user@host# set tcp-svlan0 shaping-rate 200m guaranteed-rate 100m
   user@host# set tcp-svlan0-rem shaping-rate 300m guaranteed-rate 200m
      scheduler-map smap-svlan0-rem
   user@host# set tcp-ifl1 scheduler-map smap-ifl1
   ```


Copyright © 2016, Juniper Networks, Inc.

3. Set the scheduler map.

```
[edit class-of-service scheduler-maps]
user@host# set smap-svlan0-rem forwarding-class best-effort scheduler-sched-foo
user@host# set smap-ifl1 forwarding-class best-effort scheduler-sched-bar
user@host# set smap-ifl1 forwarding-class assured-forwarding scheduler-sched-bar
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service interfaces**, **show class-of-service traffic-control-profiles** , and **show class-of-service scheduler-maps** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it. Example 2 does not include the **[edit interfaces]** configuration.

```
[edit]
user@host# show class-of-service interfaces
interface-set {
  svlan0 {
    output-traffic-control-profile tcp-svlan0; # Guarantee & shaper for svlan0
  }
}
ge-1/0/0 {
  output-traffic-control-profile-remaining tcp-svlan0-rem
    # Unit 3 and 4 are not explicitly configured, but captured by “remaining'
    unit 1 {
    output-traffic-control-profile tcp-ifl1; # Unit 1 be & ef queues
  }
}

[edit]
user@host# show class-of-service traffic-control-profiles
tcp-svlan0 {
  shaping-rate 200m;
  guaranteed-rate 100m;
}
tcp-svlan0-rem {
  shaping-rate 300m;
  guaranteed-rate 200m;
  scheduler-map smap-svlan0-rem; # This specifies queues for remaining traffic
}
tcp-ifl1 {
  scheduler-map smap-ifl1;
}

[edit]
user@host# show class-of-service scheduler-maps
smap-svlan0-rem {
  forwarding-class best-effort scheduler sched-foo;
}
smap-ifl1 {
  forwarding-class best-effort scheduler sched-bar;
  forwarding-class assured-forwarding scheduler sched-baz;
}
```

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

Copyright © 2016, Juniper Networks, Inc.

The configuration for the referenced schedulers is not given for this example.

Verification

Verifying Remaining Traffic Control

Purpose Verify that the remaining traffic is controlled properly.

Action From operational mode, enter the following commands:

- **show class-of-service interfaces**
- **show class-of-service traffic-control-profiles**
- **show class-of-service scheduler-maps**

Related Documentation

- Understanding Hierarchical Schedulers on page 1557

Copyright © 2016, Juniper Networks, Inc.

PART 19

Configuring Class of Service for IPv6

- Configuring Class of Service for IPv6 Traffic on page 1581

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 78

Configuring Class of Service for IPv6 Traffic

- CoS Functions for IPv6 Traffic Overview on page 1581
- Understanding CoS with DSCP IPv6 BA Classifier on page 1583
- Example: Configuring CoS with DSCP IPv6 BA Classifiers on page 1585
- Understanding DSCP IPv6 Rewrite Rules on page 1588
- Example: Configuring CoS with DSCP IPv6 Rewrite Rules on page 1588

CoS Functions for IPv6 Traffic Overview

Class-of-service (CoS) processing for IPv6 traffic uses the IPv6 DiffServ code point (DSCP) value. The IPv6 DSCP value is the first six bits in the 8-bit Traffic Class field of the IPv6 header. The DSCP value is used to determine the behavior aggregate (BA) classification for the packet entering the network device. You use classifier rules to map the DSCP code points to a forwarding class and packet loss priority. You use rewrite rules to map the forwarding class and packet loss priority back to DSCP values on packets exiting the device.

Figure 81 shows the components of the CoS features for Juniper Networks devices, illustrating the sequence in which they interact.

Figure 81: Packet Flow Through an SRX Series Device

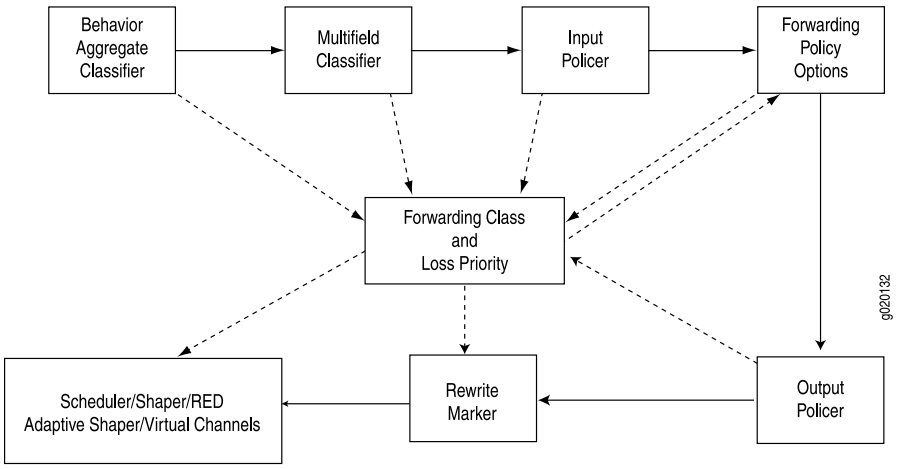

Copyright © 2016, Juniper Networks, Inc.

NOTE: Not all CoS features are supported on all devices.

- CoS components perform the following operations:

  BA classifier rules map DSCP code points to a forwarding class and loss priority. The forwarding class and loss priority determine the per-hop behavior of the packet throughout the system. The forwarding class associates a packet with an outbound transmission queue. Loss priority affects the scheduling of a packet without affecting the relative ordering of packets. BA classification is a simple way that "downstream" nodes can honor the CoS objectives that were encoded "upstream."

  See "Example: Configuring CoS with DSCP IPv6 BA Classifiers" on page 1585.

- Multifield classifier rules overwrite the initial forwarding class and loss priority determination read by the BA classifier rule. You typically use multifield classifier rules on nodes close to the content origin, where a packet might not have been encoded with the desired DSCP values in the headers. A multifield classifier rule assigns packets to a forwarding class and assigns a packet loss priority based on filters, such as source IP, destination IP, port, or application.

  See "Example: Configuring and Applying a Firewall Filter for a Multifield Classifier" on page 1445.

- Input policers meter traffic to see if traffic flow exceeds its service level. Policers might discard, change the forwarding class and loss priority, or set the packet loss priority bit of a packet. A packet for which the packet loss priority bit is set has an increased probability of being dropped during congestion.

- Scheduler maps are applied to interfaces and associate the outgoing packets with a scheduler and a forwarding class.

  The scheduler manages the output transmission queue, including:

  - Buffer size—Defines the period for which a packet is stored during congestion.
  - Scheduling priority and transmit rate—Determines the order in which a packet is transmitted.
  - Drop profile—Defines how aggressively to drop a packet that is using a particular scheduler.

  See "Example: Configuring Class-of-Service Schedulers" on page 1485.

- Output policers meter traffic and might change the forwarding class and loss priority of a packet if a traffic flow exceeds its service level.

- Rewrite rules map forwarding class and packet loss priority to DSCP values. You typically use rewrite rules in conjunction with multifield classifier rules close to the content origin, or when the device is at the border of a network and must alter the code points to meet the policies of the targeted peer.

  See "Example: Configuring CoS with DSCP IPv6 Rewrite Rules" on page 1588.

Copyright © 2016, Juniper Networks, Inc.

Only BA classification rules and rewrite rules require special consideration to support CoS for IPv6 traffic. The program logic for the other CoS features is not sensitive to differences between IPv4 and IPv6 traffic.

Related Documentation

- Understanding CoS with DSCP IPv6 BA Classifier on page 1583
- Example: Configuring CoS with DSCP IPv6 BA Classifiers on page 1585
- Understanding DSCP IPv6 Rewrite Rules on page 1588
- Example: Configuring CoS with DSCP IPv6 Rewrite Rules on page 1588

Understanding CoS with DSCP IPv6 BA Classifier

A behavior aggregate (BA) classifier rule maps DSCP code points to a forwarding class and loss priority. The forwarding class and loss priority determine the per-hop behavior of the packet throughout the system. The forwarding class associates a packet with an outbound transmission queue. Loss priority affects the scheduling of a packet without affecting the relative ordering of packets.

BA classification can be applied within one DiffServ domain or between two domains, where each domain honors the CoS results generated by the other domain. Table 148 shows the mapping for the default DSCP IPv6 BA classifier.

Table 148: Default IPv6 BA Classifier Mapping

Code Points	DSCP IPv6 Alias	Forwarding Class	Packet Loss Priority
101110	ef	expedited-forwarding	low
001010	af11	assured-forwarding	low
001100	af12	assured-forwarding	high
001110	af13	assured-forwarding	high
010010	af21	best-effort	low
010100	af22	best-effort	low
010110	af23	best-effort	low
011010	af31	best-effort	low
011100	af32	best-effort	low
011110	af33	best-effort	low
100010	af41	best-effort	low
100100	af42	best-effort	low

Copyright © 2016, Juniper Networks, Inc.

Table 148: Default IPv6 BA Classifier Mapping *(continued)*

Code Points	DSCP IPv6 Alias	Forwarding Class	Packet Loss Priority
100110	af43	best-effort	low
000000	be	best-effort	low
001000	cs1	best-effort	low
010000	cs2	best-effort	low
011000	cs3	best-effort	low
100000	cs4	best-effort	low
101000	cs5	best-effort	low
110000	nc1/cs6	network-control	low
111000	nc2/cs7	network-control	low

You can use the CLI **show** command to display the settings for the CoS classifiers. The following command shows the settings for the default DSCP IPv6 classifier:

```
user@host# show class-of-service classifier type dscp-ipv6
Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 8
  Code point         Forwarding class                  Loss priority
  000000             best-effort                       low
  000001             best-effort                       low
  000010             best-effort                       low
  000011             best-effort                       low
  000100             best-effort                       low
  000101             best-effort                       low
  011011             best-effort                       low
...
Classifier: dscp-ipv6-compatibility, Code point type: dscp-ipv6, Index: 9
  Code point         Forwarding class                  Loss priority
  000000             best-effort                       low
  000001             best-effort                       low
  000010             best-effort                       low
  000011             best-effort                       low
  000100             best-effort                       low
  000101             best-effort                       low
  000110             best-effort                       low
  000111             best-effort                       low
 ...
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: The predefined classifier named dscp-ipv6-compatibility maps all code point loss priorities to low. It maps 110000 and 111000 (typically seen in network control packets) to the network-control class and all other code points to the best-effort class. The dscp-ipv6-compatibility classifier is an implicit classifier similar to ipprec-compatibility, which is provided to map IP precedence bits in IPv4 traffic when no classifier has been configured.

Related Documentation

- Example: Configuring CoS with DSCP IPv6 BA Classifiers on page 1585
- CoS Functions for IPv6 Traffic Overview on page 1581
- Understanding DSCP IPv6 Rewrite Rules on page 1588
- Example: Configuring CoS with DSCP IPv6 Rewrite Rules on page 1588

Example: Configuring CoS with DSCP IPv6 BA Classifiers

This example shows how to associate an interface with a default or user-defined DSCP IPv6 BA classifier.

- Requirements on page 1585
- Overview on page 1585
- Configuration on page 1585
- Verification on page 1587

Requirements

Before you begin, configure the ge-0/0/0 interface on the device for IPv6 and define your user-defined DSCP IPv6 classifier settings. See "Understanding CoS with DSCP IPv6 BA Classifier" on page 1583.

Overview

In this example, you configure CoS and define forwarding classes. You create the behavior aggregate classifier for DiffServ CoS as dscp-ipv6-example and import the default DSCP IPv6 classifier.

You then specify the best-effort forwarding class as be-class, the expedited forwarding class as ef-class, the assured forwarding class as af-class, and the network control forwarding class as nc-class. Finally, you apply your user-defined classifier to interface ge-0/0/0.

Configuration

CLI Quick Configuration

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

```
set class-of-service forwarding-classes queue 0 be-class
```

Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service forwarding-classes queue 1 ef-class
set class-of-service forwarding-classes queue 2 af-class
set class-of-service forwarding-classes queue 3 nc-class
set class-of-service classifiers dscp-ipv6 dscp-ipv6-example import default
set class-of-service classifiers dscp-ipv6 dscp-ipv6-example forwarding-class be-class
  loss-priority high code-points 000001
set class-of-service classifiers dscp-ipv6 dscp-ipv6-example forwarding-class ef-class
  loss-priority high code-points 101111
set class-of-service classifiers dscp-ipv6 dscp-ipv6-example forwarding-class af-class
  loss-priority high code-points 001100
set class-of-service classifiers dscp-ipv6 dscp-ipv6-example forwarding-class nc-class
  loss-priority high code-points 110001
set class-of-service interfaces ge-0/0/0 unit 0 classifiers dscp-ipv6 dscp-ipv6-example
```

Step-by-Step Procedure

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

To configure CoS with a user-defined DSCP IPv6 BA classifier:

1. Configure CoS.

   ```
   [edit]
   user@host# edit class-of-service
   ```

2. Define forwarding classes.

   ```
   [edit class-of-service]
   user@host# set forwarding-classes queue 0 be-class
   user@host# set forwarding-classes queue 1 ef-class
   user@host# set forwarding-classes queue 2 af-class
   user@host# set forwarding-classes queue 3 nc-class
   ```

3. Create a behavior aggregate classifier for DiffServ CoS.

   ```
   [edit class-of-service]
   user@host# edit classifiers dscp-ipv6 dscp-ipv6-example
   ```

4. Import a DSCP IPv6 classifier.

   ```
   [edit class-of-service classifiers dscp-ipv6 dscp-ipv6-example]
   user@host# set import default
   ```

5. Specify a best-effort forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp-ipv6 dscp-ipv6-example]
   user@host# set forwarding-class be-class loss-priority high code-points 000001
   ```

6. Specify an expedited forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp-ipv6 dscp-ipv6-example]
   user@host# set forwarding-class ef-class loss-priority high code-points 101111
   ```

7. Specify an assured forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp-ipv6 dscp-ipv6-example]
   user@host# set forwarding-class af-class loss-priority high code-points 001100
   ```

8. Specify a network control forwarding class classifier.

   ```
   [edit class-of-service classifiers dscp-ipv6 dscp-ipv6-example]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set forwarding-class nc-class loss-priority high code-points 110001
```

9. Associate a user-defined classifier with an interface.

```
[edit class-of-service]
user@host# set interfaces ge-0/0/0 unit 0 classifiers dscp-ipv6 dscp-ipv6-example
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
classifiers {
  dscp-ipv6 dscp-ipv6-example {
  import default;
  forwarding-class be-class {
    loss-priority high code-points 000001;
  }
  forwarding-class ef-class {
    loss-priority high code-points 101111;
  }
  forwarding-class af-class {
    loss-priority high code-points 001100;
  }
  forwarding-class nc-class {
    loss-priority high code-points 110001;
    }
  }
  }
  forwarding-classes {
    queue 0 be-class;
    queue 1 ef-class;
    queue 2 af-class;
    queue 3 nc-class;
  }
  interfaces {
    ge-0/0/0 {
    unit 0 {
    classifiers {
    dscp-ipv6 dscp-ipv6-example;
      }
    }
  }
}
```

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

Verification

Verifying the CoS with DSCP IPv6 BA Classifier Configuration

Purpose Verify that the user-defined DSCP IPv6 BA classifier is associated with an interface.

Action From configuration mode, enter the **show class-of-service** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding CoS with DSCP IPv6 BA Classifier on page 1583
- CoS Functions for IPv6 Traffic Overview on page 1581
- Understanding DSCP IPv6 Rewrite Rules on page 1588
- Example: Configuring CoS with DSCP IPv6 Rewrite Rules on page 1588

Understanding DSCP IPv6 Rewrite Rules

After Junos OS CoS processing, a rewrite rule maps the forwarding class and loss priority after Junos OS CoS processing to a corresponding DSCP value specified in the rule. Typically, you use rewrite rules to alter the CoS values in outgoing packets to meet the requirements of the targeted peer.

You can use the CLI show command to display the configuration for the CoS classifiers. The following command shows the configuration of the default DSCP IPv6 rewrite rule:

```
user@host# show class-of-service rewrite-rule type dscp-ipv6
Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 32
  Forwarding class                Loss priority       Code point
  best-effort                     low                 000000
  best-effort                     high                000000
  expedited-forwarding            low                 101110
  expedited-forwarding            high                101110
  assured-forwarding              low                 001010
  assured-forwarding              high                001100
  network-control                 low                 110000
  network-control                 high                111000
```

Related Documentation

- Example: Configuring CoS with DSCP IPv6 Rewrite Rules on page 1588
- CoS Functions for IPv6 Traffic Overview on page 1581
- Understanding CoS with DSCP IPv6 BA Classifier on page 1583
- Example: Configuring CoS with DSCP IPv6 BA Classifiers on page 1585

Example: Configuring CoS with DSCP IPv6 Rewrite Rules

This example shows how to associate an interface with a default or user-defined DSCP IPv6 rewrite rule. Typically, you use rewrite rules to alter CoS values in outgoing packets to meet the requirements of the targeted peer.

- Requirements on page 1588
- Overview on page 1589
- Configuration on page 1589
- Verification on page 1591

Requirements

Before you begin, configure the ge-0/0/0 interface on the device for IPv6 and define your user-defined DSCP IPv6 rewrite rules.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you configure CoS and create a user-defined rewrite rule called rewrite-ipv6-dscps. You then specify rewrite rules for the best-effort forwarding class as be-class, the expedited forwarding class as ef-class, the assured forwarding class as af-class, and the network control forwarding class as nc-class. Finally, you associate interface ge-0/0/0 with the user-defined rule.

Configuration

CLI Quick Configuration

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

```
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class be-class
    loss-priority low code-point 000000
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class be-class
    loss-priority high code-point 000001
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class ef-class
    loss-priority low code-point 101110
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class ef-class
    loss-priority high code-point 101111
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class af-class
    loss-priority low code-point 001010
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class af-class
    loss-priority high code-point 001100
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class nc-class
    loss-priority low code-point 110000
set class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps forwarding-class nc-class
    loss-priority high code-point 110001
set class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules dscp rewrite-dscps
```

Step-by-Step Procedure

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

To configure a CoS with a user-defined DSCP IPv6 rewrite rule:

1. Configure CoS.

   ```
   [edit]
   user@host# edit class-of-service
   ```

2. Create a user-defined rewrite rule.

   ```
   [edit class-of-service]
   user@host# edit rewrite-rules dscp-ipv6 rewrite-ipv6-dscps
   ```

3. Specify rewrite rules for the best-effort forwarding class.

   ```
   [edit class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps]
   user@host# set forwarding-class be-class loss-priority low code-point 000000
   user@host# set forwarding-class be-class loss-priority high code-point 000001
   ```

Copyright © 2016, Juniper Networks, Inc.

4. Specify rewrite rules for the expedited-forwarding forwarding class.

```
[edit class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps]
user@host# set forwarding-class ef-class loss-priority low code-point 101110
user@host# set forwarding-class ef-class loss-priority high code-point 101111
```

5. Specify rewrite rules for the assured-forwarding forwarding class.

```
[edit class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps]
user@host# set forwarding-class af-class loss-priority low code-point 001010
user@host# set forwarding-class af-class loss-priority high code-point 001100
```

6. Specify rewrite rules for the network-control forwarding class.

```
[edit class-of-service rewrite-rules dscp-ipv6 rewrite-ipv6-dscps]
user@host# set forwarding-class nc-class loss-priority low code-point 110000
user@host# set forwarding-class nc-class loss-priority high code-point 110001
```

7. Associate an interface with a user-defined rule.

```
[edit class-of-service]
user@host# set interfaces ge-0/0/0 unit 0 rewrite-rules dscp rewrite-dscps
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show class-of-service
interfaces {
   ge-0/0/0 {
      unit 0 {
         rewrite-rules {
            dscp rewrite-dscps;
         }
      }
   }
}
rewrite-rules {
   dscp-ipv6 rewrite-ipv6-dscps {
      forwarding-class be-class {
         loss-priority low code-point 000000;
         loss-priority high code-point 000001;
      }
      forwarding-class ef-class {
         loss-priority low code-point 101110;
         loss-priority high code-point 101111;
      }
      forwarding-class af-class {
         loss-priority low code-point 001010;
         loss-priority high code-point 001100;
      }
      forwarding-class nc-class {
         loss-priority low code-point 110000;
         loss-priority high code-point 110001;
      }
   }
}
```


Copyright © 2016, Juniper Networks, Inc.

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

Verification

Verifying the CoS with DSCP IPv6 Rewrite Rule Configuration

Purpose Verify that the user-defined CoS with DSCP IPv6 rewrite rule is associated with an interface.

Action From configuration mode, enter the **show class-of-service** command.

Related Documentation

- Understanding DSCP IPv6 Rewrite Rules on page 1588
- CoS Functions for IPv6 Traffic Overview on page 1581
- Understanding CoS with DSCP IPv6 BA Classifier on page 1583
- Example: Configuring CoS with DSCP IPv6 BA Classifiers on page 1585

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 20

Configuring Class of Service for I/O Cards

- Configuring Class of Service for I/O Cards on page 1595

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 79

Configuring Class of Service for I/O Cards

- PIR-Only and CIR Mode Overview on page 1595
- Understanding Priority Propagation on page 1597
- Understanding IOC Hardware Properties on page 1598
- Understanding IOC Map Queues on page 1600
- WRED on the IOC Overview on page 1601
- MDRR on the IOC Overview on page 1605
- CoS Support on the SRX5000 Module Port Concentrator Overview on page 1607
- Example: Configuring CoS on High-End SRX Series Devices with an MPC on page 1608

PIR-Only and CIR Mode Overview

The actual behavior of many CoS parameters, especially the shaping rate and guaranteed rate, depends on whether the physical interface is operating in one of the following modes:

- PIR-only Mode on page 1595
- CIR Mode on page 1596

PIR-only Mode

In PIR-only (peak information rate) mode, one or more nodes perform shaping. The physical interface is in PIR-only mode if no child (or grandchild) node under the port has a guaranteed rate configured. The mode of the port is important because in PIR-only mode, the scheduling across the child nodes is in proportion to their shaping rates (PIRs) and not the guaranteed rates (CIRs). This can be important if the observed behavior is not what is anticipated.

In PIR-only mode, nodes cannot send if they are above the configured shaping rate. Table 149 shows the mapping between the configured priority and the hardware priority for PIR-only.

Table 149: Internal Node Queue Priority for PIR-Only Mode

Configured Priority	Hardware Priority
Strict-high	0

Copyright © 2016, Juniper Networks, Inc.

Table 149: Internal Node Queue Priority for PIR-Only Mode *(continued)*

Configured Priority	Hardware Priority
High	0
Medium-high	1
Medium-low	1
Low	2

CIR Mode

In CIR (committed information rate) mode, one or more nodes applies a guaranteed rate and might perform shaping. A physical interface is in CIR mode if at least one child (or grandchild) node has a guaranteed rate configured. In addition, any child or grandchild node under the physical interface can have a shaping rate configured. Only the guaranteed rate matters. In CIR mode, nodes that do not have a guaranteed rate configured are assumed to have a very small guaranteed rate (queuing weight).

In CIR mode, the priority for each internal node depends on whether the highest active child node is above or below the guaranteed rate. Table 150 shows the mapping between the highest active child's priority and the hardware priority below and above the guaranteed rate.

Table 150: Internal Node Queue Priority for CIR Mode

Configured Priority of Highest Active Child Node	Hardware Priority Below Guaranteed Rate	Hardware Priority Above Guaranteed Rate
Strict-high	0	0
High	0	3
Medium-high	1	3
Medium-low	1	3
Low	2	3

Related Documentation

- Understanding Priority Propagation on page 1597
- Understanding IOC Hardware Properties on page 1598
- Understanding IOC Map Queues on page 1600
- WRED on the IOC Overview on page 1601
- MDRR on the IOC Overview on page 1605

Copyright © 2016, Juniper Networks, Inc.

Understanding Priority Propagation

SRX5600 and SRX5800 devices with input/output cards (IOCs) perform priority propagation. Priority propagation is useful for mixed traffic environments when, for example, you want to make "Understanding IOC Map Queues" on page 1600 sure that the voice traffic of one customer does not suffer from the data traffic of another customer. Nodes and queues are always serviced in the order of their priority. The priority of a queue is decided by configuration (the default priority is low) in the scheduler. However, not all elements of hierarchical schedulers have direct priorities configured. Internal nodes, for example, must determine their priority in other ways.

The priority of any internal node is decided as follows:

- By the highest priority of an active child (interface sets only take the highest priority of their active children
- Whether the node is above its configured guaranteed rate (CIR) or not (this is relevant only if the physical interface is in CIR mode)

Each queue has a configured priority and a hardware priority. Table 151 shows the usual mapping between the configured priority and the hardware priority.

Table 151: Queue Priority

Configured Priority	Hardware Priority
Strict-high	0
High	0
Medium-high	1
Medium-low	1
Low	2

Figure 82 shows a physical interface with hierarchical schedulers configured. The configured priorities are shown for each queue at the top of the figure. The hardware priorities for each node are shown in parentheses. Each node also shows any configured shaping rate (PIR) or guaranteed rate (CIR) and whether or not the queues are above or below the CIR. The nodes are shown in one of the following three states:

- Above the CIR (clear)
- Below the CIR (dark)
- Condition where the CIR does not matter (gray)

Copyright © 2016, Juniper Networks, Inc.

Figure 82: Hierarchical Schedulers and Priorities

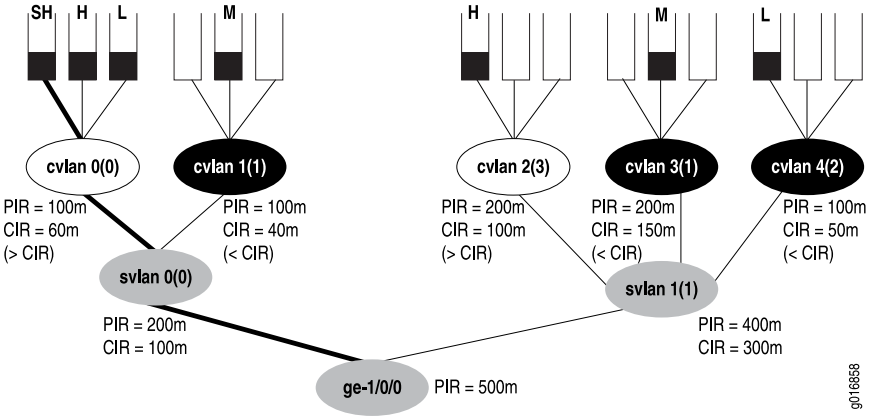

In Figure 82, the strict high queue for C-VLAN 0 (cvlan 0) receives service first, even though the C-VLAN is above the configured CIR. Once that queue has been drained, and the priority of the node has become 3 instead of 0 (because of the lack of strict-high traffic), the system moves on to the medium queues (cvlan 1 and cvlan 3), draining them in a round-robin fashion where empty queues lose their hardware priority. The low queue on cvlan 4 (priority 2) is sent next because that mode is below the CIR. Then, the high queues on cvlan 0 and cvlan2 (both now with priority 3) are drained in a round-robin fashion, and finally the low queue on cvlan 0 is drained (because svlan 0 has a priority of 3).

Related Documentation

- PIR-Only and CIR Mode Overview on page 1595
- Understanding IOC Hardware Properties on page 1598
- Understanding IOC Map Queues on page 1600
- WRED on the IOC Overview on page 1601
- MDRR on the IOC Overview on page 1605

Understanding IOC Hardware Properties

On SRX5600 and SRX5800 devices, two IOCs (40x1GE IOC and 4x10GE IOC) are supported on which you can configure schedulers and queues. You can configure 15 VLAN sets per Gigabit Ethernet (40x1GE IOC) port and 255 VLAN sets per 10-Gigabit Ethernet (4x10GE IOC) port. The IOC performs priority propagation from one hierarchy level to another, and drop statistics are available on the IOC per color per queue instead of just per queue.

SRX5600 and SRX5800 devices with IOCs have Packet Forwarding Engines that can support up to 512 MB of frame memory, and packets are stored in 512-byte frames. Table 152 compares the major properties of the Packet Forwarding Engine within the IOC.

Copyright © 2016, Juniper Networks, Inc.

Table 152: Packet Forwarding Engine Properties within 40x1GE IOC and 4x10GE IOC

Feature	PFE Within 40x1GE IOC and 4x10GE IOC
Number of usable queues	16,000
Number of shaped logical interfaces	2,000 with 8 queues each, or 4,000 with 4 queues each.
Number of hardware priorities	4
Priority propagation	Yes
Dynamic mapping	Yes: schedulers per port are not fixed.
Drop statistics	Per queue per color (PLP high, low)

Additionally, the IOC features also support hierarchical weighted random early detection (WRED).

The IOC supports the following hierarchical scheduler characteristics:

- Shaping at the physical interface level
- Shaping and scheduling at the service VLAN interface set level
- Shaping and scheduling at the customer VLAN logical interface level
- Scheduling at the queue level

The IOC supports the following features for scalability:

- 16,000 queues per PFE
- 4 PFEs per IOC
  - 4000 schedulers at logical interface level (level 3) with 4 queues each
  - 2000 schedulers at logical interface level (level 3) with 8 queues each
- 255 schedulers at the interface set level (level 2) per 1-port PFE on a 10-Gigabit Ethernet IOC (4x10GE IOC)
- 15 schedulers at the interface set level (level 2) per 10-port PFE on a 1-Gigabit Ethernet IOC (40x1GE IOC)
- About 400 milliseconds of buffer delay (this varies by packet size and if large buffers are enabled)
- 4 levels of priority (strict-high, high, medium, and low)

NOTE: **The exact option for a transmit-rate (transmit-rate rate exact) is not supported on the IOCs on SRX Series devices.**

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- PIR-Only and CIR Mode Overview on page 1595
- Understanding Priority Propagation on page 1597
- Understanding IOC Map Queues on page 1600
- WRED on the IOC Overview on page 1601
- MDRR on the IOC Overview on page 1605

Understanding IOC Map Queues

The manner in which the IOC maps a queue to a scheduler depends on whether 8 queues or 4 queues are configured. By default, a scheduler at level 3 has 4 queues. Level 3 scheduler X controls queue X*4 to X*4+3, so that scheduler 100 (for example) controls queues 400 to 403. However, when 8 queues per scheduler are enabled, the odd-numbered schedulers are disabled, allowing twice the number of queues per subscriber as before. With 8 queues, level 3 scheduler X controls queue X*4 to X*4+7, so that scheduler 100 (for example) now controls queues 400 to 407.

You configure the **max-queues-per-interface** statement to set the number of queues at 4 or 8 at the FPC level of the hierarchy. Changing this statement will result in a restart of the FPC.

The IOC maps level 3 (customer VLAN) schedulers in groups to level 2 (service VLAN) schedulers. Sixteen contiguous level 3 schedulers are mapped to level 2 when 4 queues are enabled, and 8 contiguous level 3 schedulers are mapped to level 2 when 8 queues are enabled. All the schedulers in the group should use the same queue priority mapping. For example, if the queue priorities of one scheduler are high, medium, and low, all members of the group should have the same queue priority.

Groups at level 3 to level 2 can be mapped at any time. However, a group at level 3 can only be unmapped from a level 2 scheduler, and only if all the schedulers in the group are free. Once unmapped, a level 3 group can be remapped to any level 2 scheduler. There is no restriction on the number of level 3 groups that can be mapped to a particular level 2 scheduler. There can be 256 level 3 groups, but fragmentation of the scheduler space can reduce the number of schedulers available. In other words, there are scheduler allocation patterns that might fail even though there are free schedulers.

In contrast to level 3 to level 2 mapping, the IOC maps level 2 (service VLAN) schedulers in a fixed mode to level 1 (physical interface) schedulers. On 40-port Gigabit Ethernet IOCs, there are 16 level 1 schedulers, and 10 of these are used for the physical interfaces. There are 256 level 2 schedulers, or 16 per level 1 schedulers. A level 1 scheduler uses level schedulers X*16 through X*16+15. Therefore level 1 scheduler 0 uses level 2 schedulers 0 through 15, level 1 scheduler 1 uses level 2 schedulers 16 through 31, and so on. On 4-port 10-Gigabit Ethernet PICs, there is one level 1 scheduler for the physical interface, and 256 level 2 schedulers are mapped to the single level 1 scheduler.

The maximum number of level 3 (customer VLAN) schedulers that can be used is 4076 (4 queues) or 2028 (8 queues) for the 10-port Gigabit Ethernet Packet Forwarding Engine

Copyright © 2016, Juniper Networks, Inc.

and 4094 (4 queues) or 2046 (8 queues) for the 10-Gigabit Ethernet Packet Forwarding Engine.

Related Documentation

- PIR-Only and CIR Mode Overview on page 1595
- Understanding Priority Propagation on page 1597
- Understanding IOC Hardware Properties on page 1598
- WRED on the IOC Overview on page 1601
- MDRR on the IOC Overview on page 1605

WRED on the IOC Overview

Shaping to drop out-of-profile traffic is done on the IOC at all levels except the queue level. However, weighed random early discard (WRED) is done at the queue level with much the same result. With WRED, the decision to drop or send the packet is made before the packet is placed in the queue.

WRED shaping on the IOC involves two levels. The probabilistic drop region establishes a minimum and a maximum queue depth. Below the minimum queue depth, the drop probability is 0 (send). Above the maximum level, the drop probability is 100 (certainty).

There are four drop profiles associated with each queue. These correspond to each of four loss priorities (low, medium-low, medium-high, and high). Sixty-four sets of four drop profiles are available (32 for ingress and 32 for egress). In addition, there are eight WRED scaling profiles in each direction.

An example of an IOC drop profile for expedited forwarding traffic is as follows:

```
[edit class-of-service drop-profiles]
drop-ef {
   fill-level 20 drop-probability 0; # Minimum Q depth
   fill-level 100 drop-probability 100; # Maximum Q depth
}
```

NOTE: You can specify only two fill levels for the IOC.

You can configure the **interpolate** statement, but only two fill levels are used. The **delay-buffer-rate** statement in the traffic control profile determines the maximum queue size. This delay buffer rate is converted to packet delay buffers, where one buffer is equal to 512 bytes. For example, at 10 Mbps, the IOC will allocate 610 delay buffers when the delay buffer rate is set to 250 milliseconds. The WRED threshold values are specified in terms of absolute buffer values.

The WRED scaling factor multiples all WRED thresholds (both minimum and maximum) by the value specified. There are eight values in all: 1, 2, 4, 8, 16, 32, 64, and 128. The WRED scaling factor is chosen to best match the user-configured drop profiles. This is done because the hardware supports only certain values of thresholds (all values must be a multiple of 16). So if the configured value of a threshold is 500 (for example), the multiple

Copyright © 2016, Juniper Networks, Inc.

of 16 is 256 and the scaling factor applied is 2, making the value 512, which allows the value of 500 to be used. If the configured value of a threshold is 1500, the multiple of 16 is 752 and the scaling factor applied is 2, making the value 1504, which allows the value of 1500 to be used.

Hierarchical RED is used to support the oversubscription of the delay buffers (WRED is configured only at the queue, physical interface, and PIC levels). Hierarchical RED works with WRED as follows:

- If any level accepts the packet (the queue depth is less than the minimum buffer levels), this level accepts the packet.
- If any level probabilistically drops the packet, then this level drops the packet.

However, these rules might lead to the accepting of packets under loaded conditions that might otherwise have been dropped. In other words, the logical interface will accept packets if the physical interface is not congested.

Because of the limits placed on shaping thresholds used in the hierarchy, there is a granularity associated with the IOCs. The shaper accuracies differ at various levels of the hierarchy:

- Level 3
- Level 2
- Level 1

Shapers at the logical interface level (level 3) are more accurate than shapers at the interface set level (level 2) or at the port level (level 1).

This section contains the following topics:

- Shapers at the Logical Interface Level (Level 3) on page 1602
- Shapers at the Interface Set Level (Level 2) on page 1603
- Shapers at the Port Level (Level 1) on page 1604

Shapers at the Logical Interface Level (Level 3)

Because of the limits placed on shaping thresholds used in the hierarchy, there is a granularity associated with the IOCs. The shaper accuracies differ at various levels of the hierarchy, with shapers at the logical interface level (level 3) being more accurate than shapers at the interface set level (level 2) or at the port level (level 1). Table 153 shows the accuracy of the logical interface shaper at various speeds for Ethernet ports operating at 1 Gbps.

Table 153: Shaper Accuracy of 1-Gbps Ethernet at the Logical Interface Level

Range of Logical Interface Shaper	Step Granularity
Up to 4.096 Mbps	16 Kbps

Copyright © 2016, Juniper Networks, Inc.

Table 153: Shaper Accuracy of 1-Gbps Ethernet at the Logical Interface Level *(continued)*

Range of Logical Interface Shaper	Step Granularity
4.096 to 8.192 Mbps	32 Kbps
8.192 to 16.384 Mbps	64 Kbps
16.384 to 32.768 Mbps	128 Kbps
32.768 to 65.535 Mbps	256 Kbps
65.535 to 131.072 Mbps	512 Kbps
131.072 to 262.144 Mbps	1024 Kbps
262.144 to 1 Gbps	4096 Kbps

Table 154 shows the accuracy of the logical interface shaper at various speeds for Ethernet ports operating at 10 Gbps.

Table 154: Shaper Accuracy of 10-Gbps Ethernet at the Logical Interface Level

Range of Logical Interface Shaper	Step Granularity
Up to 10.24 Mbps	40 Kbps
10.24 to 20.48 Mbps	80 Kbps
10.48 to 40.96 Mbps	160 Kbps
40.96 to 81.92 Mbps	320 Kbps
81.92 to 163.84 Mbps	640 Kbps
163.84 to 327.68 Mbps	1280 Kbps
327.68 to 655.36 Mbps	2560 Kbps
655.36 to 2611.2 Mbps	10240 Kbps
2611.2 to 5222.4 Mbps	20480 Kbps
5222.4 to 10 Gbps	40960 Kbps

Shapers at the Interface Set Level (Level 2)

Table 155 shows the accuracy of the interface set shaper at various speeds for Ethernet ports operating at 1 Gbps.

Copyright © 2016, Juniper Networks, Inc.

Table 155: Shaper Accuracy of 1-Gbps Ethernet at the Interface Set Level

Range of Interface Set Shaper	Step Granularity
Up to 20.48 Mbps	80 Kbps
20.48 Mbps to 81.92 Mbps	320 Kbps
81.92 Mbps to 327.68 Mbps	1.28 Mbps
327.68 Mbps to 1 Gbps	20.48 Mbps

Table 156 shows the accuracy of the interface set shaper at various speeds for Ethernet ports operating at 10 Gbps.

Table 156: Shaper Accuracy of 10-Gbps Ethernet at the Interface Set Level

Range of Interface Set Shaper	Step Granularity
Up to 128 Mbps	500 Kbps
128 Mbps to 512 Mbps	2 Mbps
512 Mbps to 2.048 Gbps	8 Mbps
2.048 Gbps to 10 Gbps	128 Mbps

Shapers at the Port Level (Level 1)

Table 157 shows the accuracy of the physical port shaper at various speeds for Ethernet ports operating at 1 Gbps.

Table 157: Shaper Accuracy of 1-Gbps Ethernet at the Physical Port Level

Range of Physical Port Shaper	Step Granularity
Up to 64 Mbps	250 Kbps
64 Mbps to 256 Mbps	1 Mbps
256 Mbps to 1 Gbps	4 Mbps

Table 158 shows the accuracy of the physical port shaper at various speeds for Ethernet ports operating at 10 Gbps.

Table 158: Shaper Accuracy of 10-Gbps Ethernet at the Physical Port Level

Range of Physical Port Shaper	Step Granularity
Up to 640 Mbps	2.5 Mbps

Copyright © 2016, Juniper Networks, Inc.

Table 158: Shaper Accuracy of 10-Gbps Ethernet at the Physical Port Level *(continued)*

Range of Physical Port Shaper	Step Granularity
640 Mbps to 2.56 Gbps	10 Mbps
2.56 Gbps to 10 Gbps	40 Mbps

Related Documentation

- PIR-Only and CIR Mode Overview on page 1595
- Understanding Priority Propagation on page 1597
- Understanding IOC Hardware Properties on page 1598
- Understanding IOC Map Queues on page 1600
- MDRR on the IOC Overview on page 1605

MDRR on the IOC Overview

The guaranteed rate CIR at the interface set level is implemented by using modified deficit round-robin (MDRR). The IOC hardware provides four levels of strict priority. There is no restriction on the number of queues for each priority. MDRR is used among queues of the same priority. Each queue has one priority when it is under the guaranteed rate and another priority when it is over the guaranteed rate but still under the shaping rate PIR. The IOC hardware implements the priorities with 256 service profiles. Each service profile assigns eight priorities for eight queues. One set is for logical interfaces under the guaranteed rate and another set is for logical interfaces over the guaranteed rate but under the shaping rate. Each service profile is associated with a group of 16 level 3 schedulers, so there is a unique service profile available for all 256 groups at level 3, giving 4,096 logical interfaces.

Junos OS provides three priorities for traffic under the guaranteed rate and one reserved priority for traffic over the guaranteed rate that is not configurable. Junos OS provides three priorities when there is no guaranteed rate configured on any logical interface.

Table 159 shows the relationship between Junos OS priorities and the IOC hardware priorities below and above the guaranteed rate CIR.

Table 159: Junos Priorities Mapped to IOC Hardware Priorities

Junos OS Priority	IOC Hardware Priority Below Guaranteed Rate	IOC Hardware Priority Above Guaranteed Rate
Strict-high	High	High
High	High	Low
Medium-high	Medium-high	Low
Medium-low	Medium-high	Low

Copyright © 2016, Juniper Networks, Inc.

Table 159: Junos Priorities Mapped to IOC Hardware Priorities *(continued)*

Junos OS Priority	IOC Hardware Priority Below Guaranteed Rate	IOC Hardware Priority Above Guaranteed Rate
Low	Medium-low	Low

The Junos OS parameters are set in the scheduler map:

```
[edit class-of-service schedulers]
best-effort-scheduler {
   transmit-rate percent 30; # if no shaping rate
   buffer-size percent 30;
   priority high;
}
expedited-forwarding-scheduler {
   transmit-rate percent 40; # if no shaping rate
   buffer-size percent 40;
   priority strict-high;
}
```

NOTE: The use of both a shaping rate and a guaranteed rate at the interface set level (level 2) is not supported.

MDRR is provided at three levels of the scheduler hierarchy of the IOC with a granularity of 1 through 255. There are 64 MDRR profiles at the queue level, 16 at the interface set level, and 32 at the physical interface level.

Queue transmit rates are used for queue-level MDRR profile weight calculation. The queue MDRR weight is calculated differently based on the mode set for sharing excess bandwidth. If you configure the **equal** option for excess bandwidth, then the queue MDRR weight is calculated as:

Queue weight = (255 * Transmit-rate-percentage) / 100

If you configure the **proportional** option for excess bandwidth, which is the default, then the queue MDRR weight is calculated as:

Queue weight = Queue-transmit-rate / Queue-base-rate, where

Queue-transmit-rate = (Logical-interface-rate * Transmit-rate-percentage) / 100, and

Queue-base-rate = Excess-bandwidth-proportional-rate / 255

To configure the way that the IOC should handle excess bandwidth, configure the **excess-bandwidth-share** statement at the [edit interface-set *interface-set-name*] hierarchy level. By default, the excess bandwidth is set to **proportional** with a default value of 32.64 Mbps. In this mode, the excess bandwidth is shared in the ratio of the logical interface shaping rates. If set to **equal**, the excess bandwidth is shared equally among the logical interfaces.

Copyright © 2016, Juniper Networks, Inc.

The following example sets the excess bandwidth sharing to proportional at a rate of 100 Mbps with a shaping rate of 80 Mbps:

```
[edit interface-set example-interface-set]
excess-bandwidth-share proportional 100m;
output-traffic-control-profile PIR-80Mbps;
```

Shaping rates established at the logical interface level are used to calculate the MDRR weights used at the interface set level. The 16 MDRR profiles are set to initial values, and the closest profile with rounded values is chosen. By default, the physical port MDRR weights are preset to the full bandwidth on the interface.

Related Documentation

- PIR-Only and CIR Mode Overview on page 1595
- Understanding Priority Propagation on page 1597
- Understanding IOC Hardware Properties on page 1598
- Understanding IOC Map Queues on page 1600
- WRED on the IOC Overview on page 1601

CoS Support on the SRX5000 Module Port Concentrator Overview

The SRX5000 Module Port Concentrator (SRX5K-MPC) for the SRX5600 and SRX5800 uses the Trio chipset-based queuing model, which supports class of service (CoS) characteristics that are optimized compared to CoS characteristics supported by the standard queuing model. These CoS features enable SRX5600 and SRX5800 devices to achieve end-to-end quality of service and protect the network using various security functions.

CoS features on the SRX5600 and SRX5800 devices provide differentiated services to traffic in addition to the best-effort packet processing. The main CoS features include classification, CoS field rewriting, queuing, scheduling, and traffic shaping.

When a network experiences congestion and delay, you can use the CoS features to classify packets; assign them with different levels of packet loss priority, delay, and throughput; and mark their CoS-related fields defined in Layer 2 and Layer 3 headers.

The MPC supports the following CoS features:

- BA classifier based on IEEE 802.1p for packet classification (Layer 2 headers) for priority bits of ingress packets
- Rewrite rule based on IEEE 802.1p for priority bits of egress packets

NOTE: You can configure up to 32 IEEE 802.1p rewriters on each SRX5K-MPC on the SRX5600 and SRX5800 devices.

- Eight priority queues per port with configurable schedulers at the egress physical interface

Copyright © 2016, Juniper Networks, Inc.

By default, the MPC supports eight queues. You can use the following CLI statement to change that setting to four queues:

set chassis fpc *fpc-number* pic *pic-number* max-queues-per-interface *4*

Changing to four-queue mode limits that number of configurable queues to four on the MPC. This does not have any effect on the performance.

The CoS features on the MPC have the following limitations:

- On the MPC, the per-unit-scheduler or the hierarchical-scheduler is not supported. For egress scheduling and queuing, only the default mode is supported.
- When an SPU is too busy to process every ingress packet from the MPC, some high-priority packets, such as voice packets, might be delayed or dropped by the SRX5600 or SRX5800.

NOTE: **The total number of classifiers supported on a Services Processing Unit (SPU) is 79. Three classifiers are installed on the SPU as default classifiers in the Layer 3 mode, independent of any CoS configuration, which leaves 76 classifiers that can be configured using the CoS CLI commands. The default classifiers number might vary in future releases or in different modes. You can verify the number of default classifiers installed on the SPU to determine how many classifiers can be configured using the CoS CLI commands.**

Related Documentation

- Example: Configuring CoS on High-End SRX Series Devices with an MPC on page 1608

Example: Configuring CoS on High-End SRX Series Devices with an MPC

This example shows how to configure CoS on an SRX5000 line device with an MPC.

- Requirements on page 1608
- Overview on page 1609
- Configuration on page 1610
- Verification on page 1616

Requirements

This example uses the following hardware and software components:

- SRX5600 with an SRX5K-MPC
- Junos OS Release 12.1X46-D10 or later for SRX Series

Before you begin:

- Understand CoS. See "Understanding Class of Service" on page 1401.

Copyright © 2016, Juniper Networks, Inc.

- Understand chassis cluster configuration. See *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*.
- Understand chassis cluster redundant interface configuration. See *Example: Configuring Chassis Cluster Redundant Ethernet Interfaces for IPv4 and IPv6 Addresses*.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you create a BA classifier to classify traffic based on the IEEE 802.1p value of the packet and assign forwarding-class priority queue to the traffic. You then configure the scheduler map and set the priority for the traffic.

By default, the SRX5K-MPC supports eight queues. In this example, you are configuring eight queues.

You apply the BA classifier to input interface and apply the scheduler map to the output interface.

Table 160 and Table 161 show forwarding class details with priority, assigned queue numbers, and allocated queue buffers used in this example.

Table 160: Forwarding Class Samples

Forwarding Class	Queue Number
BE	0
SIG	1
AF	2
Bronze-class	3
Silver-class	4
Gold-class	5
Control	6
VOIP	7

Table 161: Scheduler Samples

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer (Transmit Rate)
s-be	0	low	15
s-sig	1	low	15

Copyright © 2016, Juniper Networks, Inc.

Table 161: Scheduler Samples *(continued)*

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer (Transmit Rate)
s-af	2	medium-low	20
s-bronze	3	medium-low	20
s-silver	4	medium-high	10
s-gold	5	medium-high	10
s-nc	6	high	5
s-voip	7	high	5

Configuration

CLI Quick Configuration

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

```
set class-of-service classifiers ieee-802.1 c802 forwarding-class BE loss-priority low
    code-points 000
set class-of-service classifiers ieee-802.1 c802 forwarding-class SIG loss-priority low
    code-points 001
set class-of-service classifiers ieee-802.1 c802 forwarding-class AF loss-priority low
    code-points 010
set class-of-service classifiers ieee-802.1 c802 forwarding-class Bronze-Class loss-priority
    low code-points 011
set class-of-service classifiers ieee-802.1 c802 forwarding-class Silver-Class loss-priority
    low code-points 100
set class-of-service classifiers ieee-802.1 c802 forwarding-class Gold-Class loss-priority
    low code-points 101
set class-of-service classifiers ieee-802.1 c802 forwarding-class Central loss-priority low
    code-points 110
set class-of-service classifiers ieee-802.1 c802 forwarding-class VOIP loss-priority low
    code-points 111
set class-of-service forwarding-classes class BE queue-num 0
set class-of-service forwarding-classes class SIG queue-num 1
set class-of-service forwarding-classes class AF queue-num 2
set class-of-service forwarding-classes class Bronze-Class queue-num 3
set class-of-service forwarding-classes class Silver-Class queue-num 4
set class-of-service forwarding-classes class Gold-Class queue-num 5
set class-of-service forwarding-classes class Control queue-num 6
set class-of-service forwarding-classes class VOIP queue-num 7
set class-of-service scheduler-maps test forwarding-class BE scheduler s-be
set class-of-service scheduler-maps test forwarding-class SIG scheduler s-sig
set class-of-service scheduler-maps test forwarding-class AF scheduler s-af
set class-of-service scheduler-maps test forwarding-class Bronze-Class scheduler
    s-bronze
```


Copyright © 2016, Juniper Networks, Inc.

```
set class-of-service scheduler-maps test forwarding-class Silver-Class scheduler s-silver
set class-of-service scheduler-maps test forwarding-class Gold-Class scheduler s-gold
set class-of-service scheduler-maps test forwarding-class Control scheduler s-nc
set class-of-service scheduler-maps test forwarding-class VOIP scheduler s-voip
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class BE loss-priority low
  code-point 000
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class SIG loss-priority
  low code-point 001
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class AF loss-priority low
  code-point 010
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class Bronze-Class
  loss-priority low code-point 011
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class Silver-Class
  loss-priority low code-point 100
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class Gold-Class
  loss-priority low code-point 101
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class Control loss-priority
  low code-point 110
set class-of-service rewrite-rules ieee-802.1 rw802 forwarding-class VOIP loss-priority
  low code-point 111
set class-of-service schedulers s-be transmit-rate percent 15
set class-of-service schedulers s-be priority low
set class-of-service schedulers s-sig transmit-rate percent 15
set class-of-service schedulers s-sig priority low
set class-of-service schedulers s-af transmit-rate percent 20
set class-of-service schedulers s-af priority medium-low
set class-of-service schedulers s-bronze transmit-rate percent 20
set class-of-service schedulers s-bronze priority medium-low
set class-of-service schedulers s-silver transmit-rate percent 10
set class-of-service schedulers s-silver priority medium-high
set class-of-service schedulers s-gold transmit-rate percent 10
set class-of-service schedulers s-gold priority medium-high
set class-of-service schedulers s-nc transmit-rate percent 5
set class-of-service schedulers s-nc priority high
set class-of-service schedulers s-voip transmit-rate percent 5
set class-of-service schedulers s-voip priority high
set class-of-service interfaces reth0 unit 0 classifiers ieee-802.1 c802
set class-of-service interfaces reth0 unit 0 rewrite-rules ieee-802.1 rw802
set class-of-service interfaces reth0 shaping-rate 1g
```

Step-by-Step Procedure

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

To configure forwarding classes:

1. Configure a classifier.

```
[edit class-of-service]
user@host# set classifiers ieee-802.1 c802 forwarding-class BE loss-priority low
  code-points 000
user@host# set classifiers ieee-802.1 c802 forwarding-class SIG loss-priority low
  code-points 001
user@host# set classifiers ieee-802.1 c802 forwarding-class AF loss-priority low
  code-points 010
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set classifiers ieee-802.1 c802 forwarding-class Bronze-Class
  loss-priority low code-points 011
user@host# set classifiers ieee-802.1 c802 forwarding-class Silver-Class loss-priority
  low code-points 100
user@host# set classifiers ieee-802.1 c802 forwarding-class Gold-Class loss-priority
  low code-points 101
user@host# set classifiers ieee-802.1 c802 forwarding-class Central loss-priority
  low code-points 110
user@host# set classifiers ieee-802.1 c802 forwarding-class VOIP loss-priority low
  code-points 111
```

2. Assign best-effort traffic to queue.

```
[edit class-of-service forwarding-classes class]
user@host# BE queue-num 0
user@host# SIG queue-num 1
user@host# AF queue-num 2
user@host# Bronze-Class queue-num 3
user@host# Silver-Class queue-num 4
user@host# Gold-Class queue-num 5
user@host# Control queue-num 6
user@host# VOIP queue-num 7
```

3. Define mapping of forwarding classes to packet schedulers.

```
[edit class-of-service]
user@host# set scheduler-maps test forwarding-class BE scheduler s-be
user@host# set scheduler-maps test forwarding-class SIG scheduler s-sig
user@host# set scheduler-maps test forwarding-class AF scheduler s-af
user@host# set scheduler-maps test forwarding-class Bronze-Class scheduler
  s-bronze
user@host# set scheduler-maps test forwarding-class Silver-Class scheduler s-silver
user@host# set scheduler-maps test forwarding-class Gold-Class scheduler s-gold
user@host# set scheduler-maps test forwarding-class Control scheduler s-nc
user@host# set scheduler-maps test forwarding-class VOIP scheduler s-voip
```

4. Configure the CoS rewrite rules to map the forwarding class to the desired value for the 802.1p field.

```
[edit class-of-service]
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class BE loss-priority
  low code-point 000
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class SIG loss-priority
  low code-point 001
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class AF loss-priority
  low code-point 010
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class Bronze-Class
  loss-priority low code-point 011
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class Silver-Class
  loss-priority low code-point 100
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class Gold-Class
  loss-priority low code-point 101
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class Control loss-priority
  low code-point 110
user@host# set rewrite-rules ieee-802.1 rw802 forwarding-class VOIP loss-priority
  low code-point 111
```

5. Configure eight packet schedulers with scheduling priority and transmission rates.

Copyright © 2016, Juniper Networks, Inc.

```
[edit class-of-service]
user@host# set schedulers s-be transmit-rate percent 15
user@host# set schedulers s-be priority low
user@host# set schedulers s-sig transmit-rate percent 15
user@host# set schedulers s-sig priority low
user@host# set schedulers s-af transmit-rate percent 20
user@host# set schedulers s-af priority medium-low
user@host# set schedulers s-bronze transmit-rate percent 20
user@host# set schedulers s-bronze priority medium-low
user@host# set schedulers s-silver transmit-rate percent 10
user@host# set schedulers s-silver priority medium-high
user@host# set schedulers s-gold transmit-rate percent 10
user@host# set schedulers s-gold priority medium-high
user@host# set schedulers s-nc transmit-rate percent 5
user@host# set schedulers s-nc priority high
user@host# set schedulers s-voip transmit-rate percent 5
user@host# set schedulers s-voip priority high
```

6. Apply the classifier and rewrite rules to interfaces.

```
[edit class-of-service]
user@host# set interfaces reth0 unit 0 classifiers ieee-802.1 c802
user@host# set interfaces reth1 unit 0 rewrite-rules ieee-802.1 rw802
```

7. Apply the shaping rates to control the maximum rate of traffic transmitted on an interface.

```
[edit class-of-service]
user@host# set interfaces reth0 shaping-rate 1g
```

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

```
classifiers {
  ieee-802.1 c802 {
    forwarding-class BE {
      loss-priority low code-points 000;
    }
    forwarding-class SIG {
      loss-priority low code-points 001;
    }
    forwarding-class AF {
      loss-priority low code-points 010;
    }
    forwarding-class Bronze-Class {
      loss-priority low code-points 011;
    }
    forwarding-class Silver-Class {
      loss-priority low code-points 100;
    }
    forwarding-class Gold-Class {
      loss-priority low code-points 101;
    }
    forwarding-class Control {
      loss-priority low code-points 110;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        forwarding-class VOIP {
            loss-priority low code-points 111;
        }
    }
}
forwarding-classes {
    class BE queue-num 0;
    class SIG queue-num 1;
    class VOIP queue-num 7;
    class AF queue-num 2;
    class Bronze-Class queue-num 3;
    class Silver-Class queue-num 4;
    class Gold-Class queue-num 5;
    class Control queue-num 6;
}
interfaces {
    e1-0/0/0 {
        shaping-rate 1g;
        unit 0 {
            scheduler-map test;
        }
    }
    reth0 {
        shaping-rate 1g;
        unit 0 {
            classifiers {
                ieee-802.1 c802;
            }
            rewrite-rules {
                ieee-802.1 rw802;
            }
        }
    }
}
rewrite-rules {
    ieee-802.1 rw802 {
        forwarding-class BE {
            loss-priority low code-point 000;
        }
        forwarding-class SIG {
            loss-priority low code-point 001;
        }
        forwarding-class AF {
            loss-priority low code-point 010;
        }
        forwarding-class Bronze-Class {
            loss-priority low code-point 011;
        }
        forwarding-class Silver-Class {
            loss-priority low code-point 100;
        }
        forwarding-class Gold-Class {
            loss-priority low code-point 101;
        }
        forwarding-class Control {
```


Copyright © 2016, Juniper Networks, Inc.

```
        loss-priority low code-point 110;
      }
      forwarding-class VOIP {
        loss-priority low code-point 111;
      }
    }
  }
  scheduler-maps {
    test {
      forwarding-class BE scheduler s-be;
      forwarding-class VOIP scheduler s-voip;
      forwarding-class Gold-Class scheduler s-gold;
      forwarding-class SIG scheduler s-sig;
      forwarding-class AF scheduler s-af;
      forwarding-class Bronze-Class scheduler s-bronze;
      forwarding-class Silver-Class scheduler s-silver;
      forwarding-class Control scheduler s-nc;
    }
  }
  schedulers {
    s-be {
      transmit-rate percent 15;
      priority low;
    }
    s-nc {
      transmit-rate percent 5;
      priority high;
    }
    s-gold {
      transmit-rate percent 10;
      priority medium-high;
    }
    s-sig {
      transmit-rate percent 15;
      priority low;
    }
    s-af {
      transmit-rate percent 20;
      priority medium-low;
    }
    s-bronze {
      transmit-rate percent 20;
      priority medium-low;
    }
    s-silver {
      transmit-rate percent 10;
      priority medium-high;
    }
    s-voip {
      transmit-rate percent 5;
      priority high;
    }
  }
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Class-of-Service Configuration on page 1616
- Verifying the Number of Dedicated Queues Configured on MPC Interfaces on page 1616

Verifying Class-of-Service Configuration

Purpose Verify that CoS is configured.

Action From operational mode, enter the **show class-of-service classifier** command.

```
user@host> show class-of-service classifier type ieee-802.1

Forwarding class          ID    Queue Restricted queue  Fabric priority  Policing
 priority   SPU priority
  BE                       0       0          0             low
normal            low
  SIG                      1       1          1             low
normal            low
  AF                       2       2          2             low
normal            low
  Bronze-Class             3       3          3             low
normal            low
  Silver-Class             4       4          0             low
normal            low
  Gold-Class               5       5          1             low
normal            low
  Control                  6       6          2             low
normal            low
  VOIP                     7       7          3             low
normal            low
```

Verifying the Number of Dedicated Queues Configured on MPC Interfaces

Purpose Display the number of dedicated queue resources that are configured for the interfaces on a port.

Action From operational mode, enter the **show class-of-service interface** command.

```
user@host> show class-of-service interface reth0

Physical interface: reth0, Index: 129
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: reth0.0, Index: 71
    Object                 Name                    Type                    Index

    Classifier             dscp-ipv6-compatibility dscp-ipv6                9

    Classifier             ipprec-compatibility    ip                       13
```


Copyright © 2016, Juniper Networks, Inc.

```
Logical interface: reth1.32767, Index: 70
```

Related Documentation

- Understanding IOC Hardware Properties on page 1598
- CoS Support on the SRX5000 Module Port Concentrator Overview on page 1607

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 21

Configuration Statements and Operational Commands

- Configuration Statements on page 1621
- Operational Commands on page 1653

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 80

Configuration Statements

- Class-of-Service Configuration Statement Hierarchy on page 1622
- adaptive-shaper on page 1626
- adaptive-shapers on page 1627
- application-traffic-control on page 1628
- buffer-size (Schedulers) on page 1629
- classifiers (CoS) on page 1630
- code-points (CoS) on page 1631
- copy-outer-dscp on page 1631
- default (CoS) on page 1632
- forwarding-classes (CoS) on page 1633
- frame-relay-de (CoS Interfaces) on page 1634
- frame-relay-de (CoS Loss Priority) on page 1635
- frame-relay-de (CoS Rewrite Rule) on page 1635
- ingress-policer-overhead on page 1636
- interfaces (CoS) on page 1638
- logical-interface-policer on page 1639
- loss-priority (CoS Loss Priority) on page 1639
- loss-priority (CoS Rewrite Rules) on page 1640
- loss-priority-maps (CoS Interfaces) on page 1641
- loss-priority-maps (CoS) on page 1641
- rate-limiters on page 1642
- rewrite-rules (CoS) on page 1643
- rewrite-rules (CoS Interfaces) on page 1644
- rule-sets (CoS AppQoS) on page 1645
- scheduler-map (CoS Virtual Channels) on page 1647
- shaping-rate (CoS Adaptive Shapers) on page 1647
- shaping-rate (CoS Interfaces) on page 1648
- shaping-rate (CoS Virtual Channels) on page 1649

Copyright © 2016, Juniper Networks, Inc.

- trigger (CoS) on page 1649
- tunnel-queuing on page 1650
- virtual-channels on page 1650
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652

Class-of-Service Configuration Statement Hierarchy

Use the statements in the **class-of-service** configuration hierarchy to configure class-of-services (CoS) features.

```
class-of-service {
  adaptive-shapers adaptive-shaper-name {
    trigger becn {
      shaping-rate (absolute-rate | percent percent);
    }
  }
  application-traffic-control {
    rate-limiters rate-limiter-name {
      bandwidth-limit kbps;
      burst-size-limit bytes;
    }
    rule-sets rule-set-name {
      rule rule-name {
          match {
            application [application-name];
            application-any;
            application-group [application-group-name];
            application-known;
            application-unknown;
          }
          then {
            dscp-code-point dscp-value;
            forwarding-class class-name;
            log;
            loss-priority (high | low |medium-high | medium-low);
            rate-limit {
              loss-priority-high;
              client-to-server rate-limiter;
              server-to-client rate-limiter;
            }
          }
        }
      }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier name
    {
    forwarding-class class-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-points [alias-or-bit-string ];
      }
    }
    import (classifier-name | default);
  }
}
code-point-aliases {
  (dscp | dscp-ipv6 |exp |ieee-802.1 |ieee-802.1ad |inet-precedence) alias-name{
    dscp-bits;
    }
}
drop-profiles profile-name {
  fill-level percent {
    drop-probability number;
  }
  interpolate {
    drop-probability [number];
    fill-level [percent];
  }
}
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map next-hop-map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [lsp-regular-expression];
      next-hop [next-hop-identifier];
      non-lsp-next-hop;
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
fragmentation-maps fragmentation-map-name {
  forwarding-class forwarding-class-name {
    drop-timeout milliseconds;
      (fragment-threshold bytes |no-fragmentation) ;
      multilink-class number;
    }
  }
host-outbound-traffic {
  dscp-code-point static-dscp-code-point;
  forwarding-class class-name;
  tcp {
    raise-internet-control-priority;
  }
}
interfaces interface-name {
  input-traffic-control-profile profile-name;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
    scheduler-map scheduler-map;
    shaping-rate bps;
    unit  logical-unit-number  {
      adaptive-shaper adaptive-shaper-name;
      classifiers {
        (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
      }
      forwarding-class class-name;
      input-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      loss-priority-maps {
        frame-relay-de {
          (lpmap-name | default);
        }
      }
      output-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      rewrite-rules {
        (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad
          |inet-precedence)
      }
      scheduler-map scheduler-map-name;
        shaping-rate {
        rate;
      }
      vc-shared-scheduler;
      virtual-channel-group group-name;
    }
  }
}
loss-priority-maps {
  frame-relay-de loss-priority-map-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-points [bit-string];
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
rewrite-rules {
  (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
    rewrite-rule-name {
    forwarding-class forwarding-class-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-point alias-or-bit-string;
      }
      import (default | rewrite-rule-name);
    }
  }
}
scheduler-maps scheduler-map-name {
  forwarding-class class-name {
    scheduler scheduler-name;
    }
  }
schedulers scheduler-name {
  buffer-size {
    exact;
    (percent percent | remainder percent  | temporal microseconds) ;
  }
  drop-profile-map {
    loss-priority (any | high | low | medium-high | medium-low);
    protocol any ;
    drop-profile profile;
  }
  priority (high | low | medium-high | medium-low | strict-high);
  shaping-rate (absolute-rate | percent percent);
  transmit-rate <exact> (percent percent | rate bits | remainder percent);
}
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
traffic-control-profiles profile-name {
  delay-buffer-rate ( absolute-rate | cps cells-per-second | percent percent);
    guaranteed-rate (absolute-rate | percent percent);
    overhead-accounting (bytes bytes | cell-mode | frame-mode);
    scheduler-map scheduler-map-name;
    shaping-rate (absolute-rate | percent percent);
  }
tri-color;
virtual-channel-groups virtual-channel-group-name {
  virtual-channel-name {
    default;
    scheduler-map scheduler-map-name;
```

```
            shaping-rate (absolute-rate | percent percent);
          }
        }
        virtual-channels virtual-channel-name;
    }
```

Related Documentation

- SSL Proxy Overview on page 531
- Understanding Interfaces on page 2441

adaptive-shaper

Syntax adaptive-shaper *adaptive-shaper-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Assign an adaptive shaper to an interface.

Adaptive shapers enable bandwidth limits on Frame Relay interfaces when the device receives frames containing the backward explicit congestion notification (BECN) bit.

Options ***adaptive-shaper-name***—Name of the adaptive shaper.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- adaptive-shapers on page 1627

Copyright © 2016, Juniper Networks, Inc.

adaptive-shapers

Syntax

```
adaptive-shapers {
  adaptive-shaper-name {
    trigger type shaping-rate (percent percentage | rate);
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define trigger types and associated rates. Adaptive shapers enable bandwidth limits on Frame Relay interfaces when the device receives frames containing the backward explicit congestion notification (BECN) bit.

Options ***adaptive-shaper-name***—Name of the adaptive shaper.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface— view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- adaptive-shaper on page 1626

Copyright © 2016, Juniper Networks, Inc.

application-traffic-control

Syntax

```
application-traffic-control {
  rate-limiters rate-limiter-name{
    bandwidth-limit kbps;
      burst-size-limit bytes;
    }
  }
  rule-sets ruleset-name {
    rule rule-name {
      match {
        application application-name1;
        application-any;
        application-group application-group-name;
        application-known;
        application-unknown;
      }
      then {
        dscp-code-point dscp-value;
        forwarding-class forwarding-class-name;
        log;
        loss-priority [ high | medium-high | medium-low | low ];
        rate-limit {
          loss-priority-high;
          client-to-server rate-limiter-name;
          server-to-client rate-limiter-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 11.4.

Description Mark DSCP values for outgoing packets or apply rate limits based on the specified Layer 7 application types.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

buffer-size (Schedulers)

Syntax buffer-size (percent *percentage* | remainder | temporal *microseconds*);

Hierarchy Level [edit class-of-service schedulers *scheduler-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for ACX Series Routers.

Description Specify buffer size.

NOTE: On PTX Series Packet Transport Routers, buffer-size cannot be configured on rate-limited queues.

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

Options **percent *percentage***—Buffer size as a percentage of the total buffer.

Range: 0 through 100

remainder—Remaining buffer available.

temporal *microseconds*—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

Range: The ranges vary by platform as follows:

- For SRX Series Services Gateways: 1 through 2,000,000 microseconds.
- For vSRX instances: 1 through 32,000,000 microseconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Scheduler Buffer Size Overview on page 1489

Copyright © 2016, Juniper Networks, Inc.

classifiers (CoS)

Syntax

```
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
    forwarding-class forwarding-class-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-point alias-or-bit-string ;
    }
    import (default | user-defined;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.2

Description Configure a user-defined behavior aggregate (BA) classifier.

Options

- *classifier-name*—User-defined name for the classifier.
- import (default | *user-defined*)—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type **dscp** and you specify **import default**, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify **import mymap**, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named **mymap**.
- forwarding-class *class-name*—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
- loss-priority *level*—Specify a loss priority for this forwarding class: **high**, **low**, **medium-high**, **medium-low**.
- code-points (*alias* | *bits*)—Specify a code-point alias or the code points that map to this forwarding class.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

code-points (CoS)

Syntax code-points [*aliases*] [*6–bit-patterns*];

Hierarchy Level [edit class-of-service classifiers *type classifier-name* forwarding-class *class-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.

Options *aliases*—Name of the DSCP alias.

6–bit patterns—Value of the code-point bits, in decimal form.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

copy-outer-dscp

Syntax copy-outer-dscp;

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30.

Description Enable copying of Differentiated Services Code Point (DSCP) (outer DSCP+ECN) field from the outer IP header encrypted packet to the inner IP header plain text message on the decryption path. The benefit in enabling this feature is that after IPsec decryption, clear text packets can follow the inner CoS (DSCP+ECN) rules.

Default By default, the copy outer dscp feature is disabled.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *VPN Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

default (CoS)

Syntax default;

Hierarchy Level [edit class-of-service virtual-channel-groups *group-name virtual-channel-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the default channel. You must configure one of the virtual channels in the group to be the default. Any traffic not explicitly directed to a virtual channel is transmitted by way of this default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- scheduler-map (CoS Virtual Channels) on page 1647
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652
- virtual-channels on page 1650

Copyright © 2016, Juniper Networks, Inc.

forwarding-classes (CoS)

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- *class-name*—Display the forwarding class name assigned to the internal queue number.

  NOTE: This option is supported only on high-end SRX Series devices, including the SRX1500, SRX5400, SRX5600, and SRX5800.

  NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **policing-priority**—Layer 2 policing. One forwarding class can be configured as **premium** and others are configured as **normal**.
- **priority**—Fabric priority value:
  - **high**—Forwarding class's fabric queuing has high priority.
  - **low**—Forwarding class's fabric queuing has low priority.
- *queue-number*—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, either **high** or **low**.

  NOTE: The spu-priority option is only supported on SRX1500, and SRX5000 line devices.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppQoS on page 589

frame-relay-de (CoS Interfaces)

Syntax
frame-relay-de (*name* | default);

Hierarchy Level
[edit class-of-service interfaces *interface-name* unit *logical-unit-number* loss-priority-maps],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Assign the loss priority map or the rewrite rule to a logical interface.

Options
- **default**—Apply default loss priority map or default rewrite rule. The default loss priority map contains the following settings:

  ```
  loss-priority low code-point 0;
  loss-priority high code-point 1;
  ```

- The default rewrite rule contains the following settings:

  ```
  loss-priority low code-point 0;
  loss-priority medium-low code-point 0;
  loss-priority medium-high code-point 1;
  loss-priority high code-point 1;
  ```

- ***name***—Name of loss priority map or rewrite rule to be applied.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Defining a Custom Frame Relay Rewrite Rule on page 1468

Copyright © 2016, Juniper Networks, Inc.

frame-relay-de (CoS Loss Priority)

Syntax

```
loss-priority-maps {
  frame-relay-de loss-priority-map-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-points [bit-string];
    }
  }
```

Hierarchy Level [edit class-of-service loss-priority-maps]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define a Frame Relay discard-eligible bit loss-priority map.

Options

- ***level***—Level of loss priority to be applied based on the specified CoS values. The level can be **low**, **medium-low**, **medium-high**, or **high**.
- ***map-name***—Name of the loss-priority map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

frame-relay-de (CoS Rewrite Rule)

Syntax

```
frame-relay-de rewrite-name {
  forwarding-class class-name {
    loss-priority level code-point (0 | 1);
  }
  import (default | rewrite-name);
}
```

Hierarchy Level [edit class-of-service rewrite-rules]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define a Frame Relay discard-eligible bit rewrite rule.

Options

- ***level***—Level of loss priority on which to base the rewrite rule. The loss priority level can be **low**, **medium-low**, **medium-high**, or **high**.
- ***rewrite-name***—Name of a rewrite-rules mapping configured at the **[edit class-of-service rewrite-rules]** hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

ingress-policer-overhead

Syntax ingress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.
Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Add the configured number of bytes to the length of a packet entering the interface.

Configure a policer overhead to control the rate of traffic received on an interface. Use this feature to help prevent denial-of-service (DoS) attacks or to enforce traffic rates to conform to the service-level agreement (SLA). When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate-limiting action.

Traffic policing combines the configured policy bandwidth limits and the burst size to determine how to meter the incoming traffic. If you configure a policer overhead on an interface, Junos OS adds those bytes to the length of incoming Ethernet frames. This added overhead fills each frame closer to the burst size, allowing you to control the rate of traffic received on an interface.

You can configure the policer overhead to rate-limit queues and Layer 2 and Layer 3 policers, for standalone (SA) and high-avalability (HA) deployments. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

NOTE: vSRX supports policer overhead on Layer 3 policers only.

The policer overhead applies to all interfaces on the PIC. In the following example, Junos OS adds 10 bytes of overhead to all incoming Ethernet frames on ports ge-0/0/0 through ge-0/0/4.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 10
```

NOTE: vSRX only supports fpc 0 pic 0. When you commit the ingress-policer-overhead statement, the vSRX takes the PIC offline and then back online.

You need to craft the policer overhead size to match your network traffic. A value that is too low will have minimal impact on traffic bursts. A value that is too high will rate-limit too much of your incoming traffic.

In this example, the policer overhead of 255 bytes is configured for ge-0/0/0 through ge-0/0/4. The firewall policer is configured to discard traffic when the burst size is over

Copyright © 2016, Juniper Networks, Inc.

1500 bytes. This policer is applied to ge-0/0/0 and ge 0/0/1. Junos OS adds 255 bytes to every Ethernet frame that comes into the configured ports. If, during a burst of traffic, the combined length of incoming frames and the overhead bytes exceeds 1500 bytes, the policer starts to discard further incoming traffic.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 255
set interfaces ge-0/0/0 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/0 unit 0 family inet address 10.9.1.2/24
set interfaces ge-0/0/1 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/1 unit 0 family inet address 10.9.2.2/24
set firewall policer overhead_policer if-exceeding bandwidth-limit 32k
set firewall policer overhead_policer if-exceeding burst-size-limit 1500
set firewall policer overhead_policer then discard
```

Options

bytes—Number of bytes added to a frame entering an interface.

Range: 0–255 bytes

Default: 0

```
[edit chassis fpc 0 pic 0]
user@host# set ingress-policer-overhead 10;
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- set firewall policer

Copyright © 2016, Juniper Networks, Inc.

interfaces (CoS)

Syntax

```
interfaces interface-name {
  input-traffic-control-profile profile-name;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
    scheduler-map scheduler-map;
    shaping-rate bps;
    unit  logical-unit-number  {
      adaptive-shaper adaptive-shaper-name;
      classifiers {
        (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
      }
      forwarding-class class-name;
      input-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      loss-priority-maps {
        frame-relay-de {
          (lpmap-name | default);
        }
      }
      output-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      rewrite-rules {
        (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      }
      scheduler-map scheduler-map-name;
        shaping-rate {
        rate;
      }
      vc-shared-scheduler;
      virtual-channel-group group-name;
    }
  }
}
```

Hierarchy Level [edit class-of-service interface *interface-name* unit *number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Associate the class-of-service configuration elements with an interface.

Options interface *interface-name* unit *number*—The user-specified interface name and unit number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

logical-interface-policer

Syntax logical-interface-policer;

Hierarchy Level [edit firewall policer *policer-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a logical interface (aggregate) policer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Two-Color Policer Configuration Overview on page 1434
- Example: Configuring a Two-Color Logical Interface (Aggregate) Policer on page 1436

loss-priority (CoS Loss Priority)

Syntax loss-priority *level* code-points[*values*];

Hierarchy Level [edit class-of-service loss-priority-maps frame-relay-de *map-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Map CoS values to a loss priority.

Options ***level*** can be one of the following:

- **high**—Packet has high loss priority.
- **medium-high**—Packet has medium-high loss priority.
- **medium-low**—Packet has medium-low loss priority.
- **low**—Packet has low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

loss-priority (CoS Rewrite Rules)

Syntax loss-priority *level*;

Hierarchy Level [edit class-of-service rewrite-rules *type rewrite-name* forwarding-class *class-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and packet loss priority (PLP). The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern.

Options ***level*** can be one of the following:

- **high**—The rewrite rule applies to packets with high loss priority.
- **low**—The rewrite rule applies to packets with low loss priority.
- **medium-high**—The rewrite rule applies to packets with medium-high loss priority.
- **medium-low**—The rewrite rule applies to packets with medium-low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

 Copyright © 2016, Juniper Networks, Inc.

loss-priority-maps (CoS Interfaces)

Syntax

```
loss-priority-maps {
  frame-relay-de (map-name | default);
}
```

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Assign the loss priority map to a logical interface.

Options

- **default**—Apply default loss priority map. The default map contains the following:

  ```
  loss-priority low code-point 0;
  loss-priority high code-point 1;
  ```

- ***map-name***—Name of loss priority map to be applied.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

loss-priority-maps (CoS)

Syntax

```
loss-priority-maps {
  frame-relay-de loss-priority-map-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-points [bit-string];
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.2.

Description Map the loss priority of incoming packets based on CoS values.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

rate-limiters

Syntax

```
rate-limiters {
  rate-limiter-name {
    bandwidth-limit value-in-kbps;
    burst-size-limit value-in-bytes;
  }
}
```

Hierarchy Level [edit class-of-service application-traffic-control]

Release Information Statement introduced in Junos OS Release 11.4.

Description Share the available bandwidth and burst size of a device's PICs by defining rate limiter profiles and applying them in AppQoS rules.

Options

- ***rate-limiter-name***—Name of the rate limiter. It is applied in AppQoS rules to share device resources based on quality of service requirements.

  The combination of rate limiting parameters, namely bandwidth- limit and burst-size-limit rate limit, make up the rate limiter profile. A maximum of 16 profiles are allowed per device. The same profile can be used by multiple rate limiters. For example, a profile with a bandwidth-limit of 200 Kbps and a burst-limit of 130,000 bytes, could be used in several rate limiters.

  A maximum of 1000 rate limiters can be created. Rate limiters are defined for the device, and are assigned in rules in a rule set. A single rate limiter can be used multiple times within the same rule set. However, the rate limiter cannot be used in another rule set.

- **bandwidth-limit *value-in-Kbps***—Maximum number of kilobits to be transmitted per second for this rate limiter. Up to 2 GB of bandwidth can be provisioned among multiple rate limiters to share the resource proportionally.

- **burst-size-limit *value-in-bytes***—Maximum number of bytes to be transferred in a single burst or time-slice. This limit ensures that a high-priority transmission does not keep a lower priority transmission from transmitting.

NOTE: The number of bandwidth-limit and burst-size-limit combinations cannot exceed 16.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589

Copyright © 2016, Juniper Networks, Inc.

rewrite-rules (CoS)

Syntax

```
rewrite-rules {
  (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
    rewrite-rule-name {
    forwarding-class forwarding-class-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-point alias-or-bit-string ;
      }
      import (default | rewrite-rule-name);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. **ieee-802.1ad** option introduced in Junos OS Release 9.2.

Description Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

Options

- ***rewrite-name***—Name of a **rewrite-rules** mapping.
- ***type***—Traffic type.

  Values: dscp, **dscp-ipv6**, **exp**, **frame-relay-de**.**ieee-802.1**, **ieee-802.1ad**, **inet-precedence**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- rewrite-rules (CoS Interfaces) on page 1644

Copyright © 2016, Juniper Networks, Inc.

rewrite-rules (CoS Interfaces)

Syntax

```
rewrite-rules {
    (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
}
```

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Associate a rewrite-rules configuration or default mapping with a specific interface.

Options

- ***rewrite-name***—Name of a **rewrite-rules** mapping configured at the **[edit class-of-service rewrite-rules]** hierarchy level.
- **default**—The default mapping.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- rewrite-rules (CoS) on page 1643

Copyright © 2016, Juniper Networks, Inc.

rule-sets (CoS AppQoS)

Syntax

```
rule-sets {
  rule-set-name {
    rule rule-name {
      match {
        application application-name;
        application-any;
        application-group application-group-name;
        application-known;
        application-unknown;
      }
      then {
        dscp-code-point dscp-value ;
        forwarding-class forwarding-class-name;
        log;
        loss-priority [ high | medium-high | medium-low | low ];
        rate-limit {
          loss-priority-high;
          client-to-server rate-limiter-name;
          server-to-client rate-limiter-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service application-traffic-control]

Release Information Statement introduced in Junos OS Release 11.4.

Description Defines AppQoS rule sets and the rules that establish priorities based on quality of service requirements for the associated applications. AppQoS rules can be included in policy statements to implement application-aware quality of service control.

Options

- ***rule-set-name***—Name used to refer to a collection of AppQoS rules.
- **rule *rule-name***—Name applied to the match criteria and resulting actions that control the quality of service provided to any matching applications.
- **application *application-name***—Name of the application to be used as match criteria for the rule.
- **application-any** —Any application encountering this rule. Note that when you use this specification, all application matching ends. Any application rule following this one will never be encountered.
- **application-group *application-group-name***—Group of applications to be used as match criteria for the rule. Both applications and application groups can be match criteria for a single rule.
- **application-known**—Match criteria specifying any session that is identified, but its corresponding application is not specified.

Copyright © 2016, Juniper Networks, Inc.

- **application-unknown**—Match criteria specifying any session that is not identified.
- **forwarding-class *forwarding-class-name***—The AppQoS class with which matching applications will be marked. This field identifies the rewriter that has marked the DSCP value . Therefore, the AppQoS forwarding class must be different from those used by IDP or firewall filters. With this class specified, firewall filter class will not overwrite the existing DSCP value.
- **dscp-code-point**—DSCP alias or bit map with which matching applications will be marked to establish the output queue. This value can be marked by rewriters from IDP, AppQoS, or a firewall filter. The forwarding-class value identifies which rewriter has re-marked the packet with the current DSCP value. If a packet triggers all three rewriters, IDP takes precedence over AppQoS, which takes precedence over a firewall filter.
- **loss-priority**—Loss priority with which matching applications will be marked. This value is used to determine the likelihood that a packet would be dropped when encountering congestion. A high loss priority means that there is an 80% chance of packet loss in congestion. Possible values are high, medium-high, medium-low and low.
- **rate-limit**—Rate limiters to be associated with client-to-server and with server-to-client traffic for this application. The rate limiter profile defines maximum speed and volume limits for matching applications.
- **log**—AppQoS event logging.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppQoS on page 589

Copyright © 2016, Juniper Networks, Inc.

scheduler-map (CoS Virtual Channels)

Syntax scheduler-map *map-name*;

Hierarchy Level [edit class-of-service virtual-channel-groups *group-name virtual-channel-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Apply a scheduler map to this virtual channel.

Options *map-name*—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652
- virtual-channels on page 1650

shaping-rate (CoS Adaptive Shapers)

Syntax shaping-rate (percent *percentage* | *rate*);

Hierarchy Level [edit class-of-service adaptive-shapers *adaptive-shaper-name* trigger *type*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the list of trigger types and associated rates.

Options

- **percent** *percentage*—Shaping rate as a percentage of the available interface bandwidth.

  Range: 0 through 100 percent

- *rate*—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

  Range: 3200 through 32,000,000,000 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- trigger (CoS) on page 1649

Copyright © 2016, Juniper Networks, Inc.

shaping-rate (CoS Interfaces)

Syntax shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

Logical and physical interface traffic shaping is mutually exclusive. This means you can include the **shaping-rate** statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level or the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, but not both.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

Default If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options *rate*—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: For logical interfaces, 1000 through 32,000,000,000 bps.

For physical interfaces, 1000 through 160,000,000,000 bps.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

shaping-rate (CoS Virtual Channels)

Syntax shaping-rate (percent *percentage* | *rate*);

Hierarchy Level [edit class-of-service virtual-channel-groups *group-name virtual-channel-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the shaping rates to be associated with the virtual channel.

Options

- **percent** ***percentage***—Shaping rate as a percentage of the available interface bandwidth.

  Range: 0 through 100 percent

- ***rate***—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

  Range: 3200 through 32,000,000,000 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- scheduler-map (CoS Virtual Channels) on page 1647
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652
- virtual-channels on page 1650

trigger (CoS)

Syntax trigger *type* shaping-rate (percent *percentage* | *rate*);

Hierarchy Level [edit class-of-service adaptive-shapers *adaptive-shaper-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a trigger type and its associated rate.

Options ***type***—The type of trigger. Currently, the trigger type can be **becn** only.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- shaping-rate (CoS Adaptive Shapers) on page 1647

Copyright © 2016, Juniper Networks, Inc.

tunnel-queuing

Syntax tunnel-queuing;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement modified in Release 9.0 of Junos OS.

Description Enable class-of-service (CoS) queuing for generic routing encapsulation (GRE) and IP-IP tunnels.

NOTE: The tunnel-queuing option is not supported in chassis cluster mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

virtual-channels

Syntax virtual-channels *virtual-channel-name*;

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a list of virtual channels.

Each virtual channel has eight transmission queues.

Options ***virtual-channel-name***—Name of the virtual channel.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- scheduler-map (CoS Virtual Channels) on page 1647
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652

Copyright © 2016, Juniper Networks, Inc.

virtual-channel-group (CoS Interfaces)

Syntax virtual-channel-group *group-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Assign a virtual channel group to a logical interface.

If you apply a virtual channel group to multiple logical interfaces, separate queues are created on each of the interfaces. The same virtual channel names are used on all the interfaces. You can specify the scheduler and shaping rates in the virtual channels in percentages so that you can apply the same virtual channel group to logical interfaces with different available bandwidths.

Options *group-name*—Name of the virtual channel group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- scheduler-map (CoS Virtual Channels) on page 1647
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channels on page 1650
- virtual-channel-groups on page 1652

Copyright © 2016, Juniper Networks, Inc.

virtual-channel-groups

Syntax

```
virtual-channel-groups virtual-channel-group-name{
  virtual-channel-name {
    scheduler-map map-name;
    shaping-rate (percent percent | rate);
    default;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5.

Description Associate a virtual channel with a scheduler map and a shaping rate. Virtual channels and virtual channel groups enable you to direct traffic into a virtual channel and apply bandwidth limits to the channel.

Options *group-name*—Name of the virtual channel group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- scheduler-map (CoS Virtual Channels) on page 1647
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channels on page 1650

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 81

Operational Commands

- show class-of-service application-traffic-control counter
- show class-of-service application-traffic-control statistics rate-limiter
- show class-of-service application-traffic-control statistics rule
- show class-of-service forwarding-class

Copyright © 2016, Juniper Networks, Inc.

show class-of-service application-traffic-control counter

Syntax show class-of-service application-traffic-control counter

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS DSCP marking and honoring statistics based on Layer 7 application classifiers.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control counter on page 1654

Output Fields Table 41 lists the output fields for the **show class-of-service application-traffic-control counter** command. Output fields are listed in the approximate order in which they appear.

Table 162: show class-of-service application-traffic-control counter Output Fields

Field Name	Field Description
pic	PIC number of the accumulated statistics. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Sessions processed	The number of sessions where the class of service was checked.
Sessions marked	The number of sessions marked based on application-aware DSCP marking.
Sessions honored	The number of sessions honored based on application-aware traffic honoring.
Sessions rate limited	The number of sessions that have been rate limited.
Client-to-server flows rate limited	The number of client-to-server flows that have been rate limited.
Server-to-client flows rate limited	The number of server-to-client flows that have been rate limited.

Sample Output

show class-of-service application-traffic-control counter

```
user@host> show class-of-service application-traffic-control counter
pic: 2/1
  Counter type                                  Value
  Sessions processed                            300
  Sessions marked                               200
  Sessions honored                              0
  Sessions rate limited                         100
  Client-to-server flows rate limited           100
  Server-to-client flows rate limited           70
```


Copyright © 2016, Juniper Networks, Inc.

```
pic: 2/0
  Counter type                                    Value
  Sessions processed                              400
  Sessions marked                                 300
  Sessions honored                                0
  Sessions rate limited                           200
  Client-to-server flows rate limited             200
  Server-to-client flows rate limited             100
```

show class-of-service application-traffic-control statistics rate-limiter

Syntax show class-of-service application-traffic-control statistics rate-limiter

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS real-time run information about application rate limiting of current or recent sessions.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control statistics rate-limiter on page 1656

Output Fields Table 42 lists the output fields for the **show class-of-service application-traffic-control statistics rate-limiter** command. Output fields are listed in the approximate order in which they appear.

Table 163: show class-of-service application-traffic-control statistics rate-limiter Output Fields

Field Name	Field Description
pic	PIC number. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Ruleset	The rule set applied on the session.
Application	The application match for applying the rule set.
Client-to-server	The rate limiter applied from client to server.
Rate(kbps)	The rate in the client-to-server direction
Server-to-client	The rate limiter applied from server to client.
Rate(kbps)	The rate in the server-to-client direction.

Sample Output

show class-of-service application-traffic-control statistics rate-limiter

```
user@host> show class-of-service application-traffic-control statistics rate-limiter
pic: 2/1
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client
Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000     my-http-s2c-rl
20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000     my-http-s2c-rl-2
30000000
```


Copyright © 2016, Juniper Networks, Inc.

```
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl
50000
  ...

pic: 2/0
  Ruleset       Application  Client-to-server   Rate(kbps)   Server-to-client
Rate(kbps)
  my-ruleset-1  HTTP         my-http-c2s-rl     10000000    my-http-s2c-rl
20000000
  my-ruleset-2  HTTP         my-http-c2s-rl-2   20000000    my-http-s2c-rl-2
30000000
  my-ruleset-2  FTP          my-ftp-c2s-rl      50000       my-ftp-s2c-rl
50000
```

show class-of-service application-traffic-control statistics rule

Syntax show class-of-service application-traffic-control statistics rule

Release Information Command introduced in Junos OS Release 11.4.

Description Display AppQoS counters identifying rule hits.

Required Privilege Level view

Related Documentation
- Example: Configuring AppQoS on page 589

List of Sample Output show class-of-service application-traffic-control statistics rule on page 1658

Output Fields Table 43 lists the output fields for the **show class-of-service application-traffic-control statistics rule** command. Output fields are listed in the approximate order in which they appear.

Table 164: show class-of-service application-traffic-control statistics rule Output Fields

Field Name	Field Description
pic	PIC number where the rule is applied. NOTE: The PIC number is always displayed as **0** for branch SRX Series devices.
Ruleset	The rule set containing the rule.
Rule	The rule to which the statistic applies.
Hits	The number of times a match for the rule was encountered.

Sample Output

show class-of-service application-traffic-control statistics rule

```
user@host> show class-of-service application-traffic-control statistics rule
pic: 2/0
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       100
  my-ruleset-1     httpp-rule     100
  my-ruleset-2     telnet-rule    300
  my-ruleset-2     smtp-rule      300
  ...

pic: 2/1
  Ruleset          Rule           Hits
  my-ruleset-1     ftp-rule       200
  my-ruleset-1     httpp-rule     300
  my-ruleset-2     telnet-rule    400
  my-ruleset-2     smtp-rule      500
```


Copyright © 2016, Juniper Networks, Inc.

show class-of-service forwarding-class

Syntax show class-of-service forwarding-class

Release Information Command introduced before Junos OS Release 12.1.

Description Display mapping of forwarding class names to queues.

Required Privilege Level view

Related Documentation
- Forwarding Classes Overview on page 1453

List of Sample Output show class-of-service forwarding-class on page 1659

Output Fields Table 165 lists the output fields for the **show class-of-service forwarding-class** command. Output fields are listed in the approximate order in which they appear.

Table 165: show class-of-service forwarding-class Output Fields

Field Name	Field Description
Forwarding class	Forwarding class name.
ID	ID number assigned to the forwarding class.
Queue	Queue number.
Restricted queue	Restricted queue number.
Fabric priority	Fabric priority, either low or high.
Policing priority	Layer 2 policing, either premium or normal.
SPU priority	Services Processing Unit (SPU) priority queue, either high or low.

Sample Output

show class-of-service forwarding-class

```
user@host> show class-of-service forwarding-class
Forwarding class        ID  Queue  Restricted queue  Fabric priority  Policing
priority  SPU priority
  best-effort            0   0          0             low              normal
          low
  expedited-forwarding   1   1          1             low              normal
          high
  assured-forwarding     2   2          2             low              normal
          low
  network-control        3   3          3             low              normal
          low
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Flow-Based and Packet-Based Processing Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 22

Overview

- Introduction to Processing on Security Devices on page 1665

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 82

Introduction to Processing on Security Devices

- Juniper Networks Devices Processing Overview on page 1665
- Understanding SRX Series Services Gateways Central Point Architecture on page 1669
- Understanding Enhancements to Central Point Architecture for the SRX5000 Line on page 1670
- SRX5000 Line Devices Processing Overview on page 1671
- Understanding Central Point Session Limit Performance Enhancements on page 1682
- SRX320 Services Gateway Processing Overview on page 1683
- Understanding Central Point Architecture Flow Support for GTP and SCTP on page 1685

Juniper Networks Devices Processing Overview

Junos OS for security devices integrates the world-class network security and routing capabilities of Juniper Networks. Junos OS includes a wide range of packet-based filtering, class-of-service (CoS) classifiers, and traffic-shaping features as well as a rich, extensive set of flow-based security features including policies, screens, network address translation (NAT), and other flow-based services.

Traffic that enters and exits a security device is processed according to features you configure, such as packet filters, security policies, and screens. For example, the software can determine:

- Whether the packet is allowed into the device
- Which firewall screens to apply to the packet
- The route the packet takes to reach its destination
- Which CoS to apply to the packet, if any
- Whether to apply NAT to translate the packet's IP address
- Whether the packet requires an Application Layer Gateway (ALG)

Copyright © 2016, Juniper Networks, Inc.

Packets that enter and exit a device undergo both packet-based and flow-based processing:

- Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that were established for the first packet of the packet stream, which is referred to as a flow.

  For the distributed processing architecture of the services gateway, all flow-based processing occurs on the SPU and sampling is multi-thread aware. Packet sequencing is maintained for the sampled packets.

- Packet-based, or stateless, packet processing treats packets discretely. Each packet is assessed individually for treatment.

  For the distributed processing architecture of the services gateway, some packet-based processing, such as traffic shaping, occurs on the NPU. Some packet-based processing, such as application of classifiers to a packet, occurs on the SPU.

This topic includes the following sections:

- Understanding Flow-Based Processing on page 1666
- Understanding Packet-Based Processing on page 1667

Understanding Flow-Based Processing

A packet undergoes flow-based processing after packet-based filters and some screens have been applied to it. All flow-based processing for a single flow occurs on a single System Processing Unit (SPU). An SPU processes the packets of a flow according to the security features and other services configured for the session.

Figure 83 shows a conceptual view of how flow-based traffic processing occurs on services gateway.

Figure 83: Traffic Flow for Flow-Based Processing

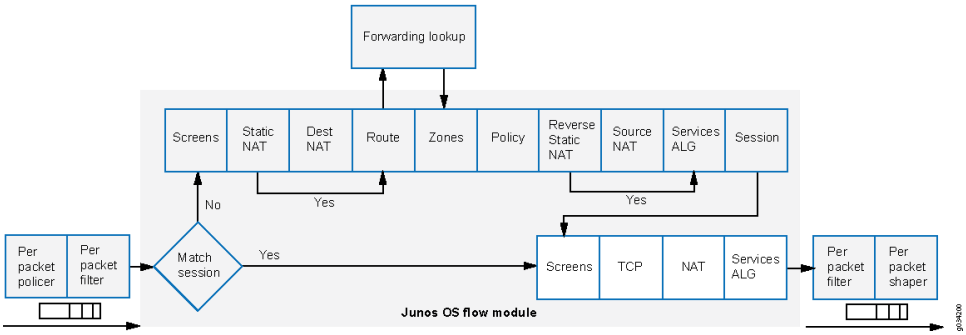

A flow is a stream of related packets that meet the same matching criteria and share the same characteristics. Junos OS treats packets belonging to the same flow in the same manner.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, if it requires an Application Layer Gateway (ALG), if NAT is applied to

Copyright © 2016, Juniper Networks, Inc.

translate the packet's source and/or destination IP address—are assessed for the first packet of a flow.

To determine if a flow exists for a packet, the NPU attempts to match the packet's information to that of an existing session based on the following match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique session token number for a given zone and virtual router

Zones and Policies

The security policy to be used for the first packet of a flow is cached in a flow table for use with the same flow and closely related flows. Security policies are associated with zones. A zone is a collection of interfaces that define a security boundary. A packet's incoming zone, as determined by the interface through which it arrived, and its outgoing zone, as determined by the forwarding lookup, together determine which policy is used for packets of the flow.

Flows and Sessions

Flow-based packet processing, which is stateful, requires the creation of sessions. A session is created for the first packet of a flow for the following purposes:

- To store most of the security measures to be applied to the packets of the flow.
- To cache information about the state of the flow.

  For example, logging and counting information for a flow is cached in its session. (Some stateful firewall screens rely on threshold values that pertain to individual sessions or across all sessions.)
- To allocate required resources for the flow for features such as NAT.
- To provide a framework for features such as ALGs and firewall features.

Most packet processing occurs in the context of a flow, including:

- Management of policies, NAT, zones, and most screens.
- Management of ALGs and authentication.

Understanding Packet-Based Processing

A packet undergoes packet-based processing when it is removed from the queue from its input interface and before it is added to the queue on its output interface.

Packet-based processing applies stateless firewall filters, CoS features, and some screens to discrete packets.

Copyright © 2016, Juniper Networks, Inc.

- When a packet arrives at an interface, sanity checks, packet-based filters, some CoS features, and some screens are applied to it.
- Before a packet leaves the device, any packet-based filters, some CoS features, and some screens associated with the interface are applied to the packet.

Filters and CoS features are typically associated with one or more interfaces to influence which packets are allowed to transit the system and to apply special actions to packets as necessary.

The following topics describe the kinds of packet-based features that you can configure and apply to transit traffic.

Configuring Packet-Based Processing

On SRX Series devices, the default mode for processing traffic is flow mode. To configure an SRX Series device as a border router, you must change the mode from flow-based processing to packet-based processing. Use the **set security forwarding-options family mpls mode packet-based** statement to configure the SRX device to packet mode. You must reboot the device for the configuration to take effect.

NOTE: Packet-based processing is supported only on the following SRX Series devices: SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500.

Stateless Firewall Filters

Also referred to as access control lists (ACLs), stateless firewall filters control access and limit traffic rates. They statically evaluate the contents of packets transiting the device from a source to a destination, or packets originating from or destined for the Routing Engine. A stateless firewall filter evaluates every packet, including fragmented packets.

You can apply a stateless firewall filter to an input or output interface, or to both. A filter contains one or more terms, and each term consists of two components—match conditions and actions. By default, a packet that does not match a firewall filter is discarded.

You can plan and design stateless firewall filters to be used for various purposes—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates. Stateless firewall filters are executed on the SPU.

Class-of-Service Features

CoS features allow you to classify and shape traffic. CoS features are executed on the SPU.

- Behavior aggregate (BA) classifiers—These classifiers operate on packets as they enter the device. Using behavior aggregate classifiers, the device aggregates different types of traffic into a single forwarding class to receive the same forwarding treatment. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Service (DiffServ) value.

Copyright © 2016, Juniper Networks, Inc.

- Traffic shaping—You can shape traffic by assigning service levels with different delay, jitter, and packet loss characteristics to particular applications served by specific traffic flows. Traffic shaping is especially useful for real-time applications, such as voice and video transmission.

Screens

Some screens, such as denial-of-service (DoS) screens, are applied to a packet outside the flow process. They are executed on the Network Processing Unit (NPU).

For details on specific stateless CoS features, see *Class of Service Feature Guide for Security Devices*.

Related Documentation

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- SRX5000 Line Devices Processing Overview on page 1671
- SRX320 Services Gateway Processing Overview on page 1683
- **forwarding-options (Security) on page 3971**

Understanding SRX Series Services Gateways Central Point Architecture

The central point in the architecture has two basic flow functionalities: load balancing and traffic identification (global session matching). The central point forwards a packet to its Services Processing Unit (SPU) upon session matching, or distributes traffic to an SPU for security processing if the packet does not match any existing session.

On some SRX Series devices, an entire SPU cannot be dedicated for central point functionality, but a certain percentage of the SPU is automatically allocated for central point functionality and the rest is allocated for normal flow processing. When an SPU performs the function of central point as well as normal flow processing, it is said to be in combination, or *combo,* mode.

The percentage of SPU dedicated to the central point functionality depends on the number of SPUs in the device. Based on the number of SPUs, there are three modes available on the SRX Series devices— small central point, medium central point, and large central point.

In small central point mode, a small percentage of an SPU is dedicated to central point functionality and the rest is dedicated to the normal flow processing. In medium central point mode, an SPU is almost equally shared for central point functionality and normal flow processing. In large central point mode, an entire SPU is dedicated to central point functionality. In combo mode, the central point and SPU share the same load-balancing thread (LBT) and packet-ordering thread (POT) infrastructure.

This topic includes the following sections:

- Load Distribution in Combo Mode on page 1670
- Sharing Processing Power and Memory in Combo Mode on page 1670

Copyright © 2016, Juniper Networks, Inc.

Load Distribution in Combo Mode

The central point maintains SPU mapping table (for load distribution) that lists live SPUs with the logic SPU IDs mapped to the physical Trivial Network Protocol (TNP) addresses mapping. In combo mode, the SPU that hosts the central point is included in the table. The load distribution algorithm is adjusted based on session capacity and processing power to avoid overloading of sessions.

Sharing Processing Power and Memory in Combo Mode

The CPU processing power in a combo-mode SPU is shared based on the platform and the number of SPUs in the system. Similarly, the CPU memory is also shared between the central point and SPU.

An SPU has multiple cores (CPUs) for networking processing. In "small" SPU combo mode, CPU functionality takes a small portion of the cores, whereas "medium" SPU combo mode requires a larger portion of cores. The processing power for central point functionalities and flow processing is shared, based on the number of Services Processing Cards (SPC), as shown in Table 166.

Table 166: Combo Mode Processing

SRX Series device	Central point mode with 1 SPC	Central point mode with 2 or More than 2 SPCs
SRX5600	Medium	Large
SRX5800	Medium	Large

NOTE: The combo mode processing only exists with SPCI on SRX5000 line devices.

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Understanding Session Characteristics for SRX Series Services Gateways on page 1689

Understanding Enhancements to Central Point Architecture for the SRX5000 Line

Previously, for the SRX5000 line of services gateways, the central point was a bottleneck in device performance and scaling. When more Services Processing Cards (SPCs) were integrated into the system, the overall processing power increased linearly, but the system connections per second (cps) remained constant and could not be improved because of the single centralized point in the system. This severely impacted the overall system utilizations in both capacity and cps.

Beginning with Junos OS Release 15.1X49-D30, on SRX5000 line devices, the central point architecture is enhanced to handle higher connections per second (cps). The new

Copyright © 2016, Juniper Networks, Inc.

central point architecture prevents data packets from going through the central point by offloading session management functionalities to the Services Processing Unit (SPU). Therefore, data packets are directly forwarded from the network processing unit to the SPU instead of going through the central point.

The central point architecture is divided into two modules, the application central point and the distributed central point. The is responsible for global resource management and loading balancing, while the distributed central point is responsible for traffic identification (global session matching). The application central point functionality runs on the dedicated central point SPU, while the distributed central point functionality is distributed to the rest of the SPUs. Now the central point sessions are no longer on the dedicated central point SPU, but with distributed central point on other flow SPUs.

NOTE: The central point for SRX5000 line refers to the application central point, or the distributed central point or both, with respect to global resource management and load balancing, it refers to the application central point, whereas with respect to traffic identification and session management, it refers to the distributed central point (sometimes referred to the SPU as well).

NOTE: The SNMP log and SNMP trap were generated by the central point with rate limit. Now, the SNMP log and SNMP trap are generated by the SPU or central point. As there is more than one SPU, the number of SNMP log and traps generated are more. To verify the number of connections per second (CPS) on the device run SNMP MIB walk nxJsNodeSessionCreationPerSecond command. The SNMP polling mechanism calculates the CPS value based on the average number of CPS in the past 96 seconds. So, if the CPS is not constant, the number of CPS reported is inaccurate.

Related Documentation

- SRX5000 Line Devices Processing Overview on page 1671

SRX5000 Line Devices Processing Overview

Junos OS is a distributed, parallel processing, high-throughput and high-performance system. The distributed parallel processing architecture of the SRX5000 line of services gateways includes multiple processors to manage sessions and run security and other services processing. This architecture provides greater flexibility and allows for high throughput and fast performance.

The SRX5000 line devices include I/O cards (IOCs) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. One or more Services Processing Units (SPUs) run on an SPC.

These processing units have different responsibilities. All flow-based services for a packet are executed on a single SPU. Otherwise, however, the responsibilities of these NPUs are

Copyright © 2016, Juniper Networks, Inc.

not clearly delineated in regard to the other kind of services that run on them. (For details on flow-based processing, see "Juniper Networks Devices Processing Overview" on page 1665.)

For example:

- An NPU processes packets discretely. It performs sanity checks and applies some screens that are configured for the interface, such as denial-of-service (DoS) screens, to the packet.
- An SPU manages the session for the packet flow and applies security features and other services to the packet. It also applies packet-based stateless firewall filters, classifiers, and traffic shapers to the packet.
- An NPU forwards a packet to the SPU using the hash algorithm. However, for some applications, like ALG, the system will need to query the application central point to determine on which SPU the packet should be processed.

These discrete, cooperating parts of the system, including the central point, each store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

This architecture allows the device to distribute processing of all sessions across multiple SPUs. It also allows an NPU to determine if a session exists for a packet, to check the packet, and to apply screens to the packet. How a packet is handled depends on whether it is the first packet in a flow.

The following sections describe the processing architecture using SRX5600 and SRX5800 devices as an example:

- Understanding First-Packet Processing on page 1672
- Understanding Fast-Path Processing on page 1673
- Understanding the Data Path for Unicast Sessions on page 1674
- Understanding Services Processing Units on page 1681
- Understanding Scheduler Characteristics on page 1681
- Understanding Network Processor Bundling on page 1681

Understanding First-Packet Processing

If the packet matches an existing flow, processing for the packet is assessed in the context of its flow state. The SPU maintains the state for each session, and the settings are then applied to the rest of the packets in the flow. If the packet does not match an existing flow, it is used to create a flow state and a session is allocated for it.

Figure 84 illustrates the path the first packet in a flow takes as it enters the device—the NPU determines that no session exists for the packet, and the NPU sends the packet to the distributed central point to set up a distributed central point session. The distributed central point then sends a message to the application central point to select the SPU to set up a session for the packet and to process the packet. The distributed central point then sends the packet to that SPU. The SPU processes the packet and sends it to the

Copyright © 2016, Juniper Networks, Inc.

NPU for transmission from the device. (This high-level description does not address application of features to a packet.)

Figure 84: First-Packet Processing

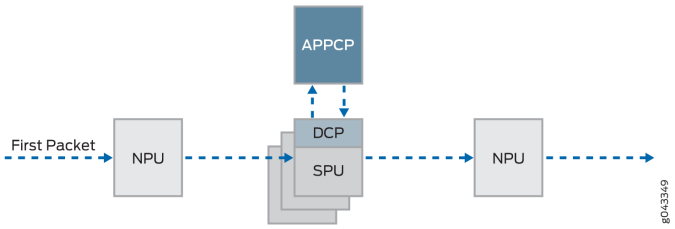

For details on session creation for the first packet in a flow, see "Understanding Session Creation: First-Packet Processing" on page 1675.

After the first packet in a flow has traversed the system and a session has been established for it, it undergoes fast-path processing.

Subsequent packets in the flow also undergo fast-path processing; in this case, after each packet enters the session and the NPU finds a match for it in its session table, the NPU forwards the packet to the SPU that manages its session.

Figure 85 illustrates fast-path processing. This is the path a packet takes when a flow has already been established for its related packets. (It is also the path that the first packet in a flow takes after the session for the flow that the packet initiated has been set up.) After the packet enters the device, the NPU finds a match for the packet in its session table, and it forwards the packet to the SPU that manages the packet's session. Note that the packet bypasses interaction with the central point.

Figure 85: Fast-Path Processing

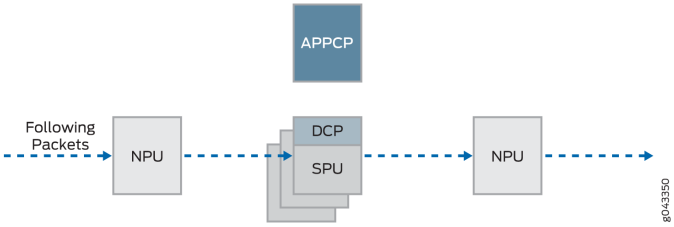

The following section explains how a session is created and the process a packet undergoes as it transits the device.

Understanding Fast-Path Processing

Here is an overview of the main components involved in setting up a session for a packet and processing packets both discretely and as part of a flow as they transit the SRX5600 and SRX5800 devices:

Copyright © 2016, Juniper Networks, Inc.

- Network Processing Units (NPUs)—NPUs reside on IOCs. They handle packet sanity checking and application of some screens. NPUs maintain session tables that they use to determine if a session exists for an incoming packet or for reverse traffic.

  The NPU session table contains an entry for a session if the session is established on an SPU for a packet that had previously entered the device via the interface and was processed by this NPU. The SPU installs the session in the NPU table when it creates the session.

  An NPU determines if a session exists for a packet by checking the packet information against its session table. If the packet matches an existing session, the NPU sends the packet and the metadata for it to the SPU. If there is no session, the NPUs sends the packet to one SPU which is calculated using the hash algorithm.

- Services Processing Units (SPUs)—The main processors of the SRX5600 and SRX5800 devices reside on SPCs. SPUs establish and manage traffic flows and perform most of the packet processing on a packet as it transits the device. Each SPU maintains a hash table for fast session lookup. The SPU applies stateless firewall filters, classifiers, and traffic shapers to traffic. An SPU performs all flow-based processing for a packet and most packet-based processing. Each multicore SPU processes packets independently with minimum interaction among SPUs on the same or different SPC. All packets that belong to the same flow are processed by the same SPU.

  The SPU maintains a session table with entries for all sessions that it established and whose packets it processes. When an SPU receives a packet from an NPU, it checks its session table to ensure that the packet belongs to it. It also checks its session table when it receives a packet from the distributed central point and sends a message to establish a session for that packet to verify that there is not an existing session for the packet.

- Central point—The central point architecture is divided into two modules, the application central point and the distributed central point. The application central point is responsible for global resource management and loading balancing, while the distributed central point is responsible for traffic identification (global session matching). The application central point functionality runs on the dedicated central point SPU, while the distributed central point functionality is distributed to the rest of the SPUs. Now the central point sessions are no longer on the dedicated central point SPU, but with the distributed central point on other flow SPUs.

- Routing Engine—The Routing Engine runs the control plane.

Understanding the Data Path for Unicast Sessions

This section describes the process of establishing a session for packets belonging to a flow that transits the device.

To illustrate session establishment and the packet “walk” including the points at which services are applied to the packets in a flow, this example uses the simple case of a unicast session.

This packet “walk” brings together the packet-based processing and flow-based processing that Junos OS performs on the packet.

Copyright © 2016, Juniper Networks, Inc.

Session Lookup and Packet-Match Criteria

To determine if a packet belongs to an existing flow, the device attempts to match the packet's information to that of an existing session based on the following six match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique token from a given zone and virtual router

Understanding Session Creation: First-Packet Processing

This section explains how a session is set up to process the packets composing a flow. To illustrate the process, this section uses an example with a source "a" and a destination "b". The direction from source to destination for the packets of the flow is referred to as (a ->b). The direction from destination to source is referred to as (b->a).

Step 1. A Packet Arrives at an Interface on the Device And the NPU Processes It.

This section describes how a packet is handled when it arrives at an SRX Series device ingress IOC.

1. The packet arrives at the device's IOC and is processed by the NPU on the IOC.
2. The NPU performs basic sanity checks on the packet and applies some screens configured for the interface to the packet.
3. The NPU checks its session table for an existing session for the packet. (It checks the packet's tuple against those of packets for existing sessions in its session table.)
   a. If no existing session is found, the NPU forwards the packet to the hash SPU.
   b. If a session match is found, the session has already been created on an SPU that was assigned to it, so the NPU forwards the packet to the SPU for processing along with the session ID. (See "Understanding Fast-Path Processing" on page 1678.)

Example: Packet (a ->b) arrives at NPU1. NPU1 performs sanity checks and applies DoS screens to the packet. NPU1 checks its session table for a tuple match, and no existing session is found. NPU1 forwards the packet to an SPU.

Step 2. The Distributed Central Point Creates a Session with a "Pending" State.

When an NPU receives a packet, the NPU send it to the distributed central point, based on the hash algorithm. The distributed central point then looks up the distributed central point session table and creates an entry if needed.

This process entails the following parts:

Copyright © 2016, Juniper Networks, Inc.

1. The distributed central point checks its session table to determine if a session exists for the packet received from the NPU. (An NPU forwards a packet to the distributed central point because it cannot find an existing session for the packet)
2. If there is no entry that matches the packet in the distributed central point session table, the distributed central point creates a pending wing for the session. The distributed central point then sends a query message to the application central point to select an SPU to be used for the session.
3. On receiving the query message, the application central point checks its gate table to determine if a gate exists for the packet. If a gate is matched or some other session distribution algorithm is triggered, the application central point selects another SPU to process the packet; otherwise, the SPU (that is, the distributed central point SPU) is selected. Finally, the application central point sends a query response to the distributed central point.
4. On receiving the query response, the distributed central point forwards the first packet in flow to the selected SPU in a message directing the SPU to set up a session locally to be used for the packet flow. For example, the distributed central point creates a pending wing (a ->b) for the session. The application central point selects SPU1 to be used for it. The distributed central point sends SPU1 the (a->b) packet along with a message to create a session for the distributed central point.

Example: The distributed central point creates a pending wing (a ->b) for the session. It selects SPU1 to be used for it. It sends SPU1 the (a->b) packet along with a message to create a session for it.

Step 3. The SPU Sets Up the Session.

Each SPU, too, has a session table, which contains information about its sessions. When the SPU receives a message from the distributed central point to set up a session, it checks its session table to ensure that a session does not already exist for the packet.

1. If there is no existing session for the packet, the SPU sets up the session locally.
2. The SPU sends a message to the distributed central point directing it to install the session.

NOTE: During first-packet processing, if NAT is enabled, the SPU allocates IP address resources for NAT. In this case, the first-packet processing for the session is suspended until the NAT allocation process is completed.

The SPU adds to the queue any additional packets for the flow that it might receive until the session has been installed.

Example: SPU1 creates the session for (a ->b) and sends a message back to the distributed central point directing it to install the pending session.

Step 4. The Distributed Central Point Installs the Session.

The distributed central point receives the install message from the SPU.

Copyright © 2016, Juniper Networks, Inc.

1. The distributed central point sets the state for the session's pending wing to active.
2. The distributed central point installs the reverse wing for the session as an active wing.

NOTE: For some cases, such as NAT, the reverse wing may be installed on a different distributed central point from the init wing distributed central point.

3. It sends an acknowledge (ACK) message to the SPU, indicating that the session is installed.

Example: The distributed central point receives a message from SPU1 to install the session for the (a->b) wing. It sets the session state for the (a->b) wing to active. It installs the reverse wing (b->a) for the session and makes it active; this allows for delivery of packets from the reverse direction of the flow: destination (b) to be delivered to the source (a).

Step 5. The SPU Sets Up the Session on the Ingress and Egress NPUs.

NPUs maintain information about a session for packet forwarding and delivery. Session information is set up on the egress and ingress NPUs (which sometimes are the same) so that packets can be sent directly to the SPU that manages their flows and not to the distributed central point for redirection.

Step 6. Fast-Path Processing Takes Place.

For the remainder of the steps entailed in packet processing, proceed to Step 1 in "Understanding Fast-Path Processing" on page 1678.

Figure 86 illustrates the first part of the process that the first packet in a flow undergoes after it reaches the device. At this point a session is set up to process the packet and the rest of the packets belonging to its flow. Subsequently, it and the rest of the packets in the flow undergo fast-path processing.

Copyright © 2016, Juniper Networks, Inc.

Figure 86: Session Creation: First-Packet Processing

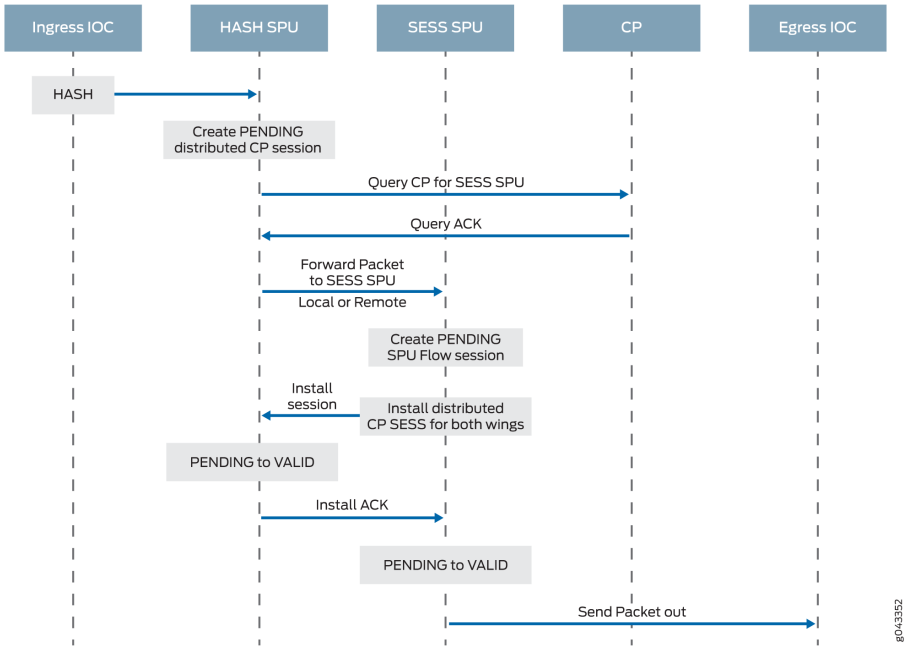

Understanding Fast-Path Processing

All packets undergo fast-path processing. However, if a session exists for a packet, the packet undergoes fast-path processing and bypasses the first-packet process. When there is already a session for the packet's flow, the packet does not transit the central point.

Here is how fast-path processing works: NPUs at the egress and ingress interfaces contain session tables that include the identification of the SPU that manages a packet's flow. Because the NPUs have this session information, all traffic for the flow, including reverse traffic, is sent directly to that SPU for processing.

To illustrate the fast-path process, this section uses an example with a source "a" and a destination "b". The direction from source to destination for the packets of the flow is referred to as (a->b). The direction from destination to source is referred to as (b->a).

Step 1. A Packet Arrives at the Device and the NPU Processes It.

This section describes how a packet is handled when it arrives at a services gateway's IOC.

1. The packet arrives at the device's IOC and is processed by the NPU on the card.

   The NPU performs sanity checks and applies some screens, such as denial-of-service (DoS) screens, to the packet.

2. The NPU identifies an entry for an existing session in its session table that the packet matches.

Copyright © 2016, Juniper Networks, Inc.

3. The NPU forwards the packet along with metadata from its session table, including the session ID and packet tuple information, to the SPU that manages the session for the flow, applies stateless firewall filters and CoS features to its packets, and handles the packet's flow processing and application of security and other features.

Example: Packet (a ->b) arrives at NPU1. NPU1 performs sanity checks on the packet, applies DoS screens to it, and checks its session table for a tuple match. It finds a match and that a session exists for the packet on SPU1. NPU1 forwards the packet to SPU1 for processing.

Step 2. The SPU for the Session Processes the Packet.

Most of a packet's processing occurs on the SPU to which its session is assigned. The packet is processed for packet-based features such as stateless firewall filters, traffic shapers, and classifiers, if applicable. Configured flow-based security and related services such as firewall features, NAT, ALGs, and so on, are applied to the packet. (For information on how security services are determined for a session, see "Juniper Networks Devices Processing Overview" on page 1665.)

1. Before it processes the packet, the SPU checks its session table to verify that the packet belongs to one of its sessions.
2. The SPU processes the packet for applicable features and services.

Example: SPU1 receives packet (a->b) from NPU1. SPU1 checks its session table to verify that the packet belongs to one of its sessions. Then it processes packet (a ->b) according to input filters and CoS features that apply to its input interface. The SPU applies the security features and services that are configured for the packet's flow to it, based on its zone and policies. If any are configured, it applies output filters, traffic shapers and additional screens to the packet.

Step 3. The SPU Forwards the Packet to the NPU.

1. The SPU forwards the packet to the NPU.
2. The NPU applies any applicable screens associated with the interface to the packet.

Example: SPU1 forwards packet (a ->b) to NPU2, and NPU2 applies DoS screens.

Step 4. The Interface Transmits the Packet from the Device.

Example: The interface transmits packet (a->b) from the device.

Step 5. A Reverse Traffic Packet Arrives at the Egress Interface and the NPU Processes It.

This step mirrors Step 1 exactly in reverse. See Step 1 in this section for details.

Example: Packet (b->a) arrives at NPU2. NPU2 checks its session table for a tuple match. It finds a match and that a session exists for the packet on SPU1. NPU2 forwards the packet to SPU1 for processing.

Step 6. The SPU for the Session Processes the Reverse Traffic Packet.

This step is the same as Step 2 except that it applies to reverse traffic. See Step 2 in this section for details.

Copyright © 2016, Juniper Networks, Inc.

Example: SPU1 receives packet (b->a) from NPU2. It checks its session table to verify that the packet belongs to the session identified by NPU2. Then it applies packet-based features configured for the NPU1's interface to the packet. It processes packet (b->a) according to the security features and other services that are configured for its flow, based on its zone and policies. (See "Juniper Networks Devices Processing Overview" on page 1665.)

Step 7. The SPU Forwards the Reverse Traffic Packet to the NPU.

This step is the same as Step 3 except that it applies to reverse traffic. See Step 3 in this section for details.

Example: SPU1 forwards packet (b->a) to NPU1. NPU1 processes any screens configured for the interface.

8. The Interface Transmits the Packet from the Device.

This step is the same as Step 4 except that it applies to reverse traffic. See Step 4 in this section for details.

Example: The interface transmits packet (b->a) from the device.

Figure 87 illustrates the process a packet undergoes when it reaches the device and a session exists for the flow that the packet belongs to.

Figure 87: Packet Walk for Fast-Path Processing

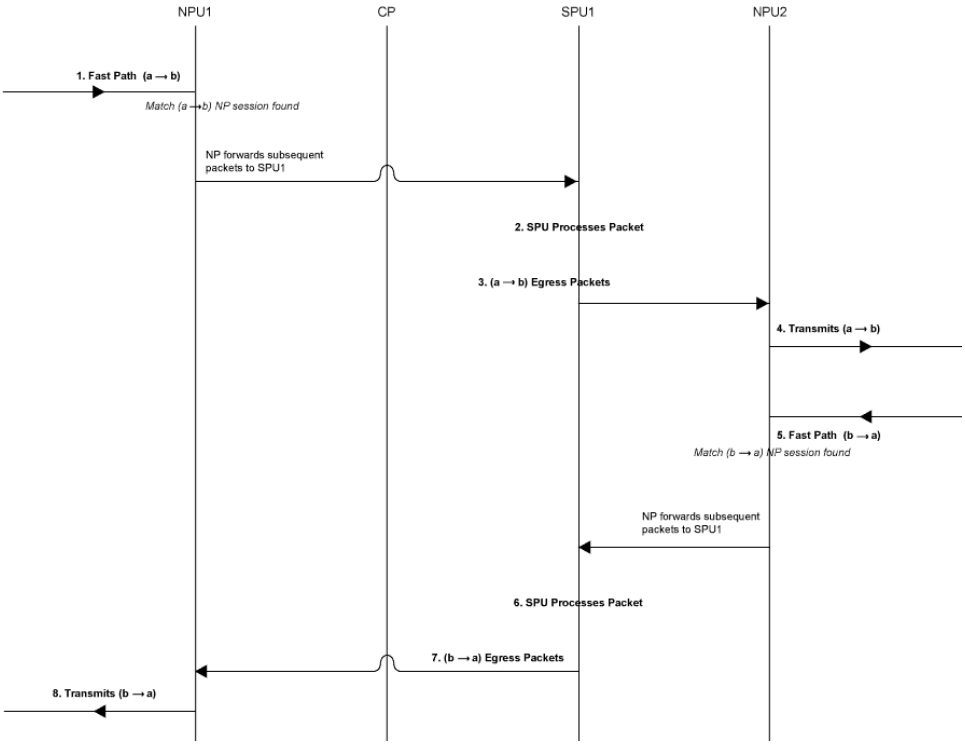

Copyright © 2016, Juniper Networks, Inc.

Understanding Services Processing Units

For a given physical interface, the SPU receives ingress packets from all network processors in the network processor bundle associated with the physical interface. The SPU extracts network processor bundle information from the physical interface and uses the same 5-tuple hash algorithm to map a flow to a network processor index. To determine the network processor, the SPU does a lookup on the network processor index in the network processor bundle. The SPU sends egress packets to the physical interface's local Physical Interface Module (PIM) for the outward traffic.

NOTE: **The network processor and the SPU use the same 5-tuple hash algorithm to get the hash values for the packets.**

Understanding Scheduler Characteristics

For SRX5600 and SRX5800 devices, the IOC supports the following hierarchical scheduler characteristics:

- IFL – The configuration of the network processor bundle is stored in the physical interface data structure. For example, SRX5600 and SRX5800 devices have a maximum of 48 PIMs. The physical interface can use a 48-bit bit-mask to indicate the PIM, or the network processor traffic from this physical interface is distributed in addition to the physical interface's primary network processor.

  On SRX5000 line devices, the iflset functionality is not supported for aggregated interfaces like *reth*.

- IFD – The logical interface associated with the physical interface of a network processor bundle is passed to all the IOCs that have a PIM in the network processor bundle.

Understanding Network Processor Bundling

The network processor bundling feature is available on SRX5000 line devices. This feature enables distribution of data traffic from one interface to multiple network processors for packet processing. A primary network processor is assigned for an interface that receives the ingress traffic and distributes the packets to several other secondary network processors. A single network processor can act as a primary network processor or as a secondary network processor to multiple interfaces. A single network processor can join only one network processor bundle.

Network Processor Bundling Limitations

Network processor bundling functionality has the following limitations:

- Network processor bundling allows a total of 16 PIMs per bundle and 8 different network processor bundle systems.
- You need to reboot the device to apply the configuration changes on the bundle.

Copyright © 2016, Juniper Networks, Inc.

- Network processor bundling is below the reth interface in the overall architecture. You can choose one or both interfaces from the network processor bundle to form the reth interface.
- If the IOC is removed from a network processor bundle, the packets forwarded to the PIM on that IOC are lost.
- When the network processor bundle is enabled, the ICMP, UDP, and TCP sync flooding thresholds no longer apply to an interface. Packets are distributed to multiple network processors for processing. These thresholds apply to each network processor in the network processor bundle.
- Network processor bundling is not supported in Layer 2 mode.
- Because of memory constraints on the network processor, the number of network processor bundled ports that are supported per PIM is limited. Within the network processor bundle, each port needs to have a global port index. The global port index is calculated using the following formula:

  Global_port_index = (global_pic * 16) + port_offset
- Link aggregation groups (LAGs) and redundant Ethernet interface LAGs in chassis cluster implementations can coexist with network processor bundling. However, neither LAGs nor redundant Ethernet interface LAGs can overlap with or share physical links with a network processor bundle.

Understanding Central Point Session Limit Performance Enhancements

The central point architecture prevents data packets from going through the central point by offloading traffic management to Services Processing Units (SPUs). The system session capacity is extended because the session on the central point is removed. The session limit handles packet processing on the SPU and the central point, and updates the session limit counter when the session ages out.

Every session that ages out triggers a central point delete message; many such delete messages consume even more time. The session limit performance is enhanced by replacing the central point delete message with a specific screen report message, thereby saving the central point memory capacity. The central point delete message contains many fields that are not useful for the session limit. So, this message is replaced with the specific screen report message, which only contains the necessary fields.

NOTE: Generating screen report message will traverse all the session limit entry and uses the changed session limit entry to generate the screen report message. For example, if 100 sessions from the same IP address age out, then all sessions have the same IP address and all sessions match the same session limit entry. Therefore this entry can be used to generate one report message instead of 100 delete messages. This method also helps reduce the number of exchange messages.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Enhancements to Central Point Architecture for the SRX5000 Line on page 1670

SRX320 Services Gateway Processing Overview

This topic describes the process that the SRX320 Services Gateway undertakes in establishing a session for packets belonging to a flow that transits the device. The flow services of the SRX320 device are single-threaded and non-distributed. Although it differs from the other SRX Series devices in this respect, the same flow model is followed and the same command line interface (CLI) is implemented.

To illustrate session establishment and the packet "walk" including the points at which services are applied to the packets of a flow, the example described in the following sections uses the simple case of a unicast session:

- Understanding Flow Processing and Session Management on page 1683
- Understanding First-Packet Processing on page 1683
- Understanding Session Creation on page 1684
- Understanding Fast-Path Processing on page 1684

Understanding Flow Processing and Session Management

This topic explains how a session is set up to process the packets composing a flow. In the following topic, the SPU refers to the data plane thread of the SRX320 Services Gateway.

At the outset, the data plane thread fetches the packet and performs basic sanity checks on it. Then it processes the packet for stateless filters and CoS classifiers and applies some screens.

Understanding First-Packet Processing

To determine if a packet belongs to an existing flow, the device attempts to match the packet's information to that of an existing session based on the following six match criteria:

- Source address
- Destination address
- Source port
- Destination port
- Protocol
- Unique token from a given zone and virtual router

The SPU checks its session table for an existing session for the packet. If no existent session is found, the SPU sets up a session for the flow. If a session match is found, the session has already been created, so the SPU performs fast-path processing on the packet.

Copyright © 2016, Juniper Networks, Inc.

Understanding Session Creation

In setting up the session, the SPU executes the following services for the packet:

- Screens
- Route lookup
- Policy lookup
- Service lookup
- NAT, if required

After a session is set up, it is used for all packets belonging to the flow. Packets of a flow are processed according to the parameters of its session. For the remainder of the steps entailed in packet processing, proceed to Step 1 in "Fast-Path Processing". All packets undergo fast-path processing.

Understanding Fast-Path Processing

If a packet matches a session, Junos OS performs fast-path processing as described in the following steps. After a session has been set up for the first packet in a flow, also undergoes fast-path processing. All packets undergo fast-path processing.

1. The SPU applies flow-based security features to the packet.
   - Configured screens are applied.
   - TCP checks are performed.
   - Flow services, such as NAT, ALG, and IPsec are applied, if required.
2. The SPU prepares the packet for forwarding and transmits it.
   - Routing packet filters are applied.
   - Traffic shaping is applied.
   - Traffic prioritizing is applied.
   - Traffic scheduling is applied.
   - The packet is transmitted.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Understanding Session Characteristics for SRX Series Services Gateways on page 1689

Copyright © 2016, Juniper Networks, Inc.

Understanding Central Point Architecture Flow Support for GTP and SCTP

Starting in Junos OS Release 15.1X49-D40, the central point architecture provides enhanced support for GPRS tunneling protocol, control (GTP-C); GPRS tunneling protocol, user plane (GTP-U); and Stream Control Transmission Protocol (SCTP).

The central point architecture is enhanced to address the GTP-C message rate-limiting to protect gateway GPRS support node (GGSN) from GTP-C message flood, to prevent GTP-C packet drop issues during SGSN handover, and to distribute GTP-U traffic handled by a GGSN and SGSN pair on all SPUs by switching to tunnel endpoint identifier (TEID)-based hash distribution. Use the **enable-gtpu-distribution** command to enable or disable GTP-U session distribution. By default, the **enable-gtpu-distribution** command is disabled.

To prevent GTP-C packets from being dropped, a new flow session is created and the GTP-C traffic is allowed to pass even if the GGSN or SGSN direction is not determined. Later, the GGSN IP is determined using the correct SPU to create the flow session; otherwise, the session is migrated to the designated SPU.

To handle load-balancing issues, tag-based hash distribution is used to ensure even distribution of SCTP traffic from different associations among all SPUs. A 32-bit connection tag is introduced that uniquely identifies the GTP-U and the SCTP sessions. The connection tag for GTP-U is the TEID and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.

Related Documentation

- Understanding Enhancements to Central Point Architecture for the SRX5000 Line on page 1670
- enable-gtpu-distribution on page 2362

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 23

Configuring Flow-Based Sessions

- Configuring Security Flow Sessions on page 1689

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 83

Configuring Security Flow Sessions

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Understanding Aggressive Session Aging on page 1690
- Example: Controlling Session Termination for SRX Series Services Gateways on page 1690
- Understanding TCP Session Checks per Policy on page 1692
- Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways on page 1693
- Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways on page 1694
- Example: Configuring TCP Packet Security Checks Per Policy on page 1696
- Configuring the Timeout Value for Multicast Flow Sessions on page 1697
- Clearing Sessions for SRX Series Services Gateways on page 1699

Understanding Session Characteristics for SRX Series Services Gateways

Sessions are created, based on routing and other classification information, to store information and allocate resources for a flow. Sessions have characteristics, some of which you can change, such as when they are terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 1800 seconds. The default timeout for UDP is 60 seconds. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 2 to 4 seconds.

If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse. You can affect the life of a session in the following ways:

- You can specify circumstances for terminating sessions by using any of the following methods:
  - Age out sessions based on how full the session table is
  - Set an explicit timeout for aging out TCP sessions
  - Configure a TCP session to be invalidated when it receives a TCP RST (reset) message

Copyright © 2016, Juniper Networks, Inc.

- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks
  - Change the maximum segment size

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Clearing Sessions for SRX Series Services Gateways on page 1699
- Example: Controlling Session Termination for SRX Series Services Gateways on page 1690

Understanding Aggressive Session Aging

The session table is a limited resource for SRX Series devices. If the session table is full, any new sessions will be rejected by the device.

The aggressive session-aging mechanism accelerates the session timeout process when the number of sessions in the session table exceeds the specified high-watermark threshold. This mechanism minimizes the likelihood that the SRX Series devices will reject new sessions when the session table becomes full.

Configure the following parameters to perform aggressive session aging:

- *high-watermark*—The device performs aggressive session aging when the number of sessions in the session table exceeds the *high-watermark* threshold.
- *low-watermark*—The device exits aggressive session aging and returns to normal when the number of sessions in the session table dips below the *low-watermark* threshold.
- *early-ageout* —During aggressive session aging, the sessions with an age-out time lower than the *early-ageout* threshold are marked as invalid.

On SRX5600, and SRX5800 devices, the SPU checks the session table, locates the sessions for which the timeout value is lower than the early-ageout time value, and then marks them as invalid.

Related Documentation

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- early-ageout on page 1850
- high-watermark on page 1858
- low-watermark on page 1865

Example: Controlling Session Termination for SRX Series Services Gateways

This example shows how to terminate sessions for SRX Series devices based on aging out after a certain period of time, or when the number of sessions in the session table is

Copyright © 2016, Juniper Networks, Inc.

full or reaches a specified percentage. You specify a timeout value or the number of sessions in the session table.

- Requirements on page 1691
- Overview on page 1691
- Configuration on page 1691
- Verification on page 1692

Requirements

Before you begin, understand the circumstances for terminating sessions. See "Understanding Session Characteristics for SRX Series Services Gateways" on page 1689.

Overview

You can control session termination in certain situations—for example, after receiving a TCP FIN Close or receiving an RST message, when encountering ICMP errors for UDP, and when no matching traffic is received before the service timeout. When sessions are terminated, their resources are freed up for use by other sessions.

In this example, you configure the following circumstances to terminate the session:

- A timeout value of 20 seconds.

  NOTE: **The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.**

- An explicit timeout value of 280 seconds, which changes the TCP session timeout during the three-way handshake.

  The command sets the initial TCP session timeout to 280 in the session table during the TCP three-way handshake. The timer is initiated when the first SYN packet is received, and reset with each packet during the three-way handshake. Once the three-way handshake is completed, the session timeout is reset to the timeout defined by the specific application. If the timer expires before the three-way handshake is complete, the session is removed from the session table.

- Any session that receives a TCP RST (reset) message is invalidated.

Configuration

Step-by-Step Procedure

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

To control session termination for SRX Series devices:

1. Specify an age-out value for the session.

   [edit security flow]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set aging early-ageout 20
```

2. Configure an aging out value.

```
[edit security flow]
user@host# set tcp-session tcp-initial-timeoout 280
```

3. Invalidate any session that receives a TCP RST message.

```
[edit security flow]
user@host# set tcp-session rst-invalidate-session
```

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

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

Verification

To verify the configuration is working properly, enter the **show security flow** command.

Related Documentation

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways on page 1693
- Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways on page 1694

Understanding TCP Session Checks per Policy

By default, the TCP SYN check and sequence check options are enabled on all TCP sessions. The Junos operating system (Junos OS) performs the following operations during TCP sessions:

- Checks for SYN flags in the first packet of a session and rejects any TCP segments with non- SYN flags that attempt to initiate a session.
- Validates the TCP sequence numbers during stateful inspection.

The TCP session check per-policy feature enables you to configure SYN and sequence checks for each policy. Currently, the TCP options flags, no-sequence-check and no-syn-check, are available at a global level to control the behavior of services gateways. To support per-policy TCP options, the following two options are available:

- sequence-check-required: The sequence-check-required value overrides the global value no-sequence-check.
- syn-check-required: The syn-check-required value overrides the global value no-syn-check.

To configure per-policy TCP options, you must turn off the respective global options; otherwise, the commit check will fail. If global TCP options are disabled and SYN flood protection permits the first packet, then the per-policy TCP options will control whether SYN and/or sequence checks are performed.

Copyright © 2016, Juniper Networks, Inc.

NOTE:

- The per-policy SYN check required option will not override the behavior of the **set security flow tcp-session no-syn-check-in-tunnel** CLI command.
- Disabling the global SYN check reduces the effectiveness of the device In defending against packet flooding.

Disabling the global SYN check and enforcing the SYN check after policy search will greatly impact the number of packets that the router can process. This in turn will result in intense CPU operations. When you disable global SYN check and enable per-policy SYN check enforcement, you should be aware of this performance impact.

Related Documentation

- Example: Configuring TCP Packet Security Checks Per Policy on page 1696

Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways

This example shows how to disable TCP packet security checks in the device.

- Requirements on page 1693
- Overview on page 1693
- Configuration on page 1694
- Verification on page 1694

Requirements

Before you begin, understand the circumstances for disabling TCP packet security checks. See "Understanding Session Characteristics for SRX Series Services Gateways" on page 1689.

Overview

Junos OS provides a mechanism for disabling security checks on TCP packets to ensure interoperability with hosts and devices with faulty TCP implementations. During no-SYN-check the Junos OS does not look for the TCP SYN packet for session creation. No-sequence check disables TCP sequence checking validation. Also, increases throughput. SYN check and sequence check are enabled by default. The set security flow command disables TCP SYN checks and TCP sequence checks on all TCP sessions thus reduces security. This may be required in scenarios with customers like big transfer files, or with applications that do not correctly work with standards.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

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

To disable TCP packet security checks:

1. Disable the checking of the TCP SYN bit before creating a session.

   ```
   [edit security flow]
   user@host# set tcp-session no-syn-check
   ```

2. Disable the checking of sequence numbers in TCP segments during stateful inspection.

   ```
   [edit security flow]
   user@host# set tcp-session no-sequence-check
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security flow** command.

Related Documentation

- Example: Controlling Session Termination for SRX Series Services Gateways on page 1690
- Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways on page 1694

Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways

This example shows how to set the maximum segment size for all TCP sessions for SRX Series devices.

- Requirements on page 1694
- Overview on page 1695
- Configuration on page 1695
- Verification on page 1696

Requirements

Before you begin, understand the circumstances for setting the maximum segment size. See "Understanding Session Characteristics for SRX Series Services Gateways" on page 1689.

Copyright © 2016, Juniper Networks, Inc.

Overview

You can terminate all TCP sessions by changing the TCP maximum segment size (TCP-MSS). To diminish the likelihood of fragmentation and to protect against packet loss, you can use the tcp-mss statement to specify a lower TCP MSS value. This statement applies to all TCP SYN packets traversing the router's ingress interfaces whose MSS value is higher than the one you specify.

If the DF bit is set, it will not fragment the packet and Junos OS will send ICMP error type 3 code 4 packet to the application server (Destination Unreachable; Fragmentation Needed and DF set). This ICMP error message contains the correct MTU (as defined in tcp-mss) to be used by the application server, which should receive this message and adjust the packet size accordingly. This is specifically required with VPNs, as IPsec has added packet overhead; thus tcp-mss must be lowered appropriately.

NOTE: When running SRX Series devices in packet mode, you use the **set system internet-options tcp-mss** statement to adjust the TCP-MSS value. All ports are affected by the TCP-MSS configuration; you cannot exclude a particular port. When running SRX Series devices in flow mode, although you can use the **set system internet-options tcp-mss** statement, we recommend using only the **set security flow tcp-mss** statement to adjust the TCP-MSS value. If both statements are configured, the lower of the two values will take effect.

Configuration

Step-by-Step Procedure

To configure the maximum segment size for all TCP sessions:

1. Set the TCP maximum segment size for all TCP sessions.

   ```
   [edit security flow]
   user@host# set tcp-mss all-tcp mss 1300
   ```

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

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

Results

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

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

```
[edit]
user@host# show security flow
...
```

Copyright © 2016, Juniper Networks, Inc.

```
tcp-mss{
  all-tcp{
    mss 1300;
  }
}
...
```

Verification

To verify the configuration is working properly, enter the **show configuration security flow** command from operational mode.

```
user@host> show configuration security flow
tcp-mss{
  all-tcp{
    mss 1300;
  }
}
```

Related Documentation

- Example: Controlling Session Termination for SRX Series Services Gateways on page 1690
- Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways on page 1693

Example: Configuring TCP Packet Security Checks Per Policy

This example shows how to configure TCP packet security checks for each policy in the device.

- Requirements on page 1696
- Overview on page 1696
- Configuration on page 1696
- Verification on page 1697

Requirements

Before you begin, you must disable the tcp options, **tcp-syn-check**, and **tcp-sequence-check** that are configured at global level. See "Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways" on page 1693.

Overview

The SYN and sequence check options are enabled by default on all TCP sessions. In environments that need to support large file transfers, or that run nonstandard applications, it might be necessary to configure sequence and sync checks differently for each policy. In this example, you configure sequence and sync check for policy **pol1**.

Configuration

Step-by-Step Procedure

To configure TCP packet security checks at the policy level:

1. Configure the checking for the TCP SYN bit before creating a session.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security policies from-zone Zone-A to-zone Zone-B policy pol1 then
  permit tcp-options syn-check-required
```

2. Configure the checking for sequence numbers in TCP segments during stateful inspection.

```
[edit]
user@host# set security policies from-zone Zone-A to-zone Zone-B policy pol1 then
  permit tcp-options sequence-check-required
```

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

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

Verification

To verify that the configuration is working properly, enter the **show security policies detail** command.

Related Documentation

- Understanding TCP Session Checks per Policy on page 1692
- Example: Disabling TCP Packet Security Checks for SRX Series Services Gateways on page 1693
- Example: Setting the Maximum Segment Size for All TCP Sessions for SRX Series Services Gateways on page 1694

Configuring the Timeout Value for Multicast Flow Sessions

You can configure the timeout value for multicast flow sessions by configuring a custom application and associating the application with a policy.

Multicast flow sessions have one template session and one or more leaf sessions. Because these sessions are linked together, they can have only one timeout value. The timeout value for multicast flow sessions is determined by considering the timeout values configured in the leaf session policies and the IP protocol timeout values. The highest of these timeout values is selected as the multicast flow session timeout.

If no leaf session timeout values are configured, the IP protocol timeout value is automatically used as the timeout value for the mulicast flow session. The IP protocol timeout is the default and is not configurable.

Configuring leaf session timeouts can be especially helpful for multicast streams that have a longer packet interval than the default IP protocol timeout. For example, multicast streams with a packet interval of more than 60 seconds would experience premature aging-out of flow sessions and packet drops with the UDP timeout value, which is always 60 seconds. For such streams, you can configure a higher leaf session timeout value and prevent packet drop.

Copyright © 2016, Juniper Networks, Inc.

To set the leaf session timeout value, configure a custom application and associate the application with a policy:

1. Create a custom application, specify its properties, and specify bypassing the application type.

   ```
   [edit]
   user@host# edit applications application my-udp

   [edit applications application my-udp]
   user@host# set protocol udp
   user@host# set destination-port 5000
   user@host# set application-protocol ignore
   ```

2. Set the timeout value for the application protocol.

   ```
   [edit applications application my-udp]
   user@host# set inactivity-timeout 500
   ```

3. Create a policy.

   ```
   [edit]
   user@host# edit security policies from-zone vr-zone-1 to-zone junos-host policy
     my-policy

   [edit security policies from-zone vr-zone-1 to-zone junos-host policy my-policy]
   user@host# set match source-address 192.0.2.1
   user@host# set match destination-address any
   ```

4. Associate the custom application (with the configured timeout) to the policy.

   ```
   [edit security policies from-zone vr-zone-1 to-zone junos-host policy my-policy]
   user@host# set match application my-udp
   user@host# set then permit
   ```

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

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

6. To verify the updated session timeout value, enter the **show security flow session** command.

   ```
   user@host> show security flow session destination-prefix 203.0.113.0

   Session ID: 2363, Policy name: N/A, Timeout: 498, Valid
     In: 18.1.1
   Session ID: 2363, Policy name: N/A, Timeout: 498, Valid
     In: 192.0.2.1/17767-->203.0.113.0/5000;udp, If: ge-0/0/1.0, Pkts:0, Bytes:0

     Out: 203.0.113.0/5000-->192.0.2.1/17767/17767;udp, If:.local..4, Pkts:0,
   Bytes:0
   .2/17767-->224.1.1.1/5000;udp, If: ge-0/0/1.0, Pkts:0, Bytes:0
     Out: 255.255.255.255/5000-->255.255.255.255/17767;udp, If:.local..4, Pkts:0,
    Bytes:0


   Session ID: 2364, Policy name: my-policy/4, Timeout: -1, Valid
     In: 192.0.2.1/17767-->203.0.113.0/5000;udp, If:ge-0/0/1.0, Pkts:1011,
   Bytes:258816
     Out: 203.0.113.0/5000-->192.0.2.1/17767;udp, If:ppe0.32769, Pkts:0, Bytes:0
   Total sessions: 2
   ```


Copyright © 2016, Juniper Networks, Inc.

In this output, the session ID 2363 section displays a template session. A timeout value of 498 indicates that the template session timeout value is ticking down from the configured value of 500 seconds.

The session ID 2364 section displays a leaf session. The timeout value of -1 essentially indicates that the session will not age out unless the template session ages out.

In this example, the configured leaf session timeout value of 500 seconds is the highest timeout value and is accepted as the template session timeout value for the multicast flow session.

Related Documentation

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- *Multicast Feature Guide for Security Devices*

Clearing Sessions for SRX Series Services Gateways

You can use the **clear** command to terminate sessions. You can clear all sessions, including sessions of a particular application type, sessions that use a specific destination port, sessions that use a specific interface or port, sessions that use a certain IP protocol, sessions that match a source prefix, and resource manager sessions.

- Terminating Sessions for SRX Series Services Gateways on page 1699
- Terminating a Specific Session for SRX Series Services Gateways on page 1699
- Using Filters to Specify the Sessions to Be Terminated for SRX Series Services Gateways on page 1699

Terminating Sessions for SRX Series Services Gateways

You can use the following command to terminate all sessions except tunnel and resource manager sessions. The command output shows the number of sessions cleared. Be aware that this command terminates the management session through which the clear command is issued.

```
user@host> clear security flow session all
```

Terminating a Specific Session for SRX Series Services Gateways

You can use the following command to terminate the session whose session ID you specify.

```
user@host> clear security flow session session-identifier 40000381
```

Using Filters to Specify the Sessions to Be Terminated for SRX Series Services Gateways

You can terminate one or more sessions based on the filter parameter you specify for the **clear** command. The following example uses the protocol as a filter.

```
user@host> clear security flow session protocol 89
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 24

Improving Flow-Based Performance

- Expanding Session Capacity by Device on page 1703
- Managing Sessions and Flow Distribution on page 1707
- Reducing Long Packet-Processing Latency by Express Path on page 1713

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 84

Expanding Session Capacity by Device

- Expanding Session Capacity by Device on page 1703
- Reverting to Default Session Capacity on an SRX5800 Device on page 1704
- Verifying the Current Session Capacity on page 1704

Expanding Session Capacity by Device

To take advantage of the processing potential of a fully loaded SRX5600, or SRX5800 device, you can expand the maximum number of concurrent sessions for these devices.

Table 167 shows the maximum number of concurrent sessions allowed on these devices by default and with expanded capacity.

Table 167: Maximum Central Point Session Increases

SRX Series Devices	Maximum Concurrent Sessions on a Fully Loaded System	
	Default	With Expanded Capacity
SRX5400	42 million	Expansion not available
SRX5600	114 million	Expansion not available
SRX5800	258 million	Expansion not available

The method used for expanding session capacity depends on the device:

- CLI optimization option on an SRX5800 device

Related Documentation

- Reverting to Default Session Capacity on an SRX5800 Device on page 1704

Copyright © 2016, Juniper Networks, Inc.

Reverting to Default Session Capacity on an SRX5800 Device

Reverting to the default session capacity on an SRX5800 device requires a CLI configuration change.

1. Enter the following command at the CLI configuration prompt to reestablish the default session capacity value:

   user@host# **set security gprs gtp enable**

2. Reboot the device to implement the new value.

Verifying the Current Session Capacity

Purpose The central point session summary includes the maximum sessions setting for the device. From this value you can determine if the session capacity has been modified as you expected.

Action To verify the current setting of the central point session capacity, enter the following CLI command.

user@host> **show security flow cp-session summary**

```
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Total sessions: 1
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747
```

Meaning The **Maximum sessions** value reflects the current session capacity on your device. A value of 14000000 means that the SRX5800 device is configured for the expanded number of central point sessions.

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 85

Managing Sessions and Flow Distribution

- Understanding Load Distribution in High-End SRX Series Devices on page 1707
- Understanding Load Distribution on SRX5000 Line Devices Using the Packet-Ordering Function on page 1709
- Disabling Packet-Ordering Mode on SRX5000 Line Devices on page 1710

Understanding Load Distribution in High-End SRX Series Devices

The load distribution algorithm is adjusted based on session capacity and processing power.

Hash-based session distribution uses a hash table. The SPU session weight table is used to assign an SPU ID to each hash index in the session distribution hash table. This way, the number of sessions created on each SPU using hash-based distribution is proportional to the SPU's weight in the SPU session weight table. Each NPU also keeps an identical SPU session weight table and session distribution hash table that it uses to select an SPU to forward packets that do not match an NPU session.

In hash-based session distribution, weights are based on session capacity. We recommend the hash session distribution mode when high session capacity is required.

NOTE: **Load distribution on SRX5000 line devices is always hash-based.**

Insertion and removal of SPCs causes recalculation of the SPU session weight table at central point initialization time because the chassis must reboot after insertion.

Hash-Based Forwarding on the SRX5K-MPC, SRX5K-MPC3-40G10G (IOC3), and the SRX5K-MPC3-100G10G (IOC3)

On a high-end SRX Series device, a packet goes through a series of events involving different components as it progresses from ingress to egress processing. With the datapath packet forwarding feature, you can obtain quick delivery of I/O traffic over the SRX 5000 line of devices.

The SRX5K-MPC, SRX5K-MPC3-40G10G (IOC3), and SRX5K-MPC3-100G10G (IOC3) are interface cards supported on the SRX5400, SRX5600, and SRX5800 devices. The

Copyright © 2016, Juniper Networks, Inc.

Modular Port Concentrator (MPC) provides load-balancing services for Services Processing Units (SPUs) by using the hash-based forwarding method.

In hash-based forwarding, the packet might be forwarded by the MPC to a selected SPU (DCP) instead of the central point. This approach enhances session scaling and prevents overloading of the central point.

Hash value calculation involves the following steps:

- For IPv4 packets, the hash-based forwarding module generates the hash value based on Layer 3 and Layer 4 information, depending on different Layer 4 protocol types.
- For Stream Control Transmission Protocol (SCTP), TCP, UDP, Authentication Header (AH), edge service provider (ESP), and Internet Control Message Protocol (ICMP) protocols, the hash module utilizes Layer 4 information to generate the hash value. For any other protocols, only Layer 3 information is used in hash generation.
- For IPv4 fragment packets, the hash value is calculated using only the Layer 3 information. This also applies to the first fragment of the packet.
- For non-IP packets, the hash-based forwarding module uses the Layer 2 information to calculate the hash value.

Once a hash value is calculated according to the packet's Layer 2, Layer 3, or Layer 4 information, an SPU ID is assigned to each hash index in the session distribution hash table.

NOTE: The SRX5K-MPC (IOC2), SRX5K-MPC3-40G10G (IOC3), and SRX5K-MPC3-100G10G (IOC3) can only be used on SRX5400, SRX5600, and SRX5800 devices that are configured for hash-based session distribution.

When the hash-based session distribution mode is enabled, the system changes its behavior to high-session-capacity-based mode when the SRX5K-MPC, SRX5K-MPC3-40G10G (IOC3), and SRX5K-MPC3-100G10G (IOC3) are installed on the device.

NOTE: On SRX5000 line devices with an SRX5K-MPC, SRX5K-MPC3-40G10G (IOC3), or SRX5K-MPC3-100G10G (IOC3) installed, during a system or an SPU reboot, when the hash-based session distribution mode is enabled, traffic will pass only when all SPUs are up after the reboot.

The MPCs on the IOC3 provide load-balancing services for SPUs by performing hash-based datapath packet forwarding to interconnect with all existing IOCs and SPCs.

The IOC3 processes ingress and egress packets. The IOC3 parses the ingress packet and sends it to the SPU for further security processing, including flow session lookup, zone and policy check, VPN, ALG, and so on.

Copyright © 2016, Juniper Networks, Inc.

The IOC3 manages packet data memory and fabric queuing for packet lookup and encapsulation functions.

NOTE: Starting with Junos OS Release 15.1X49-D10, hash-based session distribution is the default mode for the SRX5400, SRX5600, and SRX5800 devices. Selection of hash keys depends on application protocols.

The IOC3 sets up a security flow table (IPv4 and IPv6) including key, result table, and packet memory.

The following functions are provided with the flow table:

- Flow lookup
- Flow insertion and deletion
- Security flow aging out
- Security flow statistics

Understanding Load Distribution on SRX5000 Line Devices Using the Packet-Ordering Function

The packet-ordering function improves the performance of the device by activating the built-in packet-ordering function of the Packet Ordering Engine on the XLP processor on the application central point. When you activate the packet-ordering function, the load-balancing thread (LBT) and the packet-ordering thread (POT) are removed and their core is freed to run the flow thread.

The flow thread receives the packets, processes them, and sends or drops them. For packets that require no ordering, the flow thread notifies the Network Acceleration Engine (NAE) egress to send or drop the packets. For packets that require ordering, the flow thread notifies the Packet Ordering Engine to dequeue the packets from the ordering list and to send or drop the packets in order.

The packet-ordering functionality using the Packet Ordering Engine is supported on SRX5400, SRX5800 and SRX5600 devices with next-generation SPCs.

Related Documentation

- Understanding SRX Series Services Gateways Central Point Architecture on page 1669
- Disabling Packet-Ordering Mode on SRX5000 Line Devices on page 1710
- **packet-ordering-mode (Application Services) on page 1871**

Copyright © 2016, Juniper Networks, Inc.

Disabling Packet-Ordering Mode on SRX5000 Line Devices

The packet-ordering functionality using the Packet Ordering Engine is supported on SRX5400, SRX5800 and SRX5600 devices with next-generation SPCs. By default, packet-ordering mode using the Packet Ordering Engine is enabled. To disable the packet-ordering functionality using the Packet Ordering Engine, you must update the packet-ordering mode on the device.

The following packet ordering modes are supported:

- software—Disables the packet-ordering mode using the Packet Ordering Engine.
- hardware—Enables the packet-ordering mode using the Packet Ordering Engine. This is the default option.

To disable the packet-ordering mode using the Packet Ordering Engine:

1. Enter the following command at the CLI configuration prompt to specify the packet-ordering mode.

   ```
   [edit]
   user@host#  set security forwarding-process application-services
     packet-ordering-mode software
   ```

2. Use the **show security forwarding-process** command to review your configuration.

   ```
   [edit]
   user@host# show security forwarding-process
   application-services{
     packet-ordering-mode software;
   }
   ```

3. Check your changes to the configuration before committing.

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

   ```
   warning: System packet ordering mode changed, reboot is required to take effect.
   If you have deployed a cluster, be sure to reboot all nodes.
   configuration check succeeds
   ```

4. Commit the configuration.

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

   ```
   warning: System packet ordering mode changed, reboot is required to take effect.
   If you have deployed a cluster, be sure to reboot all nodes.
   commit complete
   ```

5. Reboot the device at an appropriate time.

6. Use the **show security flow status** command to verify the packet-ordering mode.

   ```
   user@host> show security flow status
   ```

   ```
   Flow forwarding mode:
     Inet forwarding mode: flow based
     Inet6 forwarding mode: drop
     MPLS forwarding mode: drop
   ```


Copyright © 2016, Juniper Networks, Inc.

```
  ISO forwarding mode: drop
Flow trace status
  Flow tracing status: off
Flow session distribution
  Distribution mode: RR-based
Flow packet ordering
  Ordering mode: Software (reboot needed to change to Software)
```

Related Documentation

- Understanding Load Distribution on SRX5000 Line Devices Using the Packet-Ordering Function on page 1709

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 86

Reducing Long Packet-Processing Latency by Express Path

- Understanding Session Cache on page 1713
- Express Path Overview on page 1717
- Understanding the Express Path Solution on page 1731
- Enabling and Disabling Express Path on page 1732
- Example: Enabling Express Path in Security Policies on page 1733
- Example: Configuring an IOC on SRX5000 Line Devices to Support Express Path on page 1735
- Example: Configuring an SRX5K-MPC on an SRX5000 Line Device to Support Express Path on page 1737
- Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path on page 1739
- Example: Configuring Express Path on an SRX5000 Line Device with IOC3 for IPv6 Traffic on page 1742
- Example: Configuring Low Latency on page 1746

Understanding Session Cache

- Overview on page 1713
- Selective Session Cache Installation on page 1714
- IPsec VPN Session Affinity Enhancement Using Session Cache on page 1716
- Fragmentation Packet Ordering Using NP Session Cache on page 1716

Overview

The SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), and SRX5K-MPC3-40G10G (IOC3) on SRX5400, SRX5600, and SRX5800 devices support session cache and selective installation of the session cache.

Session cache is used to cache a conversation between the network processor (NP) and the SPU on an IOC. A conversation could be a session, GTP-U tunnel traffic, IPsec VPN tunnel traffic, and so on. A conversation has two session cache entries, one for incoming traffic and the other for reverse traffic. Depending on where the traffic ingress and egress

Copyright © 2016, Juniper Networks, Inc.

ports are, two entries might reside in the same network processor or in different network processors. IOCs support session cache for IPv6 sessions.

A session cache entry is also called a *session wing*.

Session cache on the IOC leverages Express Path (formerly known as *services offloading*) functionality and helps prevent issues such as high latency and IPsec performance drop.

A session cache entry records:

- To which SPU the traffic of the conversion should be forwarded
- To which egress port the traffic of the conversion should be forwarded in Express Path mode
- What processing to do for egress traffic, for example, NAT translation in Express Path mode

Prior to Junos OS Release 15.1X49-D10, on the IOC2, session cache was used for Express Path sessions only. Other traffic was hashed to SPUs based on their 5-tuple key information. VPN traffic employed the concept of the anchored SPU, which did not necessarily coincide with the functions of the flow SPU. The network processor could only forward the packets to the flow SPU based on the 5-tuple hash. The flow SPU then forwarded the packet to the anchored SPU. This created an extra hop for VPN traffic, which wasted the switch fabric bandwidth and reduced the VPN throughput roughly by half. This performance reduction occurred because the traffic still had to go back to the flow SPU after processing on the anchored SPU.

Starting with Junos OS Release 15.1X49-D10, the session cache of the sessions in the IOC helps solve these issues. The SPU can now instruct the IOC session cache to forward subsequent traffic to a specific anchor SPU.

The session cache table is extended on IOC to support the NP sessions. Express Path traffic and NP traffic share the same session cache table on IOCs. Express Path traffic is forwarded by the IOC itself either locally or to another IOC, because the traffic does not require any services from the SPU. NP traffic is forwarded to the SPU specified in the session cache for further processing. All the session cache entries are shared by both Express Path session traffic and NP traffic.

To enable session cache on the IOCs you need to run the **set chassis fpc *<fpc-slot>* np-cache** command.

NOTE: The IOC2 and the IOC3 utilize the delay sessions delete mechanism. The same sessions (sessions with the same five tuples) that are deleted and then reinstalled immediately are not cached on the IOCs.

Selective Session Cache Installation

To avoid high latency, improve IPSec performance, and to better utilize the valuable resources, certain priority mechanisms are applied to both flow module and the IOC.

Copyright © 2016, Juniper Networks, Inc.

The IOCs maintain and monitor session cache usage threshold levels. The IOCs also communicate the session cache usage to the SPU, so that when a certain session cache usage threshold is reached, the SPU only sends session cache installation requests for selective high-priority traffic sessions.

The following three priority levels are used to determine which type of traffic can install session cache on the IOCs:

- **Priority 1 (P1)**—Express Path qualified traffic
- **Priority 2 (P2)**—IPsec tunneling, Fragmentation ordering, and NAT/SZ traffic
- **Priority 3 (P3)**—All other types of traffic

The IOCs maintain and monitor session cache usage threshold levels. Session cache usage less than 25 percent is defined as “green,” 26 to 50 percent is “yellow,” 51 to 99 percent is “orange,” and 99 percent and above is defined as “red.” The IOCs update current real-time session cache usage to the SPU. The SPU requests the IOC to install the session cache for selective high-priority traffic sessions. When the session cache usage is green, the session cache will be installed for all types of traffic. When the usage is yellow, session can be installed only for certain sessions including IPsec, Fragmentation, and NAT/SZ traffic sessions, and when the usage is orange, only Express Path qualified sessions will be installed. When the session cache usage is red, NP cache and the Express Path sessions are not allowed to install from SPU to the IOCs.

Table 168 shows the NP cache installation bars for different types of traffic.

Table 168: Session Cache Installation Bars

Traffic Type	Green	Yellow	Orange	Red
Express Path traffic	Yes	Yes	Yes	No
IPsec, Fragmentation, and NAT/SZ traffic	Yes	Yes	No	No
Other traffic	Yes	No	No	No

To conserve session entries on the IOC, the flow module selectively installs sessions on the IOC. To facilitate the session install selection, the IOC maintains corresponding thresholds to provide an indication to the flow module (on how full the session cache table is on the IOCs).Two bits in the meta header (see Table 169) are added to indicate the current cache table utilization status. All packets going to the SPU will carry these two status bits to inform the flow module of the utilization of the cache table on the IOC.

Table 169 shows the cache table utilization (CTU) bits and the respective session cache table utilization.

Copyright © 2016, Juniper Networks, Inc.

Table 169: Session Cache Table Utilization Bits Status

Session Cache Table Utilization (CTU) Bits	IOC Session Cache/Express Path Table Utilization	Action
00	0% < utilization < 25%	Flowd installs any eligible session.
01	25% < utilization < 75%	Flowd installs only high-priority sessions, such as Express Path, IPsec, IPsec clear-text, NAT, and fragmented sessions.
10	75% < utilization < 100%	Flowd installs only Express Path sessions.
11	NP cache/Express Path table is full	Flowd stops installing sessions.

IPsec VPN Session Affinity Enhancement Using Session Cache

SRX Series devices are fully distributed systems, and an IPsec tunnel is allocated and anchored to a specific SPU. All the traffic that belongs to an IPsec tunnel is encrypted and decrypted on its tunnel-anchored SPU. In order to achieve better IPsec performance, IOC improves the flow module to create sessions for IPsec tunnel-based traffic (before encryption and after decryption) on its tunnel-anchored SPU, and installs session cache for the sessions so that the IOC can redirect the packets directly to the same SPU to minimize packet-forwarding overhead. Express Path traffic and NP traffic share the same session cache table on IOCs.

You need to enable session cache on the IOCs and set the security policy to determine whether a session is for Express Path (formerly known as *services offloading*) mode on the selected Flexible PIC Concentrator (FPC).

To enable IPsec VPN affinity use, the **set security flow load-distribution session-affinity ipsec** command.

Fragmentation Packet Ordering Using NP Session Cache

A session might consist of both normal and fragmented packets. With hash-based distribution, 5-tuple and 3-tuple key can be used to distribute normal and fragmented packets to different SPUs, respectively. On SRX Series devices, all the packets of the session are forwarded to a processing SPU. Due to forwarding and processing latency, the processing SPU might not guarantee packet ordering of the session.

Session cache on the IOCs ensure ordering of packets of a session with fragmented packets. A session cache entry is allocated for normal packets of the session and a 3-tuple key is used to find the fragmented packets. On receipt of the first fragmented packet of the session, the flow module allows the IOC to update the session cache entry to remember the fragmented packets for the SPU. Later, IOC forwards all subsequent packets of the session to the SPU to ensure ordering of packets of a session with fragmented packets.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Express Path Overview on page 1717
- Understanding VPN Session Affinity on page 6789

Express Path Overview

Understanding Express Path Functionality

Express Path (formerly known as *services offloading*) is a mechanism for processing fast-path packets in the network processor instead of in the Services Processing Unit (SPU). This method reduces the packet-processing latency that arises when packets are forwarded from network processors to SPUs for processing and back to I/O cards (IOCs) for transmission.

Express Path considerably reduces packet-processing latency by 500–600 percent.

When the first packet arrives at the interface, the network processor forwards it to the SPU. If the SPU verifies that the traffic is qualified for Express Path, an Express Path session is created on the network processor. If the traffic does not qualify for Express Path, a normal session is created on the network processor. If an Express Path session is created, the subsequent fast-path packets are processed in the network processor itself.

NOTE: **A normal session forwards packets from the network processor to the SPU for fast-path processing, whereas an Express Path session processes fast-path packets in the network processor and the packets exit out of the network processor itself.**

When an Express Path session is created on the network processor, subsequent packets of the flow match the session on the network processors. The network processor then processes and forwards the packet. The network processor also handles additional processing such as TCP sequence check, time to live (TTL) processing, Network Address Translation (NAT), and Layer 2 header translation.

The network processor forwards packets to the SPU in the following cases:

- When the first packet arrives at the interface, the network processor forwards it to the central point (CP). The central point in turn forwards the packet to the SPU. The SPU then creates a session on the network processor.
- When an SPU session exists even if no network processor session exists, the network processor forwards a packet to the central point, which in turn forwards the packet to the SPU. The SPU then creates a session on the network processor.
- When a packet matches a normal session on the network processor, it is forwarded to the SPU.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Prior to Junos OS Release 12.3X48-D10, Express Path was a licensed software feature. Starting with Junos OS Release 12.3X48-D10, the Express Path license is no longer required to enable this functionality. Your previously acquired Express Path license will not be effective anymore.

Understanding Express Path Support on SRX Series Devices

Table 170 provides details about the Express Path support on different SRX Series cards.

Table 170: Express Path Support on SRX Series Device Cards

SRX Series Device	Card Name and Model Number	Earliest Supported Release
SRX5000 Line Devices I/O Cards (IOCs)		
SRX5600, SRX5800	SRX5K-40GE-SFP	Junos OS Release 11.4
SRX5600, SRX5800	SRX5K-4XGE-XFP	Junos OS Release 11.4
SRX5600, SRX5800	SRX5K-FPC-IOC containing one of the following cards: • SRX-IOC-16GE-TX • SRX-IOC-4XGE-XFP • SRX-IOC-16GE-SFP	Junos OS Release 11.4
SRX5400, SRX5600, SRX5800	SRX5K-MPC containing one of the following MICs: • SRX-MIC-10XGE-SFFP • SRX-MIC-2X40GE-OSFP • SRX-MIC-1X100GE-CFP • SRX-MIC-20GE-SFP	Junos OS Release 12.3X48-D10
SRX5400, SRX5600, SRX5800	SRX5K-MPC3 (IOC3) containing one of the following MPCs: • SRX5K-MPC3-40G10G (24x10GE + 6x40GE MPC) • SRX5K-MPC3-100G10G (2x100GE + 4x10GE MPC)	Junos OS Release 15.1X49-D10

NOTE: Different Express Path features are supported on different cards for different Junos OS releases. See the *Junos OS Release Notes* for details.

Copyright © 2016, Juniper Networks, Inc.

NOTE: On the SRX5600 and SRX5800 Services Gateways, the Express Path sessions for traffic that traverse between legacy IOC cards and the SRX5K-MPC or the SRX5K-MPC3 are not supported.

The Express Path sessions traversing only on legacy IOC cards or only on the SRX5K-MPC (IOC2), the SRX5K-MPC3-100G10G (IOC3), or the SRX5K-MPC3-40G10G (IOC3) are supported. However, the SRX5K-MPC (IOC2), the SRX5K-MPC3-100G10G (IOC3), the SRX5K-MPC3-40G10G (IOC3), and the legacy IOCs can still be present on the same chassis.

Understanding Express Path Features

Wing Statistics Counter

The network processor in Express Path mode provides the option for each flow entry to keep a per-wing bytes counter. The counter captures the number of bytes that the network processor sends out over the wing.

When the counter is enabled, for every ingress packet, the network processor searches its flow entry (a session wing). If the packet belongs to an established flow entry, the network processor increases the byte counter of the flow entry by byte count in the packet. The network processor periodically copies a packet (called a copy-packet) of each flow entry to its associated SPU, allowing the SPU to maintain the session. The network processor sends flow byte counter values in the header of copy-packet packets. The SPU accumulates and keeps per-wing statistics counters.

NOTE: The counter value carried to the SPU is always one packet short to allow the SPU to add the current packet's byte count to the counter to get the correct total. For example, if packet N's copy carries a counter value to the SPU, the counter value is the total bytes received in the flow up to packet N-1.

The counter value does not include packets that were sent before the session was set up on the network processor. Therefore, the SPU might need to account for the three-way handshake packet and other packets sent through the SPU. The actual session byte counter shown on the SPU might be short by the amount of bytes sent by the client during the copy interval. This discrepancy results because these bytes can be counted locally by the network processor, but have not yet been reported to the SPU.

NOTE: You cannot change the statistics configuration during the life cycle of a live session. Disabling or enabling the per-wing statistics configuration while a session is alive at the network processor invalidates the session statistics on the current session. The new sessions statistics can be valid only after the configuration changes are committed . Network processor per-wing counters cannot be cleared.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Wing statistics counter configuration is not supported on the SRX5K-MPC.

Sessions per Wing Statistics

The NP-IOC has a larger static RAM (SRAM) to accommodate session resources, thus hosting more sessions per PIC. Table 171 displays the total number of session wings, including both Express Path and non-Express Path.

Table 171: Total Number of Sessions per Wing in Network Processor Express Path Configuration Mode

Total Number of Wings		Number of Express Path UDP Wings		Number of Express Path TCP Wings	
Cards and SRX Series Device	Non-Express Path Mode Sessions	Without Statistics	With Statistics	Without Statistics	With Statistics
IOC	1.3 million	1.3 million	900,000	600,000	400,000
FIOC	2.3 million	198,000	900,000	600,000	400,000
SRX5000 line device SRX5K-MPC	NA	1.8 million	1.8 million	1.8 million	1.8 million
SRX5000 line device SRX5K-MPC3 (IOC3)	NA	2.0 million	2.0 million	2.0 million	2.0 million

Cross-Network Traffic

Express Path provides additional cross-network-processor support; therefore, it is no longer restricted to the ports of the same network processor. If network processors for both the ingress and egress ports are in Express Path mode, then Express Path packets are directly forwarded from the ingress network processor to the egress network processor in the fast-flow path. Packets cross switch fabric when they are forwarded from one network processor to another, thus increasing the latency of the packet. In Express Path mode, the latency of cross-network-processor packets is higher than the packets that are forwarded within an individual network processor.

NOTE: The SRX5K-MPC receives session messages from the SPU. The session messages carry the information to support inter- and intra-Packet Forwarding Engine Express Path for IPv4.

Copyright © 2016, Juniper Networks, Inc.

LAG Support in Express Path Mode

Ethernet link aggregation groups (LAGs) combine links and provide increased bandwidth and link availability. Express Path reduces packet latency by processing and forwarding packets in the network processor instead of in the Services Processing Unit (SPU). Supporting LAG in the Express Path mode combines the benefits of both these features and provides enhanced throughput, link redundancy, and reduced packet latency.

Qualifying for Express Path Mode

You can use the links in a LAG as ingress or egress interfaces in Express Path mode. The LAG links can include links from different network processors (in case of legacy cards such as IOCs or Flex IOCs) or from the same modular port concentrator (in case of SRX5K-MPC). For a LAG link to qualify for Express Path, all its member links should be connected to Express Path-enabled network processors. If Express Path is disabled on any of the member links in a LAG, a regular session (non-Express Path session) is created. Also, LAG links are not supported between legacy cards such as IOCs or Flex IOCs and the SRX5K-MPC.

LAG and Network Processor Wings

The network processor checks the egress interface in the Express Path mode for each wing. Per-wing traffic distribution over a LAG interface is achieved by letting the SPU install wings pointing to egress interfaces with a balanced distribution.

The network processor periodically copies a packet (called a *copy-packet*) to the SPU, allowing it to maintain the session. The copy-packet contains the egress interface information, which the SPU uses to handle LAG member change cases; for example, when a link is down or disabled. When there is no member interface that can be used as an egress interface for transmitting traffic, the network processor session is updated from Express Path to non-Express Path and the packet is sent to the SPU. A new Express Path network processor session is then installed using a new, valid egress interface.

- First wing—On the egress interface, the SPU selects one LAG active member link as the outgoing interface for a specific fast-forward session. The SPU treats this active member interface just like any physical interface in the Express Path mode and records the interface to the network processor fast-forward session. After that, all traffic that matches this network processor session is directly transmitted through that member link.
- Reverse wing—the reverse network processor session is installed only when all LAG member links are connected to a single network processor. When member links of a LAG are from multiple network processors, the reverse network processor session is initiated by reverse traffic later (it is not preinstalled).

  The Express Path network processor session needs to have outgoing interface information to send traffic. If the incoming interface is a physical interface, then it can be used as an outgoing interface for the reverse wing. However, if the incoming interface is an aggregated Ethernet interface, the SPU selects a member interface to be the outgoing interface.

Copyright © 2016, Juniper Networks, Inc.

LAG and Network Processor Session Updates

Some changes in the LAG interfaces can cause network processor session updates:

- LAG interface status changes—The LAG interface status can change due to several reasons—for example, when member interfaces are deleted, when an interface is down, or when an active LAG member is removed. In such cases, a session scan is triggered and network processor sessions related to the LAG interface are removed. The packet then installs a new network processor session, which could be a fast-forward session, if it qualifies.
- An active member is deactivated or removed from the LAG—In cases when the LAG still has an active member, neither a reroute nor session scan is triggered. The traffic is redistributed on the failed LAG member by monitoring outgoing logical interface status in the SPU.
- A new member is added to the LAG—The network processor session is not updated. A new network processor session is created, which might use the newly added interface or not, depending on the member selection algorithm for the LAG.

Redistribution of Traffic on a Failed LAG Interface

When a LAG member fails, the traffic needs to be redistributed. To redistribute traffic, the system monitors the status of the egress interface in the SPU. When the system detects a failure, it updates the Routing Engine kernel, and passes the physical interface information down to all SPUs. On receiving a copy of the session, the SPU extracts the egress interface index and checks the physical interface information. If the physical interface is down, the SPU uninstalls the session and the ingress network processor deletes the session cache.

For the next ingress packet of the same conversation, the network processor forwards the packet to the SPU to select an active member interface in the LAG as an egress interface. The SPU performs the distribution algorithm to select a new egress interface. A new session with the new egress interface index is installed in the ingress network processor and a new fast-flow path is created.

End-to-End Debugging

For regular flow packets, end-to-end debugging functions are the same as in the non-Express Path mode; packet filter and action items are supported in this flow mode. For traffic that matches Express Path sessions, the end-to-end debugging function supports one packet copy to the host CPU when the filter and the action are both affirmative in the end-to-end search results.

NOTE: End-to-end debugging is not supported on the SRX5K-MPC when Express Path mode is enabled.

Per-Session Statistics CLI

To enable the per-session statistics, copy and paste the following command into the CLI at the **[edit]** hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

set chassis fpc <*fpc-slot*> pic <*pic_slot*> services-offload per-session-statistics

Verify that the **services-offload per-session-statistics** command is enabled.

show configuration chassis

```
user@host> show configuration chassis
fpc 1 {
    pic 1 {
        services-offload {
            per-session-statistics;
        }
    }
 }
```

NOTE: The services-offload per-session-statistics command is not applicable for the SRX5K modular port concentrators when Express Path is configured, because every session has statistics by default.

Use the **show chassis hardware** command to display hardware information.

show chassis fpc pic-status (SRX5600 and SRX5800 devices When Express Path [Services-Offload] is Configured)

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

Slot 0 Online SRX5k DPC 40x 1GE
PIC 0 Online 10x 1GE RichQ
PIC 1 Online 10x 1GE RichQ
PIC 2 Online 10x 1GE RichQ
PIC 3 Online 10x 1GE RichQ
Slot 2 Online SRX5k IOC II
PIC 0 Online 12x 10GE SFP+- np-cache/services-offload
Slot 3 Online SRX5k IOC II
PIC 0 Online 10x 10GE SFP+- np-cache/services-offload
PIC 2 Online 10x 10GE SFP+- np-cache/services-offload
Slot 5 Online SRX5k SPC
PIC 0 Online SPU Cp-Flow
PIC 1 Online SPU Flow
```

Express Path Limitations

The Junos OS Express Path implementation has the following limitations.

- Unsupported features—The following features are not supported with Express Path:
  - Transparent mode is not supported. If transparent mode is configured, a normal (non-Express Path) session is installed.
  - Only multicast sessions with one fan-out are supported. If a multicast session with more than one fan-out exists, a normal session is installed.
  - Only active/passive chassis cluster configuration is supported. Active/active chassis cluster configuration is not supported.
  - Fragmented packets are not supported. If fragmented packets exist, a normal session is installed.

Copyright © 2016, Juniper Networks, Inc.

- When Express Path mode is enabled on an SRX5K-MPC, you might not be able to enable the firewall filter. In general, all processes related to Jflow v5/v8/v9 in an SRX Series device takes place in SPU. When security flow session is of Express Path type, the Jflow configuration will not take effect.
- Class of Service (CoS) on egress interfaces is not supported.
- Configuration of protection against a teardrop (Screen) attack is not supported when Express Path is enabled.
- Asymmetric IOC configuration is not supported in when Express Path is enabled on a device operating in chassis cluster mode.
- Configuring different MTU size values is not supported on SRX5K-MPC when Express Path is enabled on the device.

- Performance drop—The following drops in performance occur when Express Path is enabled:
  - Normal (non-Express Path) sessions—When Express Path is enabled, for normal sessions, the performance can drop by approximately 20 percent for connections per second (CPS) and 15 percent for packets per second (pps) when compared with normal sessions.
  - Express Path sessions—When Express Path is enabled, for fast-forward sessions, the performance can drop by approximately 13 percent for connections per second (CPS).
- Chassis cluster—When the device is operating in chassis cluster mode:
  - If a child link goes down from the LACP-enabled redundant Ethernet interface of an IOC with Express Path enabled on its FPC, all traffic on this link is distributed to other active child links of the interface. If the child link comes up and rejoins the redundant Ethernet interface, then the existing traffic or sessions might not be redistributed over this newly rejoined active child link. New sessions might however traverse through this link.
  - If a new child link is added on the LACP-enabled redundant Ethernet interface of an IOC with Express Path enabled on its FPC, then the existing traffic or sessions might not be redistributed over this new child link. New sessions might however traverse through this link.
- For the normal flow sessions, the **show security flow session** command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. So the **show security flow session** command output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Express Path Support on NP-IOC Card

The NP-IOC card integrates an existing I/O card (IOC) with a Network Processing Card (NPC) in one card with simplified Layer 2 functions in the hardware. This new hardware

Copyright © 2016, Juniper Networks, Inc.

changes the way the interface is interpreted in the system. An NP-IOC can be present in any of the slots, meaning that there also can be an interface at any slot.

NOTE: Each interface in the NP-IOC card can only be attached to the network processor on the NP-IOC card. This fixed attachment setup requires the network processor to manage the interfaces as local or relative interfaces, instead of systemwide global interfaces.

Besides providing physical layer network connections, another function of the NP-IOC card is to distribute packets coming into the physical ports to the Services Processing Units (SPUs) and to forward packets out of the physical ports. For parallel security processing, flow sessions are assigned to multiple SPUs, based on a load balance algorithm. The network processor on the NP-IOC is responsible for directing traffic to the proper SPU based on the session table installed in its local memory.

In Express Path mode, the first packet is processed as is, meaning the packet is forwarded to the central point and the central point assigns an SPU and passes the packet to the SPU. For packets in fast-path, instead of forwarding all packets to the SPU, the network processor forwards the packets to an egress network processor, which can be different from or the same as the ingress network processor.

Express Path Support on SRX5K Modular Port Concentrator

The SRX5K-MPC is a Modular Port Concentrator (MPC) that is supported on the SRX5400, SRX5600, and SRX5800.

The SRX5K-MPC is an interface card with two slots that accept MICs, which add Ethernet ports to your services gateway. An MPC with MICs installed functions in the same way as a regular I/O card (IOC) but allows you to add different types of Ethernet ports to your device.

Each MPC is equipped with Trio chipsets, which perform control functions tailored to the MPC's media type.

When a Trio chipset receives the first packet, the packet is forwarded to an SPU based on the hash value (which is determined by a hash function of the 5 tuples of the session).

If the SPU verifies that the traffic is qualified for Express Path (formerly known as *services loading*), an Express Path session is created on the Trio chipset. If the traffic does not qualify for Express Path, it is forwarded by default hash-based forwarding to SPUs. If an Express Path session is created, the subsequent fast-path packets are processed in the Trio chipset itself.

The Trio chipset performs all the necessary checks to forward the packet, including TTL checking and decreasing, TCP sequence check, NAT translation, and Layer 2 header encapsulation. In addition, the Trio chipset sends a session refresh message to the SPU every second. This message is used to refresh the SPU session, detect the current state of the Trio chip set and update SPU session statistics.

Copyright © 2016, Juniper Networks, Inc.

The session table on the SRX5K-MPC, managed by the SPU, provides the following functions:

- Flow insert or delete
- Flow lookup
- Flow aging
- Flow statistics

The SPU inserts and deletes flow entries in the session table based on policy matching results.

NOTE: Configuring the screen options on an SRX5K-MPC when operating in Express Path mode is the same as when the card is operating in normal mode.

Express Path Support on SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3)

Express Path (formerly known as *services loading*) on the IOC3 is based on processing fast-path packets through the Trio chipset instead of in the SPU to offload some basic firewall functions to the IOC3.

When the Express Path feature is enabled, the IOC3 provides much lower latency, and also supports higher throughput by removing the overload on the SPU. The IOC3 supports both intra-card traffic flow and inter-card traffic flow. To achieve the best latency results, both the ingress port and egress port of a traffic flow need to be on the same XM chip of the IOC3.

The flow table on the IOC3 is managed by the SPU of the flow module. The SPU inserts and deletes flow entries in the flow table based on policy matching results. In the data plane, the IOC3 parses packets, and looks them up in the flow table. If the IOC3 finds a match in the flow table, then it forwards packets based on the instructions given in the flow table. The IOC3 can perform NAT, encapsulate the Level 2 (L2) header, and forward the packets out of the egress interface. The egress interface can be located on the same IOC3 (intra-card case) or on another IOC3 (inter-card case).

NOTE: Flow table lookup in the IOC3 occurs only in ingress. Egress datapath packet handling is the same as supported in the previous release.

When the IOC3 receives the first packet, it does not match any existing fast-forward session. The default hash-based forwarding is performed to send the first packet to the SPU. The SPU then creates the security session. If the SPU finds that the traffic is qualified for fast forwarding, and the related IOC3 supports fast forwarding, it will install fast-forward session to the IOC3. If fast forwarding cannot be applied to the traffic, no session message is sent, and the IOC3 uses the default hash-based forwarding to forward the packets to the SPU.

Copyright © 2016, Juniper Networks, Inc.

In fast-forward IOC3 processing, if a fast-forward session is matched, the packet can be directly forwarded according to the session flow result. The IOC3 takes all the necessary actions, including forwarding the packet, TTL checking and decreasing, NAT translation, L2 header encapsulation and so on.

In addition, the XL chip sends one copy of the forwarding packet to the SPU at every predefined time. This copy is used to refresh the SPU session, detect the current XL chip state, and so on. The SPU consumes this packet and does not forward it, because the real packet has been processed and transmitted.

Expres Path support on IOC3 is illustrated in Figure 88, Figure 89, and Figure 90.

Figure 88: IOC3 Intra-PFE Express Path

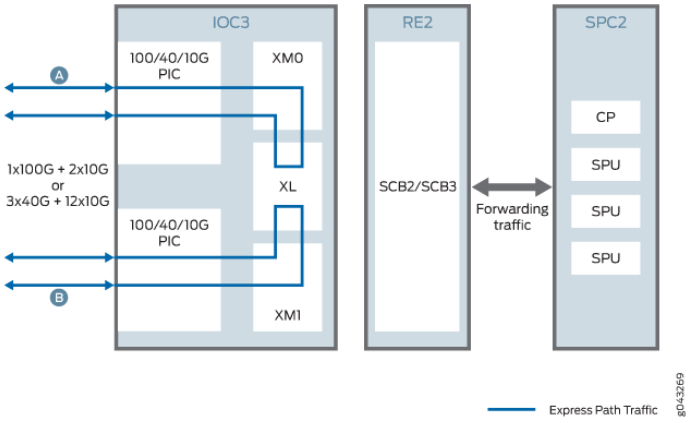

Figure 89: IOC3 Inter-PFE Express Path

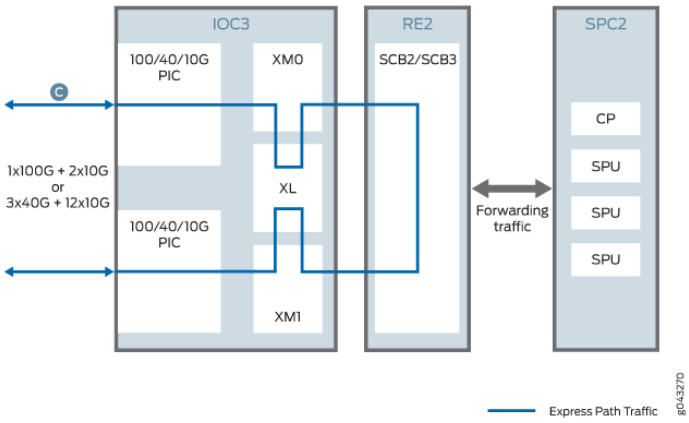

Copyright © 2016, Juniper Networks, Inc.

Figure 90: Inter-IOC3 Express Path

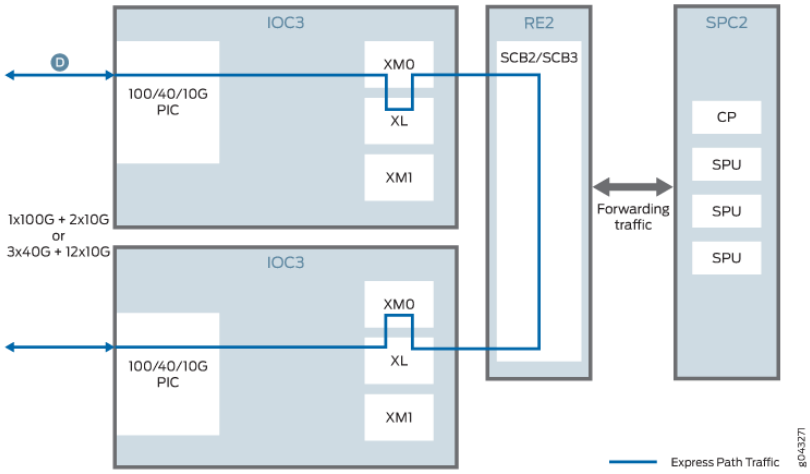

IPv6 Flow in Express Path Mode for IOC2 and IOC3

IPv6 traffic is supported on SRX5000 line devices with the SRX5K-MPC (IOC2), the SRX5K-MPC3-100G10G (IOC3), or the SRX5K-MPC3-40G10G (IOC3) in Express Path mode.

On SRX5000 line devices, Express Path for IPV6 traffic is not supported on legacy IOC cards. However, IPv6 regular flow mode is supported on legacy IOCs.

When the first packet arrives at the interface, the network processor forwards it to the SPU. If the SPU verifies that the traffic is qualified for Express Path, an Express Path session is created on the network processor. If the traffic does not qualify for Express Path, a normal session is created on the network processor. If an Express Path session is created, the subsequent fast-path packets are processed in the network processor itself.

When an Express Path session is created on the network processor, subsequent packets of the flow match the session on the network processors. The network processor then processes and forwards the packet. The network processor also handles additional processing such as TCP sequence check, time-to-live (TTL) processing, and Layer 2 header translation.

The following features are not supported in Express Path mode:

- IPv6 NAT
- Transparent mode
- Configuring different MTU size values
- Class of Service (CoS) on egress interfaces

Copyright © 2016, Juniper Networks, Inc.

Note the following limitations:

- Express Path sessions for IPv6 traffic traversing on legacy IOC cards is not supported. IOC2 and IOC3 does not support IPv6 traffic in Express Path sessions when traffic traverse between legacy IOC and IOC2 or IOC3. Normal IPv6 traffic is still supported in this scenario.
- Express Path is not supported in IPsec VPN, ALG, and IDP configurations. Normal flow sessions will be used in these scenarios.
- Express Path is not supported in logical system configurations, because logical systems use the sessions installed from the normal flow sessions.
- Multicast sessions with multiple fan-outs are not supported in Express Path mode. A normal session is installed if a multicast session with more than one fan-out exists.
- Active/active chassis cluster configuration is not supported in Express Path mode. Only active/passive chassis cluster configuration is supported.
- A redundant Ethernet interface must contain both child interfaces from the same IOC type. For example, if one child link is from 10-Gigabit Ethernet on IOC2, the second child link should also be from the IOC2. Similarly, both child interfaces can be from IOC3. Configuring child interfaces by mixing the links from both IOC2 and IOC3 is not supported.
- Fragmented packets are not supported in Express path mode. IPv6 fragmented packets will be directed to the SPU and hence they will not be qualified for Express Path.

See "Example: Configuring Express Path on an SRX5000 Line Device with IOC3 for IPv6 Traffic" on page 1742 for details on configuring Express Path for IPv6 traffic.

Disabling TCP Packet Security Checks

On an SRX Series device, you can disable security checks on TCP packets to ensure interoperability with hosts and devices with faulty TCP implementations.

The **no-sequence-check** option disables TCP sequence checks. It also increases the throughput.

The **set security flow tcp-session no-sequence-check** command disables the TCP sequence checks on all TCP sessions in default or hash-based modes. This command is also supported in Express Path mode.

IPv6 Flow in Express Path Mode

IPv6 traffic is supported on SRX5000 line devices with SRX5K-MPC (IOC2), the SRX5K-MPC3-100G10G (IOC3), or the SRX5K-MPC3-40G10G (IOC3). All IPv6 traffic is handled in regular flow mode, meaning that packets are forwarded to the SPU for flow processing. Egress IPv6 traffic is also forwarded from the SPU to the network processor, and then the network processor handles this traffic as regular flow traffic in the egress path.

When the first packet arrives at the interface, the network processor forwards it to the SPU. If the SPU verifies that the traffic is qualified for Express Path, an Express Path

Copyright © 2016, Juniper Networks, Inc.

session is created on the network processor. If the traffic does not qualify for Express Path, a normal session is created on the network processor. If an Express Path session is created, the subsequent fast-path packets are processed in the network processor itself.

When an Express Path session is created on the network processor, subsequent packets of the flow match the session on the network processors. The network processor then processes and forwards the packet. The network processor also handles additional processing such as TCP sequence check, time to live (TTL) processing, Network Address Translation (NAT), and Layer 2 header translation.

NOTE: On the SRX5000 line devices, the Express Path sessions for IPv6 traffic that traverse between legacy IOC cards and not supported.

The Express Path sessions for IPv6 traffic traversing only on legacy IOC cards or only on the SRX5K-MPC (IOC2), the SRX5K-MPC3-100G10G (IOC3), or the SRX5K-MPC3-40G10G (IOC3) are supported.

Related Documentation

- Example: Configuring Low Latency on page 1746
- Understanding the Express Path Solution on page 1731
- Enabling and Disabling Express Path on page 1732

Copyright © 2016, Juniper Networks, Inc.

Understanding the Express Path Solution

The high-end SRX Series devices have long packet-processing latency because the packets are processed through the Services Processing Unit (SPU) and through several stages of buffers in the data path.

This feature introduces a local forwarding solution where the fast-path packets are processed by the network processor on the I/O Card (IOC), without going through the switch fabric or the SPU. This solution reduces the packet-processing latency.

The behavior of the network processor in different scenarios is as follows:

- **First-path flow**—The first-path flow is the same as the current network processor flow process. When the first packet arrives at the network processor, the network processor parses the TCP or the UDP packet to extract a 5-tuple key and then performs session lookup in the flow table. The network processor then forwards the first packet to the central point. The central point cannot find a match at this time, because this is the first packet. The central point and the SPU create a session and match it against user-configured policies to determine if the session is a normal session or a services-offload session.

  If the user has specified the session to be handled with Express Path, the SPU creates a session entry in the network processor flow table, enabling the services-offload flag in the session entry table; otherwise, the SPU creates a normal session entry in the network processor without the services-offload flag.

- **Fast-path flow**—After the session entry is created in the network processor, subsequent packets of the session will match the session entry table.
  - If the services-offload flag is not set, then the network processor forwards the packet to the SPU specified in the session entry table. The packet goes through the normal flow process.
  - If the network processor finds the services-offload flag in the session entry table, it will process the packet locally and send the packet out directly.

NOTE: **The fast-forwarding function on the network processor supports one-fanout multicast sessions. The egress port in the session must also be associated with the same network processor of the ingress port. All other multicast cases need to be handled as normal sessions.**

- **NAT process**—The SPU is responsible for mapping between the internal IP address or port and the external IP addres or port. When the first packet of the session arrives, the SPU allocates the IP address or port mapping and stores the information in the network processor session entry. The network processor does the actual packet modification if the NAT flag is set.
- **Session age-out**—To improve traffic throughput for services-offload sessions, a copy of a packet is sent to the SPU at every predefined time period to reduce the packet processing demand on the SPU. To limit the number of packet copies sent to the SPU,

Copyright © 2016, Juniper Networks, Inc.

a timestamp is implemented for each services-offload session. This enables the network processor to calculate the elapsed time since the last session match. If the elapsed time is greater then the predefined time period, then the network processor sends a copy of the packet to the SPU, and updates the session timestamp.

- **Session termination and deletion**—If the network processor receives an IP packet with a FIN (finished data) or a RST (reset connection) flag, it forwards the packet to the SPU. The SPU then deletes the session. The network processor does not change session status even after the FIN or RST packet is received. It continues to receive and process any packets during the transit session.

Related Documentation

- Enabling and Disabling Express Path on page 1732

Enabling and Disabling Express Path

NOTE: Prior to Junos OS Release 12.3X48-D10, Express Path was a licensed software feature. Starting with Junos OS Release 12.3X48-D10, the Express Path license is no longer required to enable this functionality. Your previously acquired Express Path license will not be effective anymore.

You can enable Express Path mode as follows:

- Set the Express Path mode on the selected card.
- Reboot the device containing the Express Path network processor to load the Express Path firmware image on the network processors.
- Configure a policy to define the traffic that should take fast-path.

NOTE: For SRX5600 and SRX5800 devices, enable Express Path on the IOC PIC. . For SRX5K modular port concentrators (SRX5K-MPC), SRX5K-MPC3-100G10G (IOC3), and SRX5K-MPC3-40G10G (IOC3) on SRX5000 line devices, enable Express Path on the FPCs.

NOTE: When device is operating in chassis cluster mode, you need to reboot both the nodes when changing FPC(s) to Express Path mode.

To configure the Express Path mode:

- For configuring Express Path on an SRX5000 line device, use the **set chassis fpc** ***fpc-number*** **pic** ***pic-number*** **services-offload** command.
- For configuring Express Path on an SRX5000 line device with Modular Port Concentrator (MPC), enable NP cache on the IOC using the **set chassis fpc** ***fpc-number*** **np-cache**

Copyright © 2016, Juniper Networks, Inc.

command. Then configure the security policy to determine if the session is for Express Path.

NOTE: The **set chassis fpc** *fpc-number* **services-offload** command is deprecated.

- To disable Express Path on an SRX5000 line device with Modular Port Concentrator (MPC), use the **delete chassis fpc** *fpc-number* **np-cache** command.

NOTE: The **delete chassis fpc** *fpc-number* **services-offload** command is deprecated.

NOTE: You need to reboot the device when you disable Express Path.

System log files are stored locally on the device in the default **/var/log/security** directory. You can use system log files to retrieve information about the Express Path configuration changes.

Related Documentation

- Understanding the Express Path Solution on page 1731
- Example: Configuring an IOC on SRX5000 Line Devices to Support Express Path on page 1735
- Example: Configuring an SRX5K-MPC on an SRX5000 Line Device to Support Express Path on page 1737

Example: Enabling Express Path in Security Policies

This example shows how to enable Express Path (formerly known as *services offloading*) in security policies.

- Requirements on page 1733
- Overview on page 1733
- Configuration on page 1734
- Verification on page 1735

Requirements

Before you begin, see "Express Path Overview" on page 1717.

Overview

In this example, you enable Express Path in security policies to specify whether the traffic qualifies for Express Path.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

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

```
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  source-address 192.0.2.2
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  destination-address 198.51.100.10
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  application junos-http
set security policies from-zone untrust to-zone trust policy services-offload-pol1 then
  permit services-offload
```

Step-by-Step Procedure

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

To configure Express Path in policies:

1. Configure a policy to process the traffic that goes to the HTTP static ports.

   ```
   [edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
   user@host# set match source-address 192.0.2.2
   user@host# set match destination-address 198.51.100.10
   user@host# set match application junos-http
   ```

2. Enable Express Path in the security policy.

   ```
   [edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
   user@host# set then permit services-offload
   ```

Results

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

```
[edit]
user@host# show security policies
from-zone untrust to-zone trust {
   policy services-offload-pol1 {
   match {
   source-address 192.0.2.2;
      destination-address 198.51.100.10;
      application junos-http;
      }
      then {
         permit {
      services-offload;
      }
   }
   }
}
```


Copyright © 2016, Juniper Networks, Inc.

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

Verification

Confirm that the configuration is working properly.

Verifying Express Path in Policies

Purpose Verify that Express Path is enabled.

Action From operational mode, enter the **show security policies** command.

Related Documentation

- Express Path Overview on page 1717

Example: Configuring an IOC on SRX5000 Line Devices to Support Express Path

This example shows how to configure an IOC on SRX5000 line of devices to support Express Path (formerly known as *services offloading*).

- Requirements on page 1735
- Overview on page 1735
- Configuration on page 1735
- Verification on page 1736

Requirements

Before you begin, see "Express Path Overview" on page 1717.

Overview

In this example, you configure the IOC on SRX5000 line devices to perform Express Path.

Configuration

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

```
set chassis fpc 3 pic 0 services-offload
```

NOTE: For SRX5000 line devices, the IOC slot number is 3.

Step-by-Step Procedure To configure the IOC you need to run the following commands:

1. Set the services offload mode on the IOC.

   ```
   [edit]
   ```

```
user@host# set chassis fpc 3 pic 0 services-offload
```

2. Commit the configuration.

```
[edit]
user@host# commit
warning: System restart is required after fpc 3 pic 0 changed to
services-offload mode.
commit complete
```

3. Reboot the device.

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

```
[edit]
user@host> show chassis fpc pic-status
Slot 0   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
Slot 1   Online       SRX5k FIOC
  PIC 0  Online       4x 10GE XFP
  PIC 1  Online       16x 1GE SFP
Slot 3   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ- services-offload
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 5   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
```

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

Verification

Verifying the Configuration of IOC for Express Path

Purpose Verify that the IOC was configured properly for Express Path.

Action From operational mode, enter the **show chassis fpc pic-status** command.

Related Documentation

- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring an SRX5K-MPC on an SRX5000 Line Device to Support Express Path

This example shows how to configure an SRX5K-MPC on an SRX5000 line device to support Express Path (formerly known as *services offloading*).

- Requirements on page 1737
- Overview on page 1737
- Configuration on page 1737
- Verification on page 1738

Requirements

This example uses the following hardware and software components:

- One SRX5000 line device with an SRX5K-MPC
- Junos OS Release 12.3X48 or later for SRX Series devices

Before you begin, see "Express Path Overview" on page 1717.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the SRX5K-MPC on an SRX5000 line device to perform NP cache and Express Path.

Configuration

CLI Quick Configuration

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

```
set chassis fpc 2 np-cache
set chassis fpc 3 np-cache
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  source-address 198.51.100.20
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  destination-address 192.0.2.11
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  application junos-http
set security policies from-zone untrust to-zone trust policy services-offload-pol1 then
  permit services-offload
```

Step-by-Step Procedure

To configure an SRX5K-MPC on an SRX5000 line device to perform Express Path:

1. Set NP cache mode on the SRX5K-MPC on FPC 1 and FPC 2.

```
[edit]
user@host# set chassis fpc 2 np-cache
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set chassis fpc 3 np-cache
```

2. Configure a policy to process the traffic that goes to the HTTP static ports.

```
[edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
user@host# set match source-address 198.51.100.20
user@host# set match destination-address 192.0.2.11
user@host# set match application junos-http
```

3. Enable Express Path in the security policy.

```
[edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
user@host# set then permit services-offload
```

4. Commit the configuration.

```
[edit]
user@host# commit
warning: System or cluster nodes need to reboot after fpc 3 changed to
np-cache mode.
```

5. Reboot the device.

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

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

```
...
  fpc 2 {
    np-cache;
  }
  fpc 3 {
    np-cache;
  }
...
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Configuration of an SRX5K-MPC for Express Path

Purpose Verify that the SRX5K-MPC was configured properly for Express Path.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k IOC II
  PIC 0  Online       2x 40GE QSFP+- np-cache/services-offload
Slot 3   Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

Meaning The output provides the status of PICs with Express Path enabled on them.

Related Documentation

- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733
- Enabling and Disabling Express Path on page 1732

Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path

This example shows how to configure an SRX5K-MPC3-100G10G (IOC3) or an SRX5K-MPC3-40G10G (IOC3) on an SRX5000 line device to support Express Path (formerly known as *services offloading*).

- Requirements on page 1739
- Overview on page 1739
- Configuration on page 1739
- Verification on page 1741

Requirements

This example uses the following hardware and software components:

- One SRX5000 line device with an SRX5K-MPC3-40G10G (IOC3)
- Junos OS Release 15.1X49-D10 or later for SRX Series devices

Before you begin, see "Express Path Overview" on page 1717.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an SRX5K-MPC3-40G10G (IOC3) on an SRX5000 line device to perform Express Path.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis fpc 4 np-cache
set chassis fpc 5 np-cache
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  source-address 192.0.2.41
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  destination-address 203.0.113.10
set security policies from-zone untrust to-zone trust policy services-offload-pol1 match
  application junos-http
set security policies from-zone untrust to-zone trust policy services-offload-pol1 then
  permit services-offload
```

Step-by-Step Procedure

To configure an SRX5K-MPC3-40G10G (IOC3) on an SRX5000 line device to perform Express Path:

1. Set the Express Path mode on the SRX5K-MPC3 on FPC 4 and FPC 5.

   ```
   [edit]
   user@host# set chassis fpc 4 np-cache
   user@host# set chassis fpc 5 np-cache
   ```

2. Configure a policy to process the traffic that goes to the HTTP static ports.

   ```
   [edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
   user@host# set match source-address 192.0.2.41
   user@host# set match destination-address 203.0.113.10
   user@host# set match application junos-http
   ```

3. Enable Express Path in the security policy.

   ```
   [edit security policies from-zone untrust to-zone trust policy services-offload-pol1]
   user@host# set then permit services-offload
   ```

4. Commit the configuration.

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

   ```
   warning: System or cluster nodes need to reboot after fpc 3 changed to
   np-cache mode.
   ```

5. Reboot the device.

Results

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

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

```
...
  fpc 4{
    services-offload;
  }
  fpc 5{
```


Copyright © 2016, Juniper Networks, Inc.

```
    services-offload;
  }
...
...
  from-zone <fzone> to-zone <tzone> {
  policy <policy-name> {
    match {
      <match-tuples>
    }
    then {
      action (
      permit {
        ...
        services-offload
        ^^^^^^^^^^^^^^^^
      }
      reject
      deny
      log
      );
    }
    scheduler-name <scheduler-name>;
  }
  ...
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Configuration of an SRX5K-MPC3 (IOC3) for Express Path

Purpose Verify that the SRX5K-MPC3-40G10G (IOC3) was configured properly for Express Path.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Meaning The output provides the status of PICs with Express Path enabled on them.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733
- Enabling and Disabling Express Path on page 1732
- Understanding Session Cache on page 1713

Example: Configuring Express Path on an SRX5000 Line Device with IOC3 for IPv6 Traffic

This example shows how to configure Express Path (formerly known as *services offloading*) on an SRX5K-MPC3-100G10G (IOC3) or an SRX5K-MPC3-40G10G (IOC3) on an SRX5000 line device for IPv6 traffic.

- Requirements on page 1742
- Overview on page 1742
- Configuration on page 1742
- Verification on page 1745

Requirements

This example uses the following hardware and software components:

- One SRX5000 line device with an IOC3 card
- Junos OS Release 15.1X49-D40 or later for SRX Series devices

Before you begin, see "Express Path Overview" on page 1717.

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure Express Path on IOC3 on an SRX5000 line device for IPv6 traffic.

You configure two interfaces on IOC3 card and assign IPv6 addresses to them. Then you enable flow-based processing for IPv6 traffic. Next, you set up zones and add interfaces to them. Then you provide communication between the two different zones by configuring a security policy to allow traffic between two zones. You also enable Express Path in security policies to specify whether the traffic qualifies for Express Path.

Configuration

CLI Quick Configuration

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

```
set interfaces et-2/1/0 unit 0 family inet6 address 2013:1:2:3:4:5:6:1/64
set interfaces et-2/3/0 unit 0 family inet6 address 2014:1:2:3:4:5:6:1/64
```


Copyright © 2016, Juniper Networks, Inc.

```
set security forwarding-options family inet mode flow-based
set security forwarding-options family inet6 mode flow-based
set zones security-zone zone-1 host-inbound-traffic system-services all
set zones security-zone zone-1 host-inbound-traffic protocols all
set zones security-zone zone-1 interfaces et-2/1/0.0
set zones security-zone zone-2 host-inbound-traffic system-services all
set zones security-zone zone-2 host-inbound-traffic protocols all
set zones security-zone zone-2 interfaces et-2/3/0.0
security policies from-zone zone-2 to-zone zone-1 policy express-path-policy-2 match
  source-address any
security policies from-zone zone-2 to-zone zone-1 policy express-path-policy-2 match
  destination-address any
security policies from-zone zone-2 to-zone zone-1 policy express-path-policy-2 match
  application any
security policies from-zone zone-2 to-zone zone-1 policy express-path-policy-2 then
  permit then permit services-offload
security policies from-zone zone-1 to-zone zone-2 policy express-path-policy-1 match
  source-address any
security policies from-zone zone-1 to-zone zone-2 policy express-path-policy-1 match
  destination-address any
security policies from-zone zone-1 to-zone zone-2 policy express-path-policy-1 match
  application any
security policies from-zone zone-1 to-zone zone-2 policy express-path-policy-1 then permit
  then permit services-offload
set chassis fpc 2 np-cache
```

Step-by-Step Procedure

To configure an SRX5K-MPC3-40G10G (IOC3) on an SRX5000 line device to perform Express Path:

1. Configure Ethernet interface and assign an IPv6 address to it.

   ```
   [edit]
   user@host# set interfaces et-2/1/0 unit 0 family inet6 address 2013:1:2:3:4:5:6:1/64
   user@host# set interfaces et-2/3/0 unit 0 family inet6 address 2014:1:2:3:4:5:6:1/64
   ```

2. Enable flow-based processing for IPv6 traffic.

   ```
   [edit]
   user@host# set security forwarding-options family inet mode flow-based
   user@host# set security forwarding-options family inet6 mode flow-based
   ```

3. Configure security zones and add interfaces and allow all system services and interfaces. Configure a security zone and specify the types of traffic and protocols that are allowed on interface et-2/1/0.0.

   ```
   [edit]
   user@host# set zones security-zone zone-1 host-inbound-traffic system-services
     all
   user@host# set zones security-zone zone-1 host-inbound-traffic protocols all
   user@host# set zones security-zone zone-1 interfaces et-2/1/0.0
   ```

4. Configure security zones and add interfaces and allow all system services and interfaces. Configure a security zone and specify the types of traffic and protocols that are allowed on interface et-2/3/0.0.

   ```
   [edit]
   user@host# set zones security-zone zone-2 host-inbound-traffic system-services
     all
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set zones security-zone zone-2 host-inbound-traffic protocols all
user@host# set zones security-zone zone-2 interfaces et-2/3/0.0
```

5. Create a policy and specify the match criteria for that policy. The match criteria specifies that the device can allow traffic from any source, to any destination, and on any application. Enable Express Path in the security policy.

NOTE: You can specify the wildcard any-ipv6 for the source and destination address match criteria to include only IPv6 addresses. Specifying any option for the source and destination address match criteria to include both IPv4 and y IPv6 addresses.

```
[edit security policies from-zone zone-2 to-zone zone-1 policy express-path-policy-2]
user@host# set security policies from-zone zone-2 to-zone zone-1 policy
  express-path-policy-2 match source-address any
user@host# set security policies from-zone zone-2 to-zone zone-1 policy
  express-path-policy-2 match destination-address any
user@host# security policies from-zone zone-2 to-zone zone-1 policy
  express-path-policy-2 match application any
user@host# security policies from-zone zone-2 to-zone zone-1 policy
  express-path-policy-2 then permit services-offload

[edit]
user@host# set security policies from-zone zone-1 to-zone zone-2 policy
  express-path-policy-1 match source-address any
user@host# set security policies from-zone zone-1 to-zone zone-2 policy
  express-path-policy-1 match destination-address any
user@host# set security policies from-zone zone-1 to-zone zone-2 policy
  express-path-policy-1 match application any
user@host# set security policies from-zone zone-1 to-zone zone-2 policy
  express-path-policy-1 then permit services-offload
```

6. Set the Express Path mode on IOC3.

```
[edit]
user@host# set chassis fpc 2 np-cache
```

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

```
from-zone zone-1 to-zone zone-2 {
  policy express-path-policy--1 {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit {
        services-offload;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  from-zone express-path-policy--2 {
    policy policy-2 {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          services-offload;
        }
      }
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Configuration of an SRX5K-MPC3 (IOC3) for Express Path

Purpose Verify that the IOC3 was configured properly for Express Path.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k IOC3 2CGE+4XGE
  PIC 0  Online       2x 10GE SFP+- np-cache/services-offload
  PIC 1  Online       1x 100GE CFP2- np-cache/services-offload
  PIC 2  Online       2x 10GE SFP+- np-cache/services-offload
  PIC 3  Online       1x 100GE CFP2- np-cache/services-offload
Slot 3   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Offline      12x 10GE SFP+
  PIC 1  Offline      12x 10GE SFP+
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 4   Offline      SRX5k IOC3 24XGE+6XLG
```

Meaning The output provides the status of PICs with Express Path enabled on them.

Verifying All Active Sessions on the Device

Purpose Display information about all currently active Express Path security sessions on the device.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security flow session services-offload** command.

```
Flow Sessions on FPC1 PIC1:

Session ID: 50000002, Policy name: express-path-policy-2/5, Timeout: 60, Valid
  In: 2016::100/63 --> 2015::100/63;udp, If: et-2/3/0.0, Conn ID: 0x0, Pkts: 181
29505, Bytes: 1740432530, CP Session ID: 50000002
  Out: 2015::100/63 --> 2016::100/63;udp, If: et-2/1/0.0, Conn ID: 0x0, Pkts: 18
129505, Bytes: 1740432530, CP Session ID: 50000002
Total sessions: 1
```

Meaning The output provides the policy details for sessions on which Express Path was enabled.

Related Documentation

- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733
- Enabling and Disabling Express Path on page 1732
- Understanding Session Cache on page 1713

Example: Configuring Low Latency

The low latency feature allows you to configure the mode of the network processor's traffic manager (TM) on the egress path. If low latency is enabled, the network processor is initialized without the traffic manager, thus reducing the overall latency in the Express Path (formerly known as *services offloading*).

NOTE: Because all SRX Series CoS functions are supported by the traffic manager, CoS functions are not supported when low latency is enabled.

Low latency reduces the total NPC integrated with an existing IOC (NP-IOC) latency by 0.7 us. This latency reduction brings the NP-IOC card total latency to 8.7 us. The low-latency feature is supported for intra-NP-IOC card traffic only; it is not applicable to inter-NP traffic.

In the low-latency mode, the network processor does not have an egress buffer at the traffic manager. Packets are delivered directly to the system packet interface (SPI) for the field-programmable gate array (FPGA) to process.

NOTE: The low latency feature is only applicable to the NP-IOC card.

- Requirements on page 1747
- Overview on page 1747
- Configuration on page 1747
- Verification on page 1748

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, see "Express Path Overview" on page 1717.

This example uses the following software and hardware components:

- Junos OS Release 12.1X44-D10
- One SRX Series device
- One Services Processing Card (SPC)

Overview

In this example, you configure the network processor for low latency mode.

Configuration

CLI Quick Configuration

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

```
set chassis fpc 7 pic 0 services-offload low-latency
```

Step-by-Step Procedure

To enable low-latency mode:

1. Enable the Express Path mode on the NP-IOC.

   ```
   [edit]
   user@host# set chassis fpc 7 pic 0 services-offload
   ```

2. Enable low latency.

   ```
   [edit]
   user@host# low-latency
   ```

3. Commit the configuration.

   ```
   [edit]
   user@host# commit
   warning: System restart is required after fpc 7 pic 0 changed to
   services-offload mode.
   commit complete
   ```

4. Reboot the device.

Results

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

```
[edit]
user@host> show configuration chassis
fpc 7 {
     pic 0 {
         services-offload {
             low-latency;
         }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

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

Verification

Verifying Low Latency Configuration

Purpose Verify that low-latency was enabled.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
root@kg04> show chassis fpc pic-status
  Slot 0   Online       SRX3k SFB 12GE
    PIC 0  Online       8x 1GE-TX 4x 1GE-SFP
  Slot 1   Online       SRX3k 2x10GE XFP
    PIC 0  Online       2x 10GE-XFP
  Slot 2   Online       SRX3k SPC
    PIC 0  Online       SPU Cp-Flow
  Slot 7   Online       SRX1k3k 2x10GE NP-IOC
    PIC 0  Online       2x 10GE-SFP+- services-offload low-latency
```

Related Documentation

- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733

Copyright © 2016, Juniper Networks, Inc.

PART 25

Managing Flow-Based Processing for IPv6

- Enabling IPv6 Flow-Based Processing on page 1751
- Managing IPv6 Packets on page 1763
- Configuring IPv6 Dual-Stack on page 1773

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 87

Enabling IPv6 Flow-Based Processing

- IPv6 Advanced Flow on page 1751
- Understanding Sessions for IPv6 Flows on page 1753
- Understanding IPv6 Flow Processing on High-End SRX Series Devices on page 1753
- Enabling Flow-Based Processing for IPv6 Traffic on page 1756
- Using Filters to Display IPv6 Session and Flow Information for SRX Series Services Gateways on page 1757

IPv6 Advanced Flow

IPv6 advanced flow adds IPv6 support for firewall, NAT, NAT-PT, multicast (local link and transit), IPsec, IDP, JSF framework, TCP Proxy, and Session manager on SRX Series devices. MIBs are not used in the IPv6 flow.

In order to avoid the impact on the current IPv4 environment, IPv6 security is used. If IPv6 security is enabled, extended sessions and gates are allocated. The existing address fields and gates are used to store the index of extended sessions or gates. If IPv6 security is disabled, IPv6 security-related resources are not allocated.

New logs are used for IPv6 flow traffic to prevent impact on performance in the existing IPv4 system.

The behavior and implementation of the IPv6 advanced flow are the same as those of IPv4 in most cases.

Some of the differences are explained below:

- **Header Parse** IPv6 advanced flow stops parsing the headers and interprets the packet as the corresponding protocol packet if it encounters the following extension headers:
  - TCP/UDP
  - ESP/AH
  - ICMPv6

  IPv6 advanced flow continues parsing headers if it encounters the following extension headers:
  - Hop-by-Hop

Copyright © 2016, Juniper Networks, Inc.

- Routing and Destination, Fragment

  IPv6 advanced flow interprets the packets as an unknown protocol packet if it encounters the extension header **No Next Header**

- **Sanity Checks** IPv6 advanced flow supports the following sanity checks:
  - TCP length
  - UDP length
  - Hop-by-hop
  - IP data length error
  - Layer 3 sanity checks (for example, IP version and IP length)
  - **ICMPv6 Packets** In IPv6 advanced flow, the ICMPv6 packets share the same behavior as normal IPv6 traffic with the following exceptions:
    - Embedded ICMPv6 packet
    - Path MTU message
- **Host Inbound and Outbound Traffic** IPv6 advanced flow supports all route and management protocols running on the Routing Engine (RE), including OSPF v3, RIPng, Telnet, and SSH. Note that no flow label is used in the flow.
- **Tunnel Traffic** IPv6 advanced flow supports the following tunnel types:
  - IPv4 IP-IP
  - IPv4 GRE
  - IPv4 IPsec
  - Dual-stack lite
- **Events and Logs** The following logs are for IPv6-related flow traffic:
  - RT_FLOW_IPVX_SESSION_DENY
  - RT_FLOW_IPVX_SESSION_CREATE
  - RT_FLOW_IPVX_SESSION_CLOSE

The implementations of sessions, gates, ip-actions, processing of multithread, distribution, locking, synchronization, serialization, ordering, packet queuing, asynchronous messaging, IKE traffic issues, sanity check, and queues for IPv6 are similar to IPv4 implementations.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- About IPv6 Packet Header Verification Performed by the Flow Module for SRX Series Devices on page 1767

Copyright © 2016, Juniper Networks, Inc.

Understanding Sessions for IPv6 Flows

This topic gives an overview of flow-based sessions.

Most packet processing occurs in the context of a flow, including management of policies, zones, and most screens. A session is created for the first packet of a flow for the following purposes:

- To store most of the security measures to be applied to the packets of the flow.
- To cache information about the state of the flow. For example, logging and counting information for a flow is cached in its session. (Also, some stateful firewall screens rely on threshold values that pertain to individual sessions or across all sessions.)
- To allocate resources required for features for the flow.
- To provide a framework for features such as Application Layer Gateways (ALGs).

Related Documentation

- Understanding IPv6 Flow Processing on High-End SRX Series Devices on page 1753
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Understanding IPv6 Flow Processing on High-End SRX Series Devices

This topic introduces the architecture for the high-end SRX Series devices and uses it as a model to explain IP version 6 (IPv6) processing. Flow processing is similar on branch SRX Series devices

High-end SRX Series Services Gateway devices include I/O cards (IOCs) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. These processing units have different responsibilities.

- A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. An NPU processes packets discretely and performs basic flow management functions.

  When an IPv6 packet arrives at an IOC, the packet flow process begins.

  - The NPU performs the following IPv6 sanity checks for the packet:
    - For the IPv6 basic header, it performs the following header checks:
      - Version. It verifies that the header specifies IPv6 for the version.
      - Payload length. It checks the payload length to ensure that the combined length of the IPv6 packet and the Layer 2 header is shorter than the Layer 2 frame length.
      - Hop limit. It checks to ensure that the hop limit does not specify 0 (zero).
      - Address checks. It checks to ensure that the source IP address does not specify ::0 or FF::00 and that the destination IP address does not specify ::0 or ::1.
    - The NPU performs IPv6 extension header checks, including the following:

Copyright © 2016, Juniper Networks, Inc.

- Hop-by-hop options. It verifies that this is the first extension header to follow the IPv6 basic header.
- Routing extension. It verifies that there is only one routing extension header.
- Destination options. It verifies that no more than two destination options extension headers are included.
- Fragment. It verifies that there is only one fragment header.

NOTE: The NPU treats any other extension header as a Layer 4 header.

- The NPU performs Layer 4 TCP, UDP, and ICMP6 protocol checks, including the following:
  - UDP. It checks to ensure that IP Payload Length packets, other than a first-fragment packet, are at least 8 bytes long.
  - TCP. It checks to ensure that IP Payload Length packets, other than a first-fragment packet, are at least 20 bytes long.
  - ICMPv6. It checks to ensure that IP Payload Length packets, other than a first-fragment packet, are at least 8 bytes long.

- If the packet specifies a TCP or a UDP protocol, the NPU creates a tuple from the packet header data using the following information:
  - Source IP address
  - Destination IP address
  - Source port
  - Destination port
  - Protocol
  - Virtual router identifier (VRID)

    The device looks up the VRID from a VRID table.

- For Internet Control Message Protocol version 6 (ICMPv6) packets, the tuple contains the same information as used for the TCP and the UDP search key, except for the source and destination port fields. The source and destination port fields are replaced with the following information extracted from the ICMPv6 packet:
  - For ICMP error packets: The pattern "0x00010001"
  - For ICMP information packets: The type, or code, field identifier

- For packets with an Authentication Header (AH) or an Encapsulating Security Payload (ESP) header, the search key is the same as that used for the TCP and the UDP tuple, except for the source and destination port fields. In this case, the security parameter index (SPI) field value is used instead of the source and destination ports. For Encapsulating Security Payload (ESP) header and Authentication Header (AH), before enhancements to the cenral point architecture it is hashed by the 3-tuple and

Copyright © 2016, Juniper Networks, Inc.

the security parameter index (SPI) field, after enhancements to the cenral point architecture it is hashed by an IP pair.

- If a session exists for the packet's flow, the NPU sends the packet to the SPU that manages the session.
- If a matching session does not exist,
  - The NPU sends the packet information to the central point, which creates a pending session.
  - The central point selects an SPU to process the packet and create sessions for it.
  - The SPU then sends session creation messages to the central point and the ingress and egress NPUs, directing them to create a session for the packet flow.

- A central point, which can run on a dedicated SPU, or share the resources of one if there is only one SPU. A central point takes care of arbitration and allocation of resources, and it distributes sessions in an intelligent way. The central point assigns an SPU to be used for a particular session when the SPU processes the first packet of its flow.
  - For SRX5000 line devices, the central point architecture is divided into two modules—the application central point and the distributed central point (DCP). The App-CP is responsible for global resource management and loading balancing, while DCP is responsible for traffic identification (global session matching). The App-CP functionality runs on the dedicated central point SPU, while the DCP functionality is distributed to the rest of the SPUs.

- One or more SPUs that run on a Services Processing Card (SPC). All flow-based services for a packet are executed on a single SPU, within the context of a session that is set up for the packet flow.

  The SPC for SRX5000 line devices has two SPUs.

  Several SPCs can be installed in a chassis.

  Primarily, an SPU performs the following tasks:

  - It manages the session and applies security features and other services to the packet.
  - It applies packet-based stateless firewall filters, classifiers, and traffic shapers.
  - If a session does not already exist for a packet, the SPU sends a request message to the NPU that performed the search for the packet's session, to direct it to add a session for it.

These discrete, cooperating parts of the system store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

Related Documentation

- Understanding Sessions for IPv6 Flows on page 1753
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Copyright © 2016, Juniper Networks, Inc.

Enabling Flow-Based Processing for IPv6 Traffic

By default, the SRX Series device drops IP version 6 (IPv6) traffic. To enable processing by security features such as zones, screens, and firewall policies, you must enable flow-based forwarding for IPv6 traffic.

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

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

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

1. Use the **set** command to change the forwarding option mode for IPv6 to flow-based.

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

2. Use the **show** command to review your configuration.

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

3. Check your changes to the configuration before committing.

   ```
   [edit]
   user@host# commit check
   warning: You have enabled/disabled inet6 flow.
   You must reboot the system for your change to take effect.
   If you have deployed a cluster, be sure to reboot all nodes.
   configuration check succeeds
   ```

4. Commit the configuration.

   ```
   [edit]
   user@host# commit
   warning: You have enabled/disabled inet6 flow.
   You must reboot the system for your change to take effect.
   If you have deployed a cluster, be sure to reboot all nodes.
   commit complete
   ```

5. At an appropriate time, reboot the device.

Copyright © 2016, Juniper Networks, Inc.

NOTE: SRX Series devices only process IPv6 Routing Header 0 (RH0) to-self packets, the segleft field of which is zero. Other packets will be dropped.

Table 172 summarizes device status upon forwarding option configuration change.

Table 172: Device Status Upon Configuration Change

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

To process IPv6 traffic, you also need to configure IPv6 addresses for the transit interfaces that receive and forward the traffic. For information on the inet6 protocol family and procedures for configuring IPv6 addresses for interfaces, see the *Interfaces Feature Guide for Security Devices*.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Using Filters to Display IPv6 Session and Flow Information for SRX Series Services Gateways on page 1757

Using Filters to Display IPv6 Session and Flow Information for SRX Series Services Gateways

Purpose You can display flow and session information about one or more sessions with the **show security flow session** command. IPv6 sessions are included in aggregated statistics.

You can use the following filters with the **show security flow session** command: application, destination-port, destination-prefix, family, idp, interface, nat, protocol, resource-manager, session-identifier, source-port, source-prefix, and tunnel.

NOTE: Except for the session-identifier filter, the output of all the other filters can be viewed in brief, summary, and extensive mode. Brief mode is the default mode. The output of the session-identifier filter can be viewed only in the brief mode.

Copyright © 2016, Juniper Networks, Inc.

You can use the same filter options with the **clear security flow session** command to terminate sessions.

Action The following examples show how to use IPv6-related filters to display summaries and details for IPv6 sessions.

Filtered summary report based on family

```
root> show security flow session summary family ?
Possible completions:
  inet                 Show IPv4 sessions
  inet6                Show IPv6/IPv6-NATPT sessions
```

```
root> show security flow session summary family inet6
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 1
Sessions in other states: 0
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

Filtered detailed report based on family

```
root> show security flow session family ?
Possible completions:
  inet                 Show IPv4 sessions
  inet6                Show IPv6/IPv6-NATPT sessions
```

```
root> show security flow session family inet6
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000026, Policy name: default-policy-00/2, Timeout: 1794, Valid
  In: 6001::10/64712 --> 5001::2/21;tcp, If: ge-7/1/0.0, Pkts: 8, Bytes: 562, CP
 Session ID: 430000025
  Out: 5001::2/21 --> 6001::10/64712;tcp, If: ge-7/1/1.0, Pkts: 12, Bytes: 1014,
 CP Session ID: 430000025
Total sessions: 1
```

Filtered brief report based on family

```
root> show security flow session family inet brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000031, Policy name: default-policy-00/2, Timeout: 48, Valid
  In: 203.0.113.8/3 --> 198.51.100.11/43053;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes:
 84, CP Session ID: 410000039
  Out: 198.51.100.11/43053 --> 203.0.113.8/3;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes:
 0, CP Session ID: 410000039
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000034, Policy name: default-policy-00/2, Timeout: 48, Valid
  In: 203.0.113.8/4 --> 198.51.100.11/43053;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes:
 84, CP Session ID: 420000041
  Out: 198.51.100.11/43053  --> 203.0.113.8/4;icmp, If: ge-7/1/1.0, Pkts: 0,
Bytes: 0, CP Session ID: 420000041
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000042, Policy name: default-policy-00/2, Timeout: 44, Valid
  In: 203.0.113.8/2 --> 198.51.100.11/43053;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes:
 84, CP Session ID: 430000041
  Out: 198.51.100.11/43053 --> 203.0.113.8/2;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes:
 0, CP Session ID: 430000041
Total sessions: 1


2001:dbf8::6:2/32
```

Filtered detailed report based on an IPv6 source-prefix

```
root> show security flow session source-prefix 6001::10
Flow Sessions on FPC10 PIC1:
```

```
Session ID: 410000066, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 2001:dbf8::6:2/3 > 2001:dbf8:5::2/7214;icmp6, If: ge-7/1/0.0, Pkts: 1,
Bytes: 104, CP Session ID: 410000076
  Out: 2001:dbf8:5::2/7214 --> 2001:dbf8:5::2/323;icmp6, If: .local..0, Pkts: 1,
 Bytes: 104, CP Session ID: 410000076

Session ID: 410000068, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 2001:dbf8::6:2/4 --> 2001:dbf8:5::2/7214;icmp6, If: ge-7/1/0.0, Pkts: 1,
Bytes: 104, CP Session ID: 410000077
  Out: 2001:dbf8:5::2/7214 --> 2001:dbf8::6:2/4;icmp6, If: .local..0, Pkts: 1,
Bytes: 104, CP Session ID: 410000077
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000067, Policy name: default-policy-00/2, Timeout: 28, Valid
  In: 2001:dbf8::6:2/4 --> 2001:dbf8:5::3/6702;icmp6, If: ge-7/1/0.0, Pkts: 1,
Bytes: 104, CP Session ID: 420000080
  Out: 2001:dbf8:5::3/6702 --> 2001:dbf8::6:2/4 ;icmp6, If: ge-7/1/1.0, Pkts: 0,
 Bytes: 0, CP Session ID: 420000080
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000077, Policy name: default-policy-00/2, Timeout: 28, Valid
  In: 2001:dbf8::6:2/3 --> 2001:dbf8:5::3/6702;icmp6, If: ge-7/1/0.0, Pkts: 1,
Bytes: 104, CP Session ID: 430000075
  Out: 2001:dbf8:5::3/6702 --> 2001:dbf8::6:2/3;icmp6, If: ge-7/1/1.0, Pkts: 0,
Bytes: 0, CP Session ID: 430000075

Session ID: 430000078, Policy name: default-policy-00/2, Timeout: 30, Valid
  In: 2001:dbf8::6:2/5 --> 2001:dbf8:5::3/6702, If: ge-7/1/0.0, Pkts: 1, Bytes:
104, CP Session ID: 430000076
  Out: 2001:dbf8:5::3/6702 --> 2001:dbf8::6:2/5;icmp6, If: ge-7/1/1.0, Pkts: 0,
Bytes: 0, CP Session ID: 430000076

Session ID: 430000079, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 2001:dbf8::6:2/5 --> 2001:dbf8:5::1/7214;icmp6, If: ge-7/1/0.0, Pkts: 1,
Bytes: 104, CP Session ID: 430000077
  Out: 2001:dbf8:5::1/7214 --> 2001:dbf8::6:2/5;icmp6, If: .local..0, Pkts: 1,
Bytes: 104, CP Session ID: 430000077
Total sessions: 3
```

Multiple-filtered detailed report based on family, protocol and source-prefix

```
root> show security flow session family inet protocol icmp source-prefix 2001:db8::
Flow Sessions on FPC10 PIC1:

Session ID: 410000074, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 2001:dbf8::6:2/1 --> 2001:dbf8:8::2/26935;icmp, If: ge-7/1/0.0, Pkts: 1,
Bytes: 84, CP Session ID: 410000195
  Out: 2001:dbf8:8::2 --> 2001:dbf8::6:2/1;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes:
 84, CP Session ID: 410000195
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000075, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 2001:dbf8::6:2/3 --> 2001:dbf8::6:2/26935;icmp, If: ge-7/1/0.0, Pkts: 1,
Bytes: 84, CP Session ID: 420000159
  Out: 2001:dbf8::6:2/26935 --> 2001:dbf8::6:2/3;icmp, If: ge-7/1/1.0, Pkts: 1,
```


Copyright © 2016, Juniper Networks, Inc.

```
Bytes: 84, CP Session ID: 420000159
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000085, Policy name: default-policy-00/2, Timeout: 2, Valid
  In:  2001:dbf8::6:2/4 --> 2001:dbf8::6:2/26935;icmp, If: ge-7/1/0.0, Pkts: 1,
Bytes: 84, CP Session ID: 430000083
  Out: 2001:dbf8::6:2/26935 -->  2001:dbf8::6:2/4;icmp, If: ge-7/1/1.0, Pkts: 1,
 Bytes: 84, CP Session ID: 430000083
Total sessions: 1
```

Clearing all sessions, including IPv6 sessions

```
root> clear security flow session all
This command may terminate the current session too.
Continue? [yes,no] (no) yes

0 active sessions cleared
1 active sessions cleared
1 active sessions cleared
1 active sessions cleared
```

Clearing only IPv6 sessions

```
root> clear security flow session family ?
Possible completions:
  inet                 Clear IPv4 sessions
  inet6                Clear IPv6/IPv6-NATPT sessions

root> clear security flow session family inet6
0 active sessions cleared
1 active sessions cleared
1 active sessions cleared
1 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Enabling Flow-Based Processing for IPv6 Traffic on page 1756
- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787
- Clearing Sessions for SRX Series Services Gateways on page 1699

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 88

Managing IPv6 Packets

- The IPv6 Packet Header and SRX Series Overview on page 1763
- About the IPv6 Basic Packet Header on page 1764
- Understanding IPv6 Packet Header Extensions on page 1766
- About IPv6 Packet Header Verification Performed by the Flow Module for SRX Series Devices on page 1767
- Understanding Path MTU Messages for IPv6 Packets on page 1767
- Understanding How SRX Series Devices Handle Packet Fragmentation for IPv6 Flows on page 1769
- Understanding How SRX Series Devices Handle ICMPv6 Packets on page 1769

The IPv6 Packet Header and SRX Series Overview

This topic identifies the IP version 6 (IPv6) packet header and its extensions and options.

Every IPv6 packet at a minimum has a basic packet header, 40 bytes (320 bits) long. They optionally may have extension headers.

For IPv6 packets, flow processing parses the extension headers and transport layer headers in the following way:

- If the software encounters a TCP, a UDP, an ESP, an AH, or an ICMPv6 header, it parses the header and assumes that the packet payload corresponds to the specified protocol type.
- If the software encounters a hop-by-hop header, a routing and destination header, or a fragment header, it continues to parse the next extension header.
- If it encounters the no-next-header extension header, the software detects that the packet is that of an unknown protocol (protocol equals 0).
- For other extension headers, the software parses the header and identifies the packet as belonging to the protocol indicated by the extension header.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- About the IPv6 Basic Packet Header on page 1764

Copyright © 2016, Juniper Networks, Inc.

- Understanding IPv6 Packet Header Extensions on page 1766

About the IPv6 Basic Packet Header

This topic identifies the IP version 6 (IPv6) basic packet header fields with their bit lengths and uses.

Header Name	Bit Length	Purpose
Version	4	Specifies that IP version 6 is used. The IPv6 version field contains a value of 6 indicating that IPv6 is used, as opposed to 4 for IP version 4.
Traffic Class	8	Allows source nodes or routers to identify different classes (or priorities for quality of service) for IPv6 packets. (This field replaces the IPv4 Type of Service field.)
Flow Label	20	Identifies the flow to which the packet belongs. Packets in a flow share a common purpose, or belong to a common category, as interpreted by external devices such as routers or destination hosts. NOTE: For IPv6 flow-based packets, Junos OS for SRX Series devices does not use the flow label field.
Payload Length	16	Specifies the length of the IPv6 packet payload, or contents, expressed in octets.

Copyright © 2016, Juniper Networks, Inc.

Header Name	Bit Length	Purpose
Next Header	8	Identifies the type of Internet Protocol for the header that immediately follows the IPv6 header. The Next Header field replaces the IPv4 Protocol field. It is an optional field. This protocol can be one of two types: • An IPv6 extension header. For example, if the device performs IP security on exchanged packets, the Next Header value is probably 50 (ESP extension header) or 51 (AH extension header). Extension headers are optional. • An upper-layer Protocol Data Unit (PDU). For example, the Next Header value could be 6 (for TCP), 17 (for UDP), or 58 (for ICMPv6). The flow module processes these headers sequentially within the context of a packet flow. If it encounters one of the following extension headers, the software parses it and regards the packet as a corresponding protocol packet. • Internet Control Message Protocol version 6 (ICMPv6) • Transport Control Protocol (TCP) NOTE: The device checks the TCP header length as part of its sanity checks. • UDP NOTE: The device checks the UDP length as part of its sanity checks. • Enhanced Security Protocol (ESP) or Authentication Header (AH)
Hop Limit	8	Specifies the maximum number of hops the packet can make after transmission from the host device. When the Hop Limit value is zero, the device drops the packet and generates an error message. (This field is similar the to Time to Live IPv4 field.)
Source IP Address	128	Identifies the host device, or interface on a node, that generated the IPv6 packet.
Destination IP Address	128	Identifies the host device, or interface on a node, to which the IPv6 packet is to be sent. NOTE: The destination address can appear twice, the first instance after the hop limit following the source IP address and the second instance after the final extension header.

Related Documentation

- Understanding IPv6 Packet Header Extensions on page 1766
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Copyright © 2016, Juniper Networks, Inc.

Understanding IPv6 Packet Header Extensions

This topic defines IP version 6 (IPv6) packet header extensions.

IPv6 extension headers contain supplementary information used by network devices (such as routers, switches, and endpoint hosts) to decide how to direct or process an IPv6 packet. The length of each extension header is an integer multiple of 8 octets. This allows subsequent extension headers to use 8-octet structures.

Any header followed by an extension header contains a Next Header value that identifies the extension header type. Extension headers always follow the basic IPv6 header in order as shown in Table 173:

NOTE: **The destination IP address can appear twice, once after the hop-by-hop header and again after the last extension header.**

Table 173: IPv6 Extension Headers

Header Name	Purpose
Hop-by-Hop Options	Specifies delivery parameters at each hop on the path to the destination host. NOTE: A hop-by-hop option can appear only following the IPv6 basic header. If it is used, it should be the first extension header. It cannot appear after another extension header.
Destination Options	Specifies packet delivery parameters for either intermediate destination devices or the final destination host. When a packet uses this header, the Next Header value of the previous header must be 60.
Routing	Defines strict source routing and loose source routing for the packet. (With strict source routing, each intermediate destination device must be a single hop away. With loose source routing, intermediate destination devices can be one or more hops away.) When a packet uses this header, the Next Header value of the previous header must be 43.
Fragment	Specifies how to perform IPv6 fragmentation and reassembly services. When a packet uses this header, the Next Header value of the previous header must be 44. A source node uses the fragment extension header to tell the destination node the size of the packet that was fragmented so that the destination node can reassemble the packet.
Authentication	Provides authentication, data integrity, and anti-replay protection. When a packet uses this header, the Next Header value of the previous header must be 51.
Encapsulating Security Payload	Provides data confidentiality, data authentication, and anti-replay protection for Encapsulated Security Payload (ESP) packets. When a packet uses this header, the Next Header value of the previous header must be 50.
Destination IP Address	Identifies the host device, or interface on a node, to which the IPv6 packet is to be sent. NOTE: The destination address may appear twice, the first instance after the hop limit following the source IP address and the second instance after the final extension header.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- About the IPv6 Basic Packet Header on page 1764
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

About IPv6 Packet Header Verification Performed by the Flow Module for SRX Series Devices

This topic gives an overview of some of the IP version 6 (IPv6) packet header verification that the flow module for SRX Series devices performs.

To ensure the integrity of an IPv6 packet, the flow module performs the following sanity checks.

For all IPv6 packets, it checks the following parts of the header:

- TCP length
- UDP length
- Hop-by-hop extension to ensure that it follows the basic IPv6 header and does not come after another extension header
- That the IP data length error (IP length—total extension header length is not less than zero (<0)

  In addition to these verifications, the software performs other standard checks such as verifying that the correct IP version is specified and that the length of the IP address is correct.

Related Documentation
- About the IPv6 Basic Packet Header on page 1764
- Understanding IPv6 Packet Header Extensions on page 1766
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Understanding Path MTU Messages for IPv6 Packets

This topics describes path maximum transmission unit (MTU) and explains how the flow module for SRX Series devices processes and uses path MTU messages.

Every link has an MTU size that specifies the size of the largest packet the link can transmit. A larger MTU size means that fewer packets are required to transmit a certain amount of data. To achieve the best data transmission performance, IPv6 data packets sent from one node (the source) to another node (the destination) should be the largest possible size that can traverse the path between the nodes. (Larger and fewer packets constrain the cost of packet header processing and routing processes that can affect transmission performance.)

However, for a packet to successfully traverse the path from the source node to the destination node, the MTU size of the source node interface must be no larger than that

Copyright © 2016, Juniper Networks, Inc.

of the smallest MTU size of all nodes on the path between the source and destination. This value is referred to as the path maximum transmission unit (path MTU). If a packet is larger than a link's MTU size, it is likely that the link will drop it. For IPv6, an intermediate node cannot fragment a packet.

IPv6 defines a standard mechanism called path MTU discovery that a source node can use to learn the path MTU of a path that a packet is likely to traverse. If any of the packets sent on that path are too large to be forwarded by a node along the path, that node discards the packet and returns an ICMPv6 Packet Too Big message. The source node can then adjust the MTU size to be smaller than that of the node that dropped it and sent the ICMPv6 message, and then retransmit the packet. A source node might receive Packet Too Big messages repeatedly until its packet traverses all nodes along the path successfully.

NOTE: On all SRX Series devices, the Routing Engine cannot detect the path MTU of an IPv6 multicast address (with a large size packet).

After the path MTU size is determined and the appropriate MTU size is set, an outgoing packet might be routed along a different path with a node whose link MTU size is smaller than the path MTU size determined previously. In this case, the flow module engages the path MTU discovery process again.

When the flow module receives an ICMP Packet Too Big message with a destination address that belongs to it, it:

- Checks to determine if the embedded 5-tuple data of the packet is for a tunnel interface. (That is, it checks to determine if the embedded 5-tuple data matches a tunnel session.) If there is a match, the flow module updates the tunnel interface's MTU size. Then it performs post-fragment processing for the encrypted packets that follow the first packet. Afterward, the flow module delivers the packet to the ICMPv6 stack on the Routing Engine (RE) for it to continue processing it.
- If the packet is a transit one, the flow module searches for a session that matches the packet's embedded 5-tuple data. It is finds a matching session, it delivers the packet to it. If there is no matching session, it drops the packet.

When the flow module receives a packet, before it transmits it to the egress interface, it checks to determine if the MTU size of the egress interface is greater than the packet length.

- If the MTU size is greater than the packet length, it continues to process the packet.
- If the MTU size is less than the packet length, it drops the packet and sends an ICMPv6 Packet Too Big message to the source node.

NOTE: When chassis cluster is configured and the path MTU updates the MTU of the tunnel interface, the flow module does not synchronize the new MTU to peer nodes. The MTU size might be updated again by a larger packet on a peer node, which has no impact on packet transmission.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding How SRX Series Devices Handle ICMPv6 Packets on page 1769
- About the IPv6 Basic Packet Header on page 1764
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Understanding How SRX Series Devices Handle Packet Fragmentation for IPv6 Flows

This topic explains packet fragmentation for IP version 6 (IPv6).

For IPv4 Internet Control Message Protocol (IPv4 ICMP), if a node within the path between a source node and a destination node receives a packet that is larger than its MTU size, it can fragment the packet and transmit the resulting smaller packets. For IPv6, only a source node (the node that sent the packet) can fragment a packet, and this is done to accommodate a path MTU size-adjustment requirement. Nodes along the path of a packet cannot fragment the packet to transmit it.

Related Documentation

- Understanding How SRX Series Devices Handle ICMPv6 Packets on page 1769
- Understanding Path MTU Messages for IPv6 Packets on page 1767
- Understanding IPv6 Packet Header Extensions on page 1766
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Understanding How SRX Series Devices Handle ICMPv6 Packets

This topic explains Internet Control Message Protocol (ICMP), ICMP messages, and how Junos OS for SRX Series Services Gateways uses them.

ICMP provides a framework for reporting packet processing errors, for diagnostic purposes, and for implementation-specific functions. ICMP error messages make it possible for one node to inform another node that something has gone wrong during the course of data transfer. When IP version 6 (IPv6) was defined, the differences between IP version 4 (IPv4) and it were significant enough to require a new version of ICMP.

Every ICMPv6 message is preceded by an IPv6 header and zero or more IPv6 extension headers. The ICMPv6 header is identified by a Next Header value of 58 in the immediately preceding header. This is different from the value used to identify ICMP for IPv4. All ICMPv6 error messages have 32 bits of type-specific data to help the packet recipient locate the embedded invoking packet.

Most ICMPv6 packets have the same characteristics and behavior as normal IPv6 packets, and the Junos OS flow module processes them through first path and fast-path processing in the same way that it does normal IPv6 packets. Table 174 shows the ICMPv6 embedded packet types that the flow module handles differently from normal ICMPv6 packets.

For these packets, the flow module uses a tuple that it creates from the embedded ICMPv6 packet to search for a matching session. It continues to process the packet

Copyright © 2016, Juniper Networks, Inc.

without modifying the maximum transmission unit (MTU) until it finds a matching session, unless it receives an ICMPv6 Packet Too Big message for the interface. In this case, it modifies the MTU size for that interface. If the flow module does not find a matching session or if it cannot obtain a valid IPv6 header from the embedded payload, it drops the packet.

NOTE: A Packet Too Big message is the only kind of ICMPv6 packet that will cause the flow module to modify an interface.

Table 174: ICMPv6 Packets That Junos OS Handles Differently from Other ICMPv6 Packets

Message	Meaning
01-Destination Unreachable	When a packet cannot be delivered because of a problem with the way it is being sent, it is useful to have a feedback mechanism that can tell the source about the problem, including the reason why delivery of the packet failed. For IPv6, the Destination Unreachable message serves this purpose. Each message includes a code that indicates the nature of the problem that caused the packet delivery to fail. It also includes all or part of the packet that could not be delivered, to help the source device resolve the problem. When the flow module encounters a Destination Unreachable ICMP packet whose embedded packet header data matches the 5-tuple data for a session, the software terminates the session.
02-Packet Too Big	When the flow module receives an ICMPv6 Packet Too Big message intended for it, the flow module sends the packet to the ICMP protocol stack on the Routing Engine to engage the path maximum transmission unit (path MTU) discovery process. If the Packet Too Big message does not pertain to the device but rather is a transit packet, the device attempts to match the embedded 5-tuple data with a session. • If a matching session exists, the device delivers it to the source node. • If a matching session does not exist, the device drops the packet NOTE: A Packet Too Big message is the only kind of ICMPv6 packet that will cause the flow module to modify an interface.
03-Time Exceeded	When the flow module receives a packet that cannot be delivered because it has exceeded the hop count specified in the basic header hop-by-hop field, it sends this message to inform the packet's source node that the packet was discarded for this reason.
04-Parameter Problem	When the device finds a problem with a field in the IPv6 header or extension headers that makes it impossible for it to process the packet, the software discards it and sends this ICMPv6 message to the packet's source node, indicating the type and location of the problem.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Path MTU Messages for IPv6 Packets on page 1767
- Understanding How SRX Series Devices Handle Packet Fragmentation for IPv6 Flows on page 1769
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 89

Configuring IPv6 Dual-Stack

- Understanding IPv6 Dual-Stack Lite on page 1773
- Example: Configuring IPv6 Dual-Stack Lite on page 1776

Understanding IPv6 Dual-Stack Lite

IPv6 dual-stack lite (DS-Lite) is a technology that enables Internet service providers to move to an IPv6 network while simultaneously handling IPv4 address depletion.

IPv4 addresses are becoming depleted; therefore, broadband service providers (DSL, cable, and mobile) need new addresses to support new users. Providing IPv6 addresses alone is often not workable because most of the systems that make up the public Internet are still enabled and support only IPv4, and many users' systems do not yet fully support IPv6.

DS-Lite allows service providers to migrate to an IPv6 access network without changing end-user software. The device that accesses the Internet remains the same, thus allowing IPv4 users to continue accessing IPv4 internet content with minimum disruption to their home networks, while enabling IPv6 users to access IPv6 content.

Figure 91 illustrates the DS-Lite architecture which uses IPv6-only links between the provider and the user while maintaining the IPv4 (or dual-stack) hosts in the user network.

Copyright © 2016, Juniper Networks, Inc.

Figure 91: DS-Lite NAT (IPv4-in-IPv6)

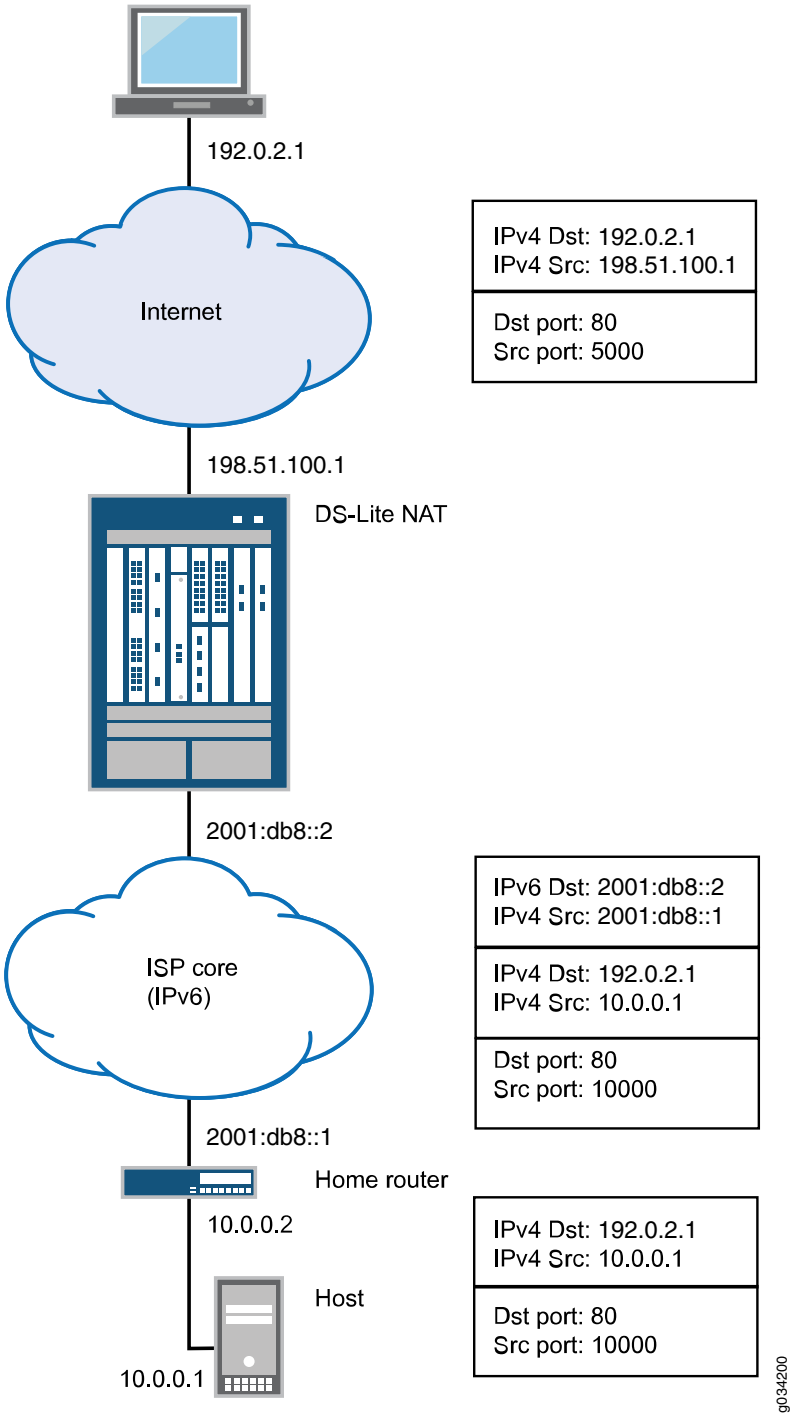

The DS-Lite deployment model consists of the following components:

Copyright © 2016, Juniper Networks, Inc.

- Softwire initiator for the DS-Lite home router--Encapsulates the IPv4 packet and transmits it across an IPv6 tunnel.
- Softwire concentrator for DS-Lite carrier-grade Network Address Translation (NAT)—Decapsulates the IPv4-in-IPv6 packet and also performs IPv4-IPv4 NAT translations.

When a user's device sends an IPv4 packet to an external destination, DS-Lite encapsulates the IPv4 packet in an IPv6 packet for transport into the provider network. These IPv4-in-IPv6 tunnels are called *softwires*. Tunneling IPv4 over IPv6 is simpler than translation and eliminates performance and redundancy concerns.

The softwires terminate in a softwire concentrator at some point in the service provider network, which decapsulates the IPv4 packets and sends them through a carrier-grade Network Address Translation (NAT) device. There, the packets undergo source NAT processing to hide the original source address.

IPv6 packets originated by hosts in the subscriber's home network are transported natively over the access network.

The DS-Lite carrier-grade NAT translates IPv4-to-IPv4 addresses to multiple subscribers through a single global IPv4 address. Overlapping address spaces used by subscribers are disambiguated through the identification of tunnel endpoints. One concentrator can be the endpoint of multiple softwires.

The IPv4 packets originated by the end hosts have private (and possibly overlapping) IP addresses. Therefore, NAT must be applied to these packets. If end hosts have overlapping addresses, Network Address Port Translation (NAPT) is needed.

Using NAPT, the system adds the source address of the encapsulating IPv6 packet in the subscriber network to the inside IPv4 source address and port. Because each user's IPv6 address is unique, the combination of the IPv6 source address with the IPv4 source address and port creates an unambiguous mapping.

The system takes the following actions when it receives a responding IPv4 packet from outside the subscriber network:

- Encapsulates the IPv4 packet in an IPv6 packet using the mapped IPv6 address as the IPv6 destination address.
- Forwards the packet to the user.

Table 175 lists the maximum number of softwire initiators and softwire concentrators per device.

Table 175: Softwire Initiator and Softwire Concentrator Capacity

Description	SRX1500	SRX5400	SRX5600	SRX5800
Maximum softwire initiators connected per device	300	100,000	100,000	100,000
Maximum softwire concentrator numbers per device	32	32	32	32

Copyright © 2016, Juniper Networks, Inc.

NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:

- The term softwire initiator has been replaced by B4.
- The term softwire concentrator has been replaced by AFTR.

Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the CLI statements used to configure DS-Lite.

For more information, see the following documents:

- draft-ietf-softwire-dual-stack-lite-06, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*, August 2010.
- RFC 2473, *Generic Packet Tunneling in IPv6 Specification*, December 1998.
- RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, August 1999.
- RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, BCP 127*, January 2007.
- RFC 4925, *Softwire Problem Statement*, July 2007.
- RFC 5382, *NAT Behavioral Requirements for TCP, BCP 142*, October 2008.
- RFC 5508, *NAT Behavioral Requirements for ICMP, BCP 148*, April 2009.
- http://www.potaroo.net/tools/ipv4/index.html
- http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml

Related Documentation

- Example: Configuring IPv6 Dual-Stack Lite on page 1776
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Understanding How SRX Series Devices Handle ICMPv6 Packets on page 1769
- About the IPv6 Basic Packet Header on page 1764

Example: Configuring IPv6 Dual-Stack Lite

When an ISP begins to allocate IPv6 addresses and IPv6-capable equipment to new subscriber homes, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 CE WAN equipment to reach the IPv4 network. DS-Lite enables IPv4 customers to continue to access the Internet using their current hardware by using a softwire initiator at the customer edge to encapsulate IPv4 packets into IPv6 packets with minimum disruption to their home network, while enabling IPv6 customers to access IPv6 content. The softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4-IPv4 NAT translations.

Copyright © 2016, Juniper Networks, Inc.

This example shows you how to configure a softwire concentrator for IPv4-in-IPv6 addresses.

- Requirements on page 1777
- Overview on page 1777
- Configuration on page 1777
- Verification on page 1778

Requirements

Before you begin:

- Review the overview section on DS-Lite. See "Understanding IPv6 Dual-Stack Lite" on page 1773.
- Review how ICMPv6 packets are handled by the SRX Series devices. See "Understanding How SRX Series Devices Handle ICMPv6 Packets" on page 1769.

Overview

This configuration example shows how to configure a softwire concentrator, the softwire name, the concentrator address, and the softwire type.

NOTE: **The softwire concentrator IPv6 address can match an IPv6 address configured on a physical interface or an IPv6 address configured on a loopback interface.**

Configuration

CLI Quick Configuration

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

```
set security softwires softwire-name my_sc1 softwire-concentrator 2001:db8::1
  softwire-type IPv4-in-IPv6
```

Step-by-Step Procedure

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

To configure a DS-Lite softwire concentrator to convert IPv4 packets into IPv6 packets:

1. Assign a name for the softwire concentrator.

   ```
   [edit security]
   user@host# edit softwires softwire-name my_sc1
   ```

2. Specify the address of the softwire concentrator.

   ```
   [edit security softwires softwire-name my_sc1]
   user@host# set softwire-concentrator 2001:db8::1
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Specify the softwire type for IPv4 to IPv6.

```
[edit security softwires softwire-name my_sc1 softwire-concentrator 2001:db8::1
user@host# set softwire-type IPv4-in-IPv6
```

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

```
[edit security softwires softwire-name my_sc1]
user@host# show
softwire-concentrator 2001:db8::1;
softwire-type ipv4-in-ipv6;
```

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

Verification

From operational mode, enter the **show security softwires** command. If a softwire is not connected, the operational output looks like the following sample:

```
user@host# show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  my-sc1                2001:db8::1         Active      0
```

If a softwire is connected, the operational output looks like the following sample:

```
user@host# show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  my-sc1                2001:db8::1         Connected      1
```

Related Documentation

- Understanding IPv6 Dual-Stack Lite on page 1773
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- About the IPv6 Basic Packet Header on page 1764

Copyright © 2016, Juniper Networks, Inc.

PART 26

Monitoring Flow-Based Sessions

- Monitoring Security Flow Sessions on page 1781
- Monitoring X2 Traffic By Configuring Mirror Filters on page 1793

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 90

Monitoring Security Flow Sessions

- Monitoring Security Flow Sessions Overview on page 1781
- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Monitoring Security Flow Sessions Overview

Junos OS allows you to configure and start the monitoring of flow sessions using operational mode commands. Thus, you can debug without having to commit or modify your running configuration. This approach can be especially useful when you do not want to change the state of your device by committing the configuration to turn on trace options.

To configure flow session monitoring, you must define flow filters, specify the output file, and start monitoring. Flow session monitoring does not start unless a filter (at least one) and an output file are specified. Also, defining the filters themselves does not trigger monitoring. You have to explicitly use the **monitor security flow start** and **monitor security flow stop** commands to enable and disable monitoring, respectively.

Copyright © 2016, Juniper Networks, Inc.

- Define flow filters—Define the flow sessions that you want to monitor using combinations of match criteria, such as source address, destination address, source port, destination port, IP protocol number, name of the incoming or outgoing interface, and the logical system name. You can delete filters using the **clear monitor security flow filter** command.

  NOTE: Unlike filters defined in the configuration mode, filters defined using operational mode commands are cleared when you reboot your system.

- Specify the output file—Create an output file in which the security flow monitoring information is to be saved. This file is saved in the **/var/log/** directory. You can view the contents of this file by using the **show log** *filename* command. Use the **monitor security flow file** command to specify output file characteristics, such as its maximum size, maximum number, and type.
- Start monitoring—Use the **monitor security flow start** command to start monitoring. Once monitoring starts, any traffic that matches the filters is saved in the specified output file in the **/var/log/** directory. The basic-datapath flag is the default flag and turns on as monitoring starts.

  Use the **monitor security flow stop** command to stop monitoring. Once monitoring stops, the basic-datapath flag is cleared.
- Display monitoring flow information—Use the **show monitoring security flow** command to display details about the monitoring operation.

  NOTE: You can configure flow session monitoring and debugging by using the monitoring operational mode commands and flow traceoptions configuration statements. These two operations cannot run in parallel. When you turn on security flow monitoring, the flow traceoption session is blocked and when the flow traceoption session is running, monitoring of the flow session is blocked.

Related Documentation

- monitor security flow start on page 1926
- monitor security flow file on page 1922
- monitor security flow filter on page 1924
- show monitor security flow on page 2029
- monitor security flow stop on page 1927
- clear monitor security flow filter on page 1899

Understanding How to Obtain Session Information for SRX Series Services Gateways

You can obtain information about the sessions and packet flows active on your device, including detailed information about specific sessions. (The SRX Series device also

Copyright © 2016, Juniper Networks, Inc.

displays information about failed sessions.) You can display this information to observe activity and for debugging purposes. For example, you can use the show security flow session command:

- To display a list of incoming and outgoing IP flows, including services
- To show the security attributes associated with a flow, for example, the policies that apply to traffic belonging to that flow
- To display the session timeout value, when the session became active, for how long it has been active, and if there is active traffic on the session

NOTE: If an interface NAT is configured and sessions are set up with the NAT using that interface IP address, whenever the interface IP address changes, the sessions set up with NAT get refreshed and new sessions will be setup with new IP address. This you can verify using **show security flow session** CLI command.

Session information can also be logged if a related policy configuration includes the logging option. See"Information Provided in Session Log Entries for SRX Series Services Gateways" on page 1787 for details about session information provided in system logs.

For the flow session log on all SRX Series devices, policy configuration has been enhanced. Information on the packet incoming interface parameter in the session log for session-init and session-close and when a session is denied by a policy or by the application firewall is provided to meet Common Criteria (CC) Medium Robustness Protection Profiles (MRPP) compliance:

Policy configuration—To configure the policy for the session for which you want to log matches as log **session-init** or **session-close** and to record sessions in syslog:

- **set security policies from-zone untrustZone to-zone trust zone policy policy13 match source-address extHost1**
- **set security policies from-zone untrustZone to-zone trust zone policy policy13 match source-address extHost1**
- **set security policies from-zone untrustZone to-zone trustZone policy policy13 match application junos-ping**
- **set security policies from-zone untrustZone to-zone trustZone policy policy13 then permit**
- **set security policies from-zone untrustZone to-zone trustZone policy policy13 then log session-init**
- **set security policies from-zone untrustZone to-zone trustZone policy policy13 then log session-close**

Example : Flow match policy13 will record the following information in the log:

<14>1 2010-09-30T14:55:04.323+08:00 mrpp-srx550-dut01 RT_FLOW - RT_FLOW_SESSION_CREATE [junos@2626.192.0.2.1.40 source-address="192.0.2.1"

Copyright © 2016, Juniper Networks, Inc.

source-port="1" destination-address="198.51.100.12" destination-port="46384" service-name="icmp" nat-source-address="192.0.2.1" nat-source-port="1" nat-destination-address="198.51.100.12" nat-destination-port="46384" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="1" policy-name="policy1" source-zone-name="trustZone" destination-zone-name="untrustZone" session-id-32="41" packet-incoming-interface="ge-0/0/1.0"] session created 192.0.2.1/1-->198.51.100.12/46384 icmp 192.0.2.1/1-->198.51.100.12/46384 None None 1 policy1 trustZone untrustZone 41 ge-0/0/1.0

<14>1 2010-09-30T14:55:07.188+08:00 mrpp-srx550-dut01 RT_FLOW - RT_FLOW_SESSION_CLOSE [junos@2626.192.0.2.1.40 reason="response received" source-address="192.0.2.1" source-port="1" destination-address="198.51.100.12" destination-port="46384" service-name="icmp" nat-source-address="192.0.2.1" nat-source-port="1" nat-destination-address=”198.51.100.12" nat-destination-port="46384" src-nat-rule-name="None" dst-nat-rule-name="None" protocol-id="1" policy-name="policy1" source-zone-name="trustZone" destination-zone-name="untrustZone" session-id-32="41" packets-from-client="1" bytes-from-client="84" packets-from-server="1" bytes-from-server="84" elapsed-time="0" packet-incoming-interface="ge-0/0/1.0"] session closed response received: 192.0.2.1/1-->198.51.100.12/46384 icmp 192.0.2.1/1-->198.51.100.12/46384 None None 1 policy1 trustZone untrustZone 41 1(84) 1(84) 0 ge-0/0/1.0

Related Documentation

- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Clearing Sessions for SRX Series Services Gateways on page 1699
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- **allow-embedded-icmp on page 1841**

Displaying Global Session Parameters for All SRX Series Services Gateways

Purpose Obtain information about configured parameters that apply to all flows or sessions.

Action To view session information in the CLI, enter the following command:

```
user@host# show security flow
```

Meaning The **show security flow** configuration command displays the following information:

- **allow-dns-reply**—Identifies if unmatched incoming Domain Name System (DNS) reply packets are allowed.
- **route-change-timeout**—If enabled, displays the session timeout value to be used on a route change to a nonexistent route.
- **tcp-mss**—Shows the current configuration for the TCP maximum segment size value to be used for all TCP packets for network traffic.
- **tcp-session**—Displays all configured parameters that control session parameters.
- **syn-flood-protection-mode**—Displays the SYN Proxy mode.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Displaying a Summary of Sessions for SRX Series Services Gateways

Purpose Determine the kinds of sessions on your device, how many of each kind there are—for example, the number of unicast sessions and multicast sessions—the number of failed sessions, the number of sessions that are currently used and the maximum number of sessions that the device supports. This command also displays the details of the sessions that are currently used. For example, valid sessions, pending sessions, invalidated sessions and sessions in other states.

Action To view session summary information in the CLI, enter the following CLI command:

```
user@host> show security flow session summary
```

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Displaying Session and Flow Information About Sessions for SRX Series Services Gateways

Purpose Display information about all sessions on your device, including the session ID, the virtual system the session belongs to, the Network Address Translation (NAT) source pool (if source NAT is used), the configured timeout value for the session and its standard timeout,

Copyright © 2016, Juniper Networks, Inc.

and the session start time and how long the session has been active. The display also shows all standard flow information, including the direction of the flow, the source address and port, the destination address and port, the IP protocol, and the interface used for the session.

Action To view session flow information in the CLI, enter the following command:

```
user@host> show security flow session
```

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways

Purpose When you know the session identifier, you can display all session and flow information for a specific session rather than for all sessions.

Action To view information about a specific session in the CLI, enter the following command:

```
user@host> show security flow session session-identifier 40000381
```

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Copyright © 2016, Juniper Networks, Inc.

Using Filters to Display Session and Flow Information for SRX Series Services Gateways

Purpose You can display flow and session information about one or more sessions by specifying a filter as an argument to the **show security flow session** command. You can use the following filters: application, destination-port, destination-prefix, family, idp, interface, nat, protocol, resource-manager, session-identifier, source-port, source-prefix and tunnel. The device displays the information for each session followed by a line specifying the number of sessions reported on. Here is an example of the command using the source-prefix filter.

Action To view information about selected sessions using filters in the CLI, enter the following command:

```
user@host> show security flow session source-prefix 10/8
```

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Information Provided in Session Log Entries for SRX Series Services Gateways on page 1787

Information Provided in Session Log Entries for SRX Series Services Gateways

Session log entries are tied to policy configuration. Each main session event—create, close, and deny—will create a log entry if the controlling policy has enabled logging.

Different fields are logged for session create, session close, and session deny events as shown in Table 176, Table 177, and Table 178. The same field name under each type indicates that the same information is logged, but each table is a full list of all data recorded for that type of session log.

The following table defines the fields displayed in session log entries.

Table 176: Session Create Log Fields

Field	Description
source-address	Source IP address of the packet that created the session.
source-port	Source port of the packet that created the session.
destination-address	Destination IP address of the packet that created the session.

Copyright © 2016, Juniper Networks, Inc.

Table 176: Session Create Log Fields *(continued)*

Field	Description
destination-port	Destination port of the packet that created the session.
service-name	Application that the packet traversed (for example, "junos-telnet" for Telnet traffic during the session allowed by a policy that permits native Telnet).
nat-source-address	The translated NAT source address if NAT was applied; otherwise, the source address as above.
nat-source-port	The translated NAT source port if NAT was applied; otherwise, the source port as above.
nat-destination-address	The translated NAT destination address if NAT was applied; otherwise, the destination address as above.
nat-destination-port	The translated NAT destination port if NAT was applied; otherwise, the destination port as above.
src-nat-rule-name	The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.*
dst-nat-rule-name	The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.*
protocol-id	The protocol ID of the packet that created the session.
policy-name	The name of the policy that permitted the session creation.
session-id-32	The 32-bit session ID.

** Note that some sessions might have both destination and source NAT applied and the information logged.*

Table 177: Session Close Log Fields

Field	Description
reason	The reason the session was closed.
source-address	Source IP address of the packet that created the session.
source-port	Source port of the packet that created the session.
destination-address	Destination IP address of the packet that created the session.
destination-port	Destination port of the packet that created the session.
service-name	Application that the packet traversed (for example, "junos-telnet" for Telnet traffic during the session allowed by a policy that permits native Telnet).

Copyright © 2016, Juniper Networks, Inc.

Table 177: Session Close Log Fields *(continued)*

Field	Description
nat-source-address	The translated NAT source address if NAT was applied; otherwise, the source address as above.
nat-source-port	The translated NAT source port if NAT was applied; otherwise, the source port as above.
nat-destination-address	The translated NAT destination address if NAT was applied; otherwise, the destination address as above.
nat-destination-port	The translated NAT destination port if NAT was applied; otherwise, the destination port as above.
src-nat-rule-name	The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.*
dst-nat-rule-name	The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.*
protocol-id	The protocol ID of the packet that created the session.
policy-name	The name of the policy that permitted the session creation.
session-id-32	The 32-bit session ID.
packets-from-client	The number of packets sent by the client related to this session.
bytes-from-client	The number of data bytes sent by the client related to this session.
packets-from-server	The number of packets sent by the server related to this session.
bytes-from-server	The number of data bytes sent by the server related to this session.
elapsed-time	The total session elapsed time from permit to close, given in seconds.
unset	During the session creation, you can set the session close reason as **unset**. The session closes with the reason **unset** if the session installation on the control point is not successful. The reason for session installation varies, for example, nonavailability of memory for nonmanagement session installation.
TCP RST	RST received from either end.

Copyright © 2016, Juniper Networks, Inc.

Table 177: Session Close Log Fields *(continued)*

Field	Description
TCP FIN	FIN received from either end.
response received	Response received for a packet request (for example, ICMP req-reply).
ICMP error	ICMP error received.
aged out	Session aged out was reached.
ALG	ALG errors closed the session (for example, remote access server (RAS) maximum limit reached).
HA	HA message closed the session.
auth	Authentication failed.
IDP	IDP closed the session because of security module (SM) internal error.
synproxy failure	SYN proxy failure closed the session.
synproxy limit	Reason for failure in allocating minor session, need to free original session.
parent closed	Parent session closed.
CLI	Session cleared by a CLI statement.
CP NACK	CP NACK response received.
CP delete	CP ACK deletion closed the session.
policy delete	Corresponding policy marked for deletion.
fwd session	Session closed because of forwarding session deletion.
multicast route change	Session closed because multicast route changed.
first path reroute, session recreated	The first path is rerouted and session is re-created.
source NAT allocation failure	SPU received ACK message from the central point but failed to receive the DIP resource. Therefore this packet is dropped and the session is closed.
other	Session closed because of all other reasons (for example, the pim reg tun needed refreshing).
error create IKE pass-through template	IKE pass-through template creation errors.

Copyright © 2016, Juniper Networks, Inc.

Table 177: Session Close Log Fields *(continued)*

Field	Description
IKE pass-through child session ageout	Session is deleted because the IKE pass through template session has no child.
sess timeout on pending state	Pending session closed because time out timer reached the pending state.
unknown	Session closed because of unknown reasons.

** Note that some sessions might have both destination and source NAT applied and the information logged.*

Table 178: Session Deny Log Fields

Field	Description
source-address	Source IP address of the packet that attempted to create the session.
source-port	Source port of the packet that attempted to create the session.
destination-address	Destination IP address of the packet that attempted to create the session.
destination-port	Destination port of the packet that attempted to create the session.
service-name	Application that the packet attempted to traverse.
protocol-id	The protocol ID of the packet that attempted to create the session.
icmp-type	The ICMP type if the denied packet was ICMP configured; otherwise, this field will be 0.
policy-name	The name of the policy that denied the session creation.

Related Documentation

- Understanding How to Obtain Session Information for SRX Series Services Gateways on page 1782
- Displaying Global Session Parameters for All SRX Series Services Gateways on page 1784
- Displaying a Summary of Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About Sessions for SRX Series Services Gateways on page 1785
- Displaying Session and Flow Information About a Specific Session for SRX Series Services Gateways on page 1786
- Clearing Sessions for SRX Series Services Gateways on page 1699
- Using Filters to Display Session and Flow Information for SRX Series Services Gateways on page 1787

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 91

Monitoring X2 Traffic By Configuring Mirror Filters

- Understanding X2 Traffic Monitoring on page 1793
- Example: Configuring Mirror Filters for X2 Traffic Monitoring on page 1796

Understanding X2 Traffic Monitoring

In an LTE mobile network, SRX Series devices act as secure gateways and connect Evolved Node Bs (eNodeBs) for signal handover, monitoring, and radio coverage. SRX Series devices use IPsec tunnels to connect eNodeBs. The user plane and control plane traffic that flows from one eNodeB to the other eNodeB is called the X2 traffic.

This topic covers X2 traffic monitoring on SRX Series devices:

- X2 Traffic Monitoring Overview on page 1793
- Limitations of X2 Traffic Monitoring on page 1795
- X2 Traffic Terminology on page 1795

X2 Traffic Monitoring Overview

The X2 traffic passing through IPsec tunnels is encrypted; thus the mobile network operators need a way to monitor X2 traffic so that they can debug handover issues across eNodeBs. The Junos OS implementation allows monitoring of the X2 traffic by snooping into the cleartext X2 traffic as it flows through the SRX Series device coming out of one IPsec tunnel and going into the other IPsec tunnel—after traffic is decrypted and before it is encrypted again.

Figure 92 shows the flow of X2 traffic within the SRX Series device. As the traffic reaches the SRX Series device on one st0.x interface, it gets decrypted. Then it is encrypted and forwarded to the destination eNodeB through its dedicated st0.y interface. Snooping is performed on the decrypted X2 traffic on the SRX Series device.

Copyright © 2016, Juniper Networks, Inc.

Figure 92: SRX Series Device in an LTE Mobile Network

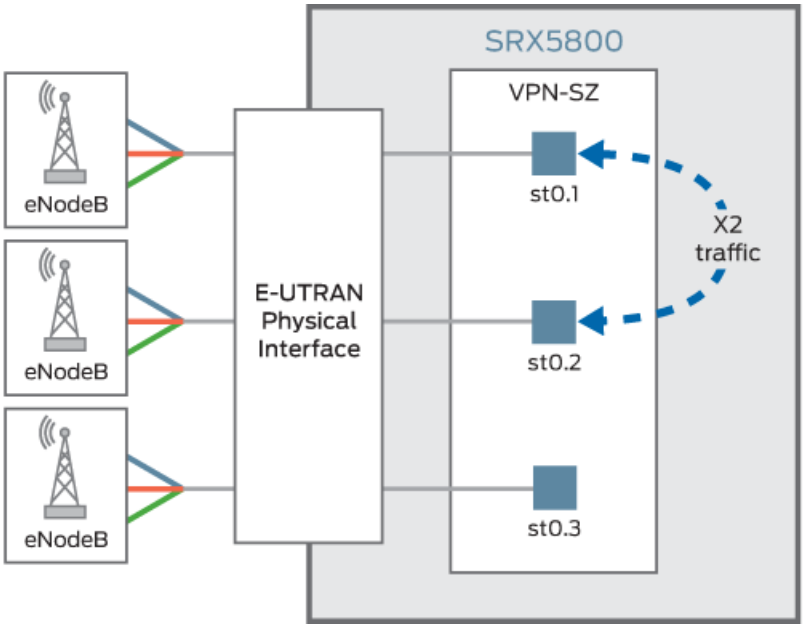

Figure 93 shows a mobile operators network with an SRX Series device providing IPsec tunnel connection between the two eNodeBs. The SRX Series device is connected to a packet analyzer (also called a *sniffing* device) that is used for collecting and checking the X2 interface traffic. The IPsec tunnel from each eNodeB terminates on a dedicated secure tunnel interface on the SRX Series device. Traffic coming out of the IPsec tunnel is decrypted while traffic in the opposite direction is encrypted.

Figure 93: Monitoring X2 Traffic

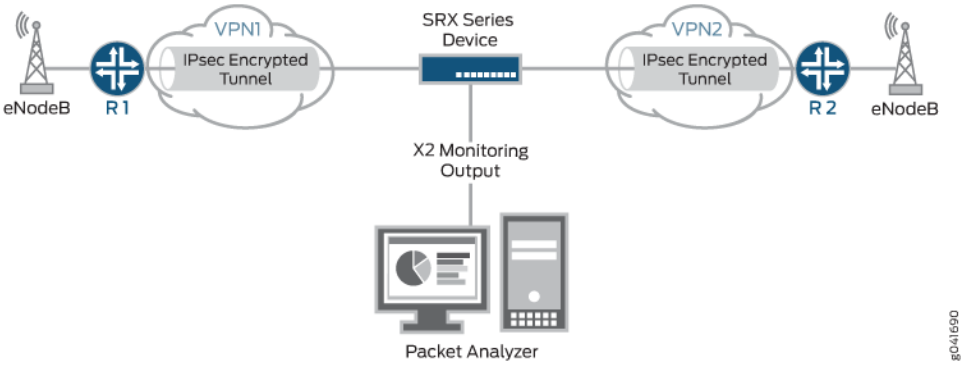

To monitor the X2 interface, you configure mirror filters on the SRX Series device and filter the traffic that you want to analyze. The filtered packets are duplicated and sent to a physical interface. You also specify the output interface on the SRX Series device and the MAC address of the packet analyzer. Because the output interface is connected to the same Layer 2 network as the packet analyzer, once the mirror filtering is turned on, the packet analyzer collects and verifies the X2 traffic.

Copyright © 2016, Juniper Networks, Inc.

Limitations of X2 Traffic Monitoring

The X2 traffic monitoring feature has a limitation—for X2 traffic in a chassis cluster setup, mirrored packets cannot traverse through the data link (fabric interface).

X2 Traffic Terminology

Table 179 lists some X2 traffic related terms and their descriptions.

Table 179: X2 Traffic Terminology

Term	Description
Evolved packet core (EPC)	Main component of System Architecture Evolution (SAE) and is also known as the SAE core. The EPC supports the IP network and serves as the equivalent of a General Packet Radio Service (GPRS) network, using the mobility management entity (MME), Serving Gateway (SGW), and Packet Data Network Gateway (PGW) subcomponents.
Evolved Universal Terrestrial Radio Access Network (E-UTRAN)	A radio access network standard. E-UTRAN is a new air interface system. It provides higher data rates and lower latency and is optimized for packet data. It uses Orthogonal Frequency-Division Multiple Access (OFDMA) for the downlink and Single-carrier Frequency Division Multiple Access for the uplink.
Evolved Node B (eNodeB)	A device connected to the mobile phone network that communicates directly with mobile handsets, like a base transceiver station in Global System for Mobile Communications (GSM) networks. An eNodeB is controlled by a radio network controller (RNC).
Long Term Evolution (LTE)	A standard for wireless communication of high-speed data for mobile phones and data terminals. It increases the capacity and speed using a different radio interface and makes core network improvements.
X2 interface	A point-to-point logical interface between two eNodeBs with the E-UTRAN. It supports the exchange of signaling information between two eNodeBs and supports the forwarding of protocol data units (PDUs) to the respective tunnel endpoints.
X2 Application Protocol (X2AP)	Protocol used by the X2 interface. It is used for handling the user equipment mobility within the E-UTRAN and provides the following functions: • Manages mobility and load • Reports general error situations • Sets and resets the X2 interface • Updates the eNodeB configuration

Related Documentation

- Example: Configuring Mirror Filters for X2 Traffic Monitoring on page 1796
- mirror-filter (Security Forwarding Options) on page 1866

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Mirror Filters for X2 Traffic Monitoring

This example shows how to configure mirror filters to monitor X2 traffic between two eNodeBs in an LTE mobile network.

- Requirements on page 1796
- Overview on page 1796
- Configuration on page 1797
- Verification on page 1798

Requirements

Before you begin:

- Understand X2 traffic monitoring. See "Understanding X2 Traffic Monitoring" on page 1793.
- Configure the interfaces, security zones, security policies, and the route-based VPN tunnels to allow data to be securely transferred between the SRX Series device and the two eNodeBs. See "Example: Configuring a Route-Based VPN" on page 6074.

Overview

In this example, an SRX Series device uses IPsec tunnels to connect two eNodeBs in an LTE mobile network. You, as a network operator, need a way to monitor the X2 traffic to debug any handover issues across eNodeBs. Traffic coming out of an IPsec tunnel is decrypted and then encrypted again to go into the other IPsec tunnel in the opposite direction.

To monitor traffic, you configure mirror filters on the SRX Series device. The traffic that matches your specified filters is mirrored and sent to an output interface that is connected to a packet analyzer (also called a *sniffing* device). The packet analyzer analyzes the traffic allowing you to monitor the X2 traffic.

Figure 94 shows the SRX Series device connecting to the eNodeBs using IPsec tunnels. The SRX Series device is also connected to a packet analyzer.

Figure 94: Configuring Mirror Filters for X2 Traffic Monitoring

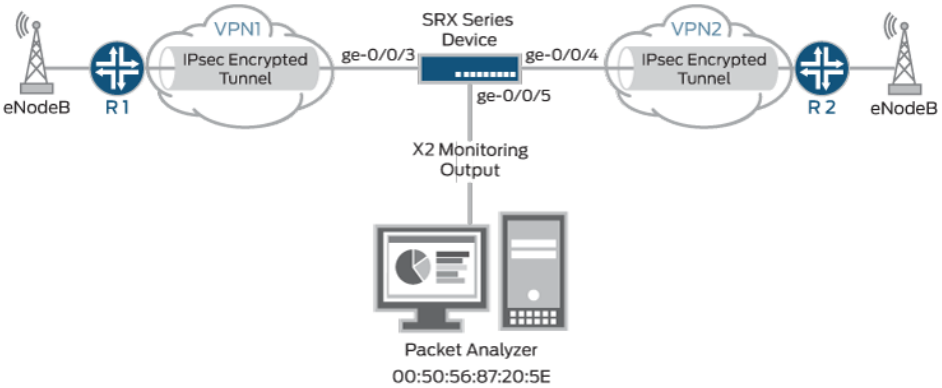

Copyright © 2016, Juniper Networks, Inc.

In this example, you specify the following match conditions for the filters:

- Protocol—ICMP
- Source-prefix—192.0.2.0/24
- Destination-prefix—203.0.113.0
- Incoming logical interface—st0.1
- Outgoing logical interface—st0.2

Packets that match these conditions are mirrored and sent through the output interface ge-0/0/5 to the packet analyzer with a MAC address 00:50:56:87:20:5E. The packet analyzer analyzes traffic that you can use for monitoring and debugging the X2 traffic on the SRX Series device.

Configuration

CLI Quick Configuration

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

```
set security forwarding-options mirror-filter m1 protocol icmp
set security forwarding-options mirror-filter m1 source-prefix 192.0.2.0/24
set security forwarding-options mirror-filter m1 destination-prefix 203.0.113.0
set security forwarding-options mirror-filter m1 interface-in st0.1
set security forwarding-options mirror-filter m1 interface-out st0.2
set security forwarding-options mirror-filter m1 output interface ge-0/0/5
set security forwarding-options mirror-filter m1 output destination-mac 00:50:56:87:20:5E
```

Step-by-Step Procedure

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

To configure mirror filters for monitoring X2 traffic:

1. Create a mirror filter.

   ```
   [edit]
   user@host# edit security forwarding-options mirror-filter m1
   ```

2. Specify the match conditions for the mirror filter.

   ```
   [edit security forwarding-options mirror-filter m1]
   user@host# set protocol icmp
   user@host# set source-prefix 192.0.2.0/24
   user@host# set destination-prefix 203.0.113.0
   user@host# set interface-in st0.1
   user@host# set interface-out st0.2
   ```

3. Specify the output interface for the mirrored packets.

   ```
   [edit security forwarding-options mirror-filter m1]
   user@host# set output interface ge-0/0/5
   ```

Copyright © 2016, Juniper Networks, Inc.

4. Specify the MAC address of the packet analyzer as a destination for all mirrored packets.

```
[edit security forwarding-options mirror-filter m1]
user@host# set output destination-mac 00:50:56:87:20:5E
```

Results From configuration mode, confirm your configuration by entering the **show security forwarding-options** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show security forwarding-options
    mirror-filter m1 {
      protocol icmp;
      source-prefix 192.0.2.0/24;
      destination-prefix 203.0.113.0;
      interface-in st0.1;
      interface-out st0.2;
      output {
        interface ge-0/0/5;
        destination-mac 00:50:56:87:20:5E;
      }
    }
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Status of Mirror Filter

Purpose Verify that mirror filter is active or not.

Action From operational mode, enter the **show security forward-options mirror-filter** command.

```
user@host> show security forward-options mirror-filter m1
Security mirror status

    mirror-filter-name: m1
    interface-in: st0.1
    interface-out: st0.2
    protocol: 1
    source-prefix: 192.0.2.0/24
    destination-prefix: 203.0.113.0
    filter-counters: 2
    output-counters: 2
```

Meaning The output provides the mirror filter status. It shows that a mirror filter named m1 is active. The match conditions for the filter are set—source prefix, destination prefix, protocol and incoming and outgoing interfaces.

This output shows that two packets were matched for mirroring and two packets were sent to the packet analyzer. It also correctly shows that the filter and output counters are the same. This means that the number of packets matched for mirroring and the number of packets sent to the packet analyzer are the same.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- clear security forward-options mirror filter on page 1921

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 27

Configuring Packet-Based Forwarding

- Configuring Selective Stateless Packet-Based Services on page 1803

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 92

Configuring Selective Stateless Packet-Based Services

- Understanding Packet-Based Processing on page 1803
- Understanding Selective Stateless Packet-Based Services on page 1804
- Selective Stateless Packet-Based Services Configuration Overview on page 1806
- Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding on page 1808
- Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding on page 1818

Understanding Packet-Based Processing

Packets that enter and exit a Juniper Networks device running Junos OS can undergo packet-based processing. Packet-based, or stateless, packet processing treats packets discretely. Each packet is assessed individually for treatment. Stateless packet-based forwarding is performed on a packet-by-packet basis without regard to flow or state information. Each packet is assessed individually for treatment.

Figure 95 shows the traffic flow for packet-based forwarding.

Figure 95: Traffic Flow for Packet-Based Forwarding

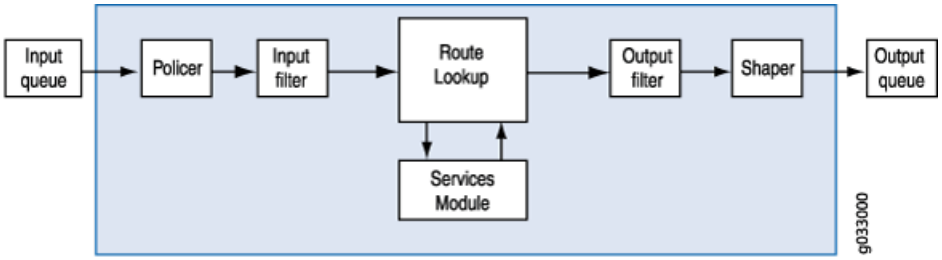

As packets enter the device, classifiers, filters and policers are applied to it. Next, the egress interface for the packet is determined through a route lookup. Once the egress interface for the packet is found, filters are applied and the packet is sent to the egress interface where it is queued and scheduled for transmission.

Packet-based forwarding does not require any information about either previous or subsequent packets that belong to a given connection, and any decision to allow or deny

Copyright © 2016, Juniper Networks, Inc.

traffic is packet specific. This architecture has the benefit of massive scaling because it forwards packets without keeping track of individual flows or state.

On SRX Series devices, the default mode for processing traffic is flow mode. To configure an SRX Series device as a border router, you must change the mode from flow-based processing to packet-based processing. Use the **set security forwarding-options family mpls mode packet-based** statement to configure the SRX device to packet mode. You must reboot the device for the configuration to take effect.

Related Documentation

- Understanding Selective Stateless Packet-Based Services on page 1804
- Selective Stateless Packet-Based Services Configuration Overview on page 1806
- Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding on page 1808
- Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding on page 1818
- **forwarding-options (Security) on page 3971**

Understanding Selective Stateless Packet-Based Services

Selective stateless packet-based services allow you to use both flow-based and packet-based forwarding simultaneously on a system. You can selectively direct traffic that requires packet-based, stateless forwarding to avoid stateful flow-based forwarding by using stateless firewall filters, also known as access control lists (ACLs). The traffic not so directed follows the default flow-based forwarding path. Bypassing flow-based forwarding can be useful for traffic for which you explicitly want to avoid flow session-scaling constraints.

By default, Juniper Networks devices running Junos OS use flow-based forwarding. Selective stateless packet-based services allows you to configure the device to provide only packet-based processing for selected traffic based on input filter terms. Other traffic is processed for flow-based forwarding. Bypassing flow-based forwarding is useful for deployments where you want to avoid session-scaling constraints and session creation and maintenance costs.

When you configure the device for selective stateless packet-based processing, packets entering the system are treated differently depending on certain conditions:

- If a packet satisfies matching conditions specified in input filter terms, it is marked for packet mode and all configured packet mode features are applied to it. No flow-based security features are applied. It bypasses them.
- If a packet has not been flagged for packet-mode, it undergoes normal processing. All services except for MPLS can be applied to this traffic.

Figure 96 shows traffic flow with selective stateless packet-based services bypassing flow-based processing.

Copyright © 2016, Juniper Networks, Inc.

Figure 96: Traffic Flow with Selective Stateless Packet-Based Services

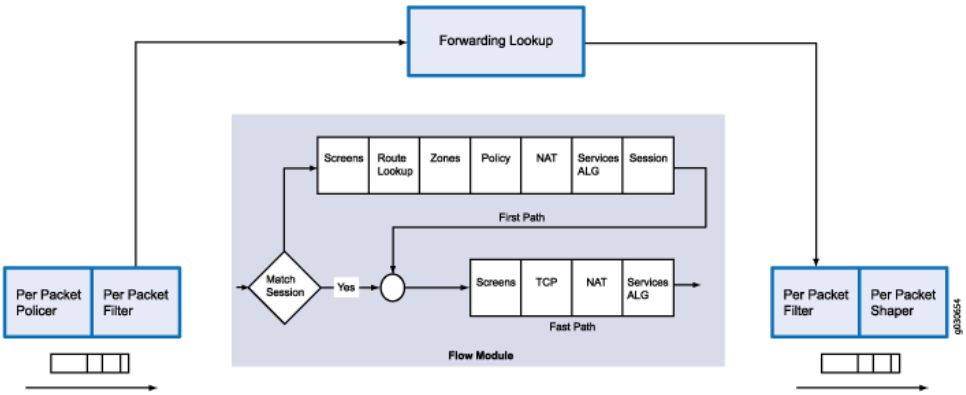

When the packet comes in on an interface, the input packet filters configured on the interface are applied.

- If the packet matches the conditions specified in the firewall filter, a **packet-mode** action modifier is set to the packet. The packet-mode action modifier updates a bit field in the packet key buffer—this bit field is used to determine if the flow-based forwarding needs to be bypassed. As a result, the packet with the packet-mode action modifier bypasses the flow-based forwarding completely. The egress interface for the packet is determined through a route lookup. Once the egress interface for the packet is found, filters are applied and the packet is sent to the egress interface where it is queued and scheduled for transmission.
- If the packet does not match the conditions specified in this filter term, it is evaluated against other terms configured in the filter. If, after all terms are evaluated, a packet matches no terms in a filter, the packet is silently discarded. To prevent packets from being discarded, you configure a term in the filter specifying an action to accept all packets.

A defined set of stateless services is available with selective stateless packet-based services:

- IPv4 routing (unicast and multicast protocols)
- Class of service (CoS)
- Link fragmentation and interleaving (LFI)
- Generic routing encapsulation (GRE)
- Layer 2 switching
- Multiprotocol Label Switching (MPLS)
- Stateless firewall filters
- Compressed Real-Time Transport Protocol (CRTP)

Although traffic requiring MPLS services must be processed in packet mode, under some circumstances it might be necessary to concurrently apply certain services to this traffic that can only be provided in flow mode, such as stateful inspection, NAT, and IPsec. To

Copyright © 2016, Juniper Networks, Inc.

direct the system to process traffic in both flow and packet modes, you must configure multiple routing instances connected through a tunnel interface. One routing instance must be configured to process the packets in flow mode and the other routing instance must be configured to process the packets in packet mode. When you use a tunnel interface to connect routing instances, traffic between those routing instances is injected again into the forwarding path and it can then be reprocessed using a different forwarding method.

Related Documentation

- Understanding Packet-Based Processing on page 1803
- Selective Stateless Packet-Based Services Configuration Overview on page 1806
- Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding on page 1808
- Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding on page 1818

Selective Stateless Packet-Based Services Configuration Overview

You configure selective stateless packet-based services using the stateless firewall filters, also known as access control lists (ACLs). You classify traffic for packet-based forwarding by specifying match conditions in the firewall filters and configure a **packet-mode** action modifier to specify the action. Once match conditions and actions are defined, firewall filters are applied to relevant interfaces.

To configure a firewall filter:

1. Define the address family—First define the address family of the packets that a firewall filter matches. To define the family name, specify **inet** to filter IPv4 packets. Specify **mpls** to filter MPLS packets. Specify **ccc** to filter Layer 2 switching cross-connects.
2. Define terms—Define one or more terms that specify the filtering criteria and the action to take if a match occurs. Each term consists of two components—match conditions and actions.
   - Match conditions—Specify certain characteristics that the packet must match for the action to be performed. You can define various match conditions, such as the IP source address field, IP destination address field, and IP protocol field.
   - Action—Specify what is to be done with the packet if it matches the match conditions. Possible actions are to accept, discard, or reject a packet; go to the next term; or take no action.

     You can specify only one **action** statement (or omit it) in a term, but you can specify any combination of action modifiers with it. Action modifiers include a default **accept** action. For example, if you specify an action modifier and do not specify an action, the specified action modifier is implemented and the packet is accepted.

     The **packet-mode** action modifier specifies traffic to bypass flow-based forwarding. Like other action modifiers, you can configure the **packet-mode** action modifier along with other actions, such as **accept** or **count**.

Copyright © 2016, Juniper Networks, Inc.

3. Apply firewall filters to interfaces—Apply the firewall filter to the interface to have the firewall filter take effect.

When the packet comes in on an interface, the input packet filters configured on the interface are applied. If the packet matches the specified conditions and **packet-mode** action is configured, the packet bypasses the flow-based forwarding completely.

When configuring filters, be mindful of the order of the terms within the firewall filter. Packets are tested against each term in the order in which it is listed in the configuration. When the first matching conditions are found, the action associated with that term is applied to the packet and the evaluation of the firewall filter ends, unless the **next term** action modifier is included. If the **next term** action is included, the matching packet is then evaluated against the next term in the firewall filter; otherwise, the matching packet is not evaluated against subsequent terms in the firewall filter.

When configuring firewall filters for selective stateless packet-based services:

- Accurately identify traffic that needs to bypass flow to avoid unnecessary packet drops.
- Make sure to apply the firewall filter with packet-mode action on all interfaces involved in the packet-based flow path.
- Make sure to configure host-bound TCP traffic to use flow-based forwarding—exclude this traffic when specifying match conditions for the firewall filter term containing the **packet-mode** action modifier. Any host-bound TCP traffic configured to bypass flow is dropped. Asynchronous flow-mode processing is not supported with selective stateless packet-based services.
- Configure input packet filters (not output) with the **packet-mode** action modifier.

NOTE: **Nested firewall filters (configuring a filter within the term of another filter) are not supported with selective stateless packet-based services.**

Some typical deployment scenarios where you can configure selective stateless packet-based services are as follows:

- Traffic flow between private LAN and WAN interfaces, such as for Intranet traffic, where end-to-end forwarding is packet-based
- Traffic flow between private LAN and not-so-secure WAN interfaces, where traffic uses packet-based and flow-based forwarding for secure and not so secure traffic respectively
- Traffic flow between the private LAN and WAN interface with failover to flow-based IPsec WAN when the private WAN link is down
- Traffic flow from flow-based LAN to packet-based MPLS WAN

Related Documentation

- Understanding Packet-Based Processing on page 1803
- Understanding Selective Stateless Packet-Based Services on page 1804

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding on page 1808
- Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding on page 1818

Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding

This example shows how to configure selective stateless packet-based services for end-to-end packet-based forwarding.

- Requirements on page 1808
- Overview on page 1808
- Configuration on page 1809
- Verification on page 1815

Requirements

Before you begin:

- Understand how to configure stateless firewall filters. See the *Junos OS Routing Protocols Library for Security Devices*.
- Establish basic connectivity. .

Overview

In this example, you configure the IP addresses for the interfaces on each of the devices. For R0 it is 10.1.1.2/24 ; for R1 they are 10.1.1.1/24, 10.2.1.1/24, and 203.0.113.1/30; for R2 it is 203.0.113.2/30; and for R3 it is 10.2.1.2/24. You create static routes and associate next-hop addresses for the devices as follows: R0 is 10.1.1.2, R1 is 198.51.100.2, R2 is 203.0.113.1, and R3 is 10.2.1.1.

Then on device R1 you configure a zone called untrust and assign it to interface ge-0/0/3. You also create a zone called trust and assign interfaces ge-0/0/1 and ge-0/0/2 to it. You configure trust and untrust zones to allow all supported application services as inbound services. You allow traffic from any source address, destination address, and application to pass between the zones.

You then create the firewall filter bypass-flow-filter and define the terms bypass-flow-term-1 and bypass-flow-term-2 that match the traffic between internal interfaces ge-0/0/1 and ge-0/0/2 and that contain the packet-mode action modifier. You define the term accept-rest to accept all remaining traffic. Finally, you apply the firewall filter bypass-flow-filter to internal interfaces ge-0/0/1 and ge-0/0/2 (not on the external interface). As a result, all internal traffic bypasses flow-based forwarding and the traffic to and from the Internet does not bypass flow-based forwarding.

Figure 97 shows the network topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 97: Intranet Traffic Using End-to-End Packet-Based Services

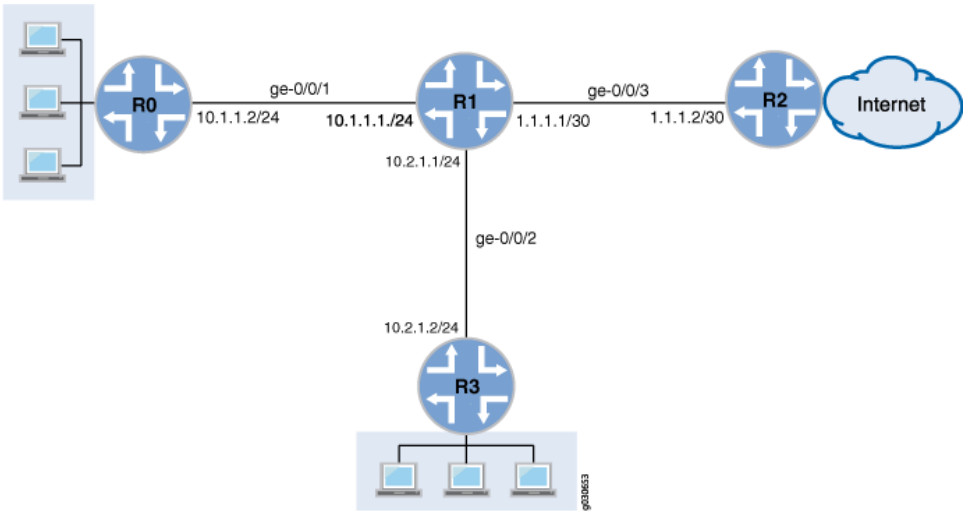

Your company's branch offices are connected to each other through a private WAN. For this internal traffic, packet forwarding is required because security is not an issue. Hence for this traffic, you decide to configure selective stateless packet-based services to bypass flow-based forwarding. The remaining traffic, to and from the Internet, uses flow-based forwarding.

Configuration

CLI Quick Configuration

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

{device R0}
[edit]
set interfaces description "Internal 1" ge-0/0/1 unit 0 family inet address 10.1.1.2/24
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1

{device R1}
set interfaces description "Internal 1" ge-0/0/1 unit 0 family inet address 10.1.1.1/24
set interfaces description "Internal 2" ge-0/0/2 unit 0 family inet address 10.2.1.1/24
set interfaces description "Internet" ge-0/0/3 unit 0 family inet address 203.0.113.1/30
set routing-options static route 0.0.0.0/0 next-hop 203.0.113.2/30
set security zones security-zone untrust interfaces ge-0/0/3
set security zones security-zone trust interfaces ge-0/0/1
set security zones security-zone trust interfaces ge-0/0/2
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic system-services all
set security policies from-zone trust to-zone untrust policy Internet-traffic match
source-address any destination-address any application any
set security policies from-zone trust to-zone untrust policy Internet-traffic then permit
set security policies from-zone untrust to-zone trust policy Incoming-traffic match
source-address any destination-address any application any
set security policies from-zone untrust to-zone trust policy Incoming-traffic then permit
set security policies from-zone trust to-zone trust policy Intrazone-traffic match
source-address any destination-address any application any

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone trust policy Intrazone-traffic then permit
set firewall family inet filter bypass-flow-filter term bypass-flow-term-1 from
  source-address 10.1.1.0/24
set firewall family inet filter bypass-flow-filter term bypass-flow-term–1 from
  destination-address 10.2.1.0/24
set firewall family inet filter bypass-flow-filter term bypass-flow-term-1 then packet-mode
set firewall family inet filter bypass-flow-filter term bypass-flow-term-2 from
  source-address 10.2.1.0/24
set firewall family inet filter bypass-flow-filter term bypass-flow-term-2 from
  destination-address 10.1.1.0/24
set firewall family inet filter bypass-flow-filter term bypass-flow-term-2 then packet-mode
set firewall family inet filter bypass-flow-filter term accept-rest then accept
set interfaces description “Internal 1” ge-0/0/1 unit 0 family inet filter input
  bypass-flow-filter
set interfaces description “Internal 2” ge-0/0/2 unit 0 family inet filter input
  bypass-flow-filter

{device R2}
set interfaces description “Internet” ge-0/0/3 unit 0 family inet address 10.1.1.2/30
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1

{device R3}
[edit]
set interfaces description “Internal 2” ge-0/0/2 unit 0 family inet address 10.2.1.2/24
set routing-options static route 0.0.0.0/0 next-hop 10.2.1.1
```

Step-by-Step Procedure

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

To configure selective stateless packet-based services for end-to-end packet-based forwarding:

1. Configure the IP addresses for the interfaces on devices R0, R1, R2, and R3.

```
{device R0}
[edit]
user@host# set interfaces description “Internal 1” ge-0/0/1 unit 0 family inet address
  10.1.1.2/24

{device R1}
[edit]
user@host# set interfaces description “Internal 1” ge-0/0/1 unit 0 family inet address
  10.1.1.1/24
user@host# set interfaces description “Internal 2” ge-0/0/2 unit 0 family inet address
  10.2.1.1/24
user@host# set interfaces description “Internet” ge-0/0/3 unit 0 family inet address
  203.0.113.1/30

{device R2}
[edit]
user@host# set interfaces description “Internet” ge-0/0/3 unit 0 family inet address
  203.0.113.2/30

{device R3}
[edit]
user@host# set interfaces description “Internal 2” ge-0/0/2 unit 0 family inet address
  10.2.1.2/24
```


Copyright © 2016, Juniper Networks, Inc.

2. Create static routes and associate the appropriate next-hop addresses for devices R0, R1, R2, and R3.

```
{device R0}
[edit]
user@host# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1

{device R1}
[edit]
user@host# set routing-options static route 0.0.0.0/0 next-hop 203.0.113.1

{device R2}
[edit]
user@host# set routing-options static route 0.0.0.0/0 next-hop 203.0.113.2

{device R3}
[edit]
user@host# set routing-options static route 0.0.0.0/0 next-hop 10.2.1.1
```

3. Configure security zones and assign interfaces.

```
{device R1}
[edit]
user@host# set security zones security-zone untrust interfaces ge-0/0/3
user@host# set security zones security-zone trust interfaces ge-0/0/1
user@host# set security zones security-zone trust interfaces ge-0/0/2
```

4. Configure application services for zones.

```
{device R1}
[edit]
user@host# set security zones security-zone trust host-inbound-traffic
  system-services all
user@host# set security zones security-zone untrust host-inbound-traffic
  system-services all
```

5. Configure a security policy

```
{device R1}
[edit]
user@host# set security policies from-zone trust to-zone untrust policy
  Internet-traffic match source-address any destination-address any application
  any
user@host# set security policies from-zone trust to-zone untrust policy
  Internet-traffic then permit
user@host# set security policies from-zone untrust to-zone trust policy
  Incoming-traffic match source-address any destination-address any application
  any
user@host# set security policies from-zone untrust to-zone trust policy
  Incoming-traffic then permit
user@host# set security policies from-zone trust to-zone trust policy Intrazone-traffic
  match source-address any destination-address any application any
user@host# set security policies from-zone trust to-zone trust policy Intrazone-traffic
  then permit
```

6. Create a firewall filter and define terms for all the packet-based forwarding traffic.

```
{device R1}
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term-1
  from source-address 10.1.1.0/24
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term–1
  from destination-address 10.2.1.0/24
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term-1
  then packet-mode
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term-2
  from source-address 10.2.1.0/24
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term-2
  from destination-address 10.1.1.0/24
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term-2
  then packet-mode
```

7. Specify another term for the remaining traffic.

```
{device R1}
[edit]
user@host# set firewall family inet filter bypass-flow-filter term accept-rest then
  accept
```

8. Apply the firewall filter to relevant interfaces.

```
{device R1}
[edit]
user@host# set interfaces description “Internal 1” ge-0/0/1 unit 0 family inet filter
  input bypass-flow-filter
user@host# set interfaces description “Internal 2” ge-0/0/2 unit 0 family inet filter
  input bypass-flow-filter
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show firewall** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
{device R0}
[edit]
user@host# show interfaces
ge-0/0/1 {
  description “Internal 1”
  unit 0 {
    family inet {
      address 10.1.1.2/24
    }
  }
}
```

```
{device R0}
[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 10.1.1.1;
}
```

```
{device R2}
[edit]
user@host# show interfaces
ge-0/0/3 {
```

 Copyright © 2016, Juniper Networks, Inc.

```
  description “Internet”
  unit 0 {
    family inet {
      address 203.0.113.2/30;
    }
  }
}
```

```
{device R2}
[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 203.0.113.1;
}
```

```
{device R3}
[edit]
user@host# show interfaces
ge-0/0/2 {
  description “Internal 2”
  unit 0 {
    family inet {
      address 10.2.1.2/24;
    }
  }
}
```

```
{device R3}
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 10.2.1.1;
}
```

```
{device R1}
[edit]
user@host# show interfaces
ge-0/0/1 {
  description “internal 1”
  unit 0 {
    family inet {
      filter {
        input bypass-flow-filter;
        }
      address 10.1.1.1/24;
    }
  }
}
ge-0/0/2 {
  description “Internal 2”
  unit 0 {
    family inet {
      filter {
        input bypass-flow-filter;
        }
      address 10.2.1.1/24;
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
ge-0/0/3 {
  description “Internet”
  unit 0 {
    family inet {
      address 203.0.113.1/30;
    }
  }
}
{device R1}
[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 203.0.113.1;
}
{device R1}
[edit]
user@host# show firewall
family inet {
  filter bypass-flow-filter {
    term bypass-flow-term-1 {
      from {
        source-address {
          10.1.1.0/24;
        }
        destination-address {
          10.2.1.0/24;
        }
      }
      then packet-mode;
    }
    term bypass-flow-term-2 {
      from {
        source-address {
          10.2.1.0/24;
        }
        destination-address {
          10.1.1.0/24;
        }
      }
      then packet-mode;
    }
    term accept-rest {
      then accept;
    }
  }
}
```

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

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying the End-to-End Packet-Based Configuration on page 1815
- Verifying Session Establishment on Intranet Traffic on page 1815
- Verifying Session Establishment on Internet Traffic on page 1816

Verifying the End-to-End Packet-Based Configuration

Purpose Verify that the selective stateless packet-based services are configured.

Action From configuration mode, enter the **show interfaces**, **show routing-options**, **show security zones**, **show security policies**, and **show firewall** commands.

Verify that the output shows the intended configuration of the firewall filter, interfaces, and policies.

Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the **insert** command.

Verifying Session Establishment on Intranet Traffic

Purpose Verify that sessions are established when traffic is transmitted to interfaces within the Intranet.

Action To verify that sessions are established, perform the following tasks:

1. On device **R1**, enter the operational mode **clear security flow session all** command to clear all existing security flow sessions.
2. On device **R0**, enter the operational mode **ping** command to transmit traffic to device **R3**.
3. On device **R1**, with traffic transmitting from devices **R0** to **R3** through **R1**, enter the operational mode **show security flow session** command.

```
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

NOTE: To verify established sessions, make sure to enter the **show security flow session** command while the **ping** command is sending and receiving packets.

Copyright © 2016, Juniper Networks, Inc.

```
  {device R0}
  user@host> ping 192.0.2.2 -c 10

PING 192.0.2.2 (192.0.2.2) 56(84) bytes of data.
64 bytes from 192.0.2.2: icmp_seq=1 ttl=63 time=6.07 ms
64 bytes from 192.0.2.2: icmp_seq=2 ttl=63 time=4.24 ms
64 bytes from 192.0.2.2: icmp_seq=3 ttl=63 time=2.85 ms
64 bytes from 192.0.2.2: icmp_seq=4 ttl=63 time=6.14 ms
...

  {device R1}
  user@host>show security flow session

Flow Sessions on FPC10 PIC1:

Session ID: 410000077, Policy name: Internet-traffic/5, Timeout: 2, Valid
  In: 198.51.100.1/3 --> 192.0.2.2/32055;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes:
84, CP Session ID: 410000198
  Out: 192.0.2.2/32055 --> 198.51.100.1/3;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes:
84, CP Session ID: 410000198
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000079, Policy name: Internet-traffic/5, Timeout: 2, Valid
  In: 198.51.100.1/5 --> 192.0.2.2/32055;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes:
84, CP Session ID: 420000163
  Out: 192.0.2.2/32055 --> 198.51.100.1/5;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes:
84, CP Session ID: 420000163
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000090, Policy name: Internet-traffic/5, Timeout: 4, Valid
  In:198.51.100.1/7 --> 192.0.2.2/32055;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
 CP Session ID: 430000088
  Out: 192.0.2.2/32055 --> 198.51.100.1/7;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes:
84, CP Session ID: 430000088
Total sessions: 1
```

The output shows traffic transmitting from **R0** to **R3** and no sessions are established. In this example, you applied the **bypass-flow-filter** with the **packet-mode** action modifier on interfaces **Internal 1** and **Internal 2** for your company's Intranet traffic. This output verifies that the traffic between the two interfaces is correctly bypassing flow-based forwarding and hence no sessions are established.

Verifying Session Establishment on Internet Traffic

Purpose Verify that sessions are established when traffic is transmitted to the Internet.

Action To verify that traffic to the Internet is using flow-based forwarding and sessions are established, perform the following tasks:

1. On device **R1**, enter the operational mode **clear security flow session all** command to clear all existing security flow sessions.

Copyright © 2016, Juniper Networks, Inc.

2. On device **R0**, enter the operational mode **ping** command to transmit traffic to device **R2**.
3. On device **R1**, with traffic transmitting from **R0** to **R2** through **R1**, enter the operational mode **show security flow session** command.

NOTE: To verify established sessions, make sure to enter the **show security flow session** command while the **ping** command is sending and receiving packets.

{device R0}
user@host> **ping 203.0.113.6**

```
PING 203.0.113.6 (203.0.113.6): 56 data bytes
64 bytes from 203.0.113.6: icmp_seq=0 ttl=63 time=2.326 ms
64 bytes from 203.0.113.6: icmp_seq=1 ttl=63 time=2.569 ms
64 bytes from 203.0.113.6: icmp_seq=2 ttl=63 time=2.565 ms
64 bytes from 203.0.113.6: icmp_seq=3 ttl=63 time=2.563 ms
64 bytes from 203.0.113.6: icmp_seq=4 ttl=63 time=2.306 ms
64 bytes from 203.0.113.6: icmp_seq=5 ttl=63 time=2.560 ms
64 bytes from 203.0.113.6: icmp_seq=6 ttl=63 time=4.130 ms
64 bytes from 203.0.113.6: icmp_seq=7 ttl=63 time=2.316 ms
...
```

{device R1}
user@host> **show security flow session**

```
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

The output shows traffic transmitting from devices **R0** to **R1** and established sessions. In this example, you did not apply the **bypass-flow-filter** with the **packet-mode** action modifier on interface **Internet** for your company's Internet traffic. This output verifies that the traffic to the Internet is correctly using flow-based forwarding and hence sessions are established.

Transmit traffic from device **R3** to **R2** and use the commands in this section to verify established sessions.

Related Documentation

- Understanding Selective Stateless Packet-Based Services on page 1804
- Selective Stateless Packet-Based Services Configuration Overview on page 1806
- Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding on page 1818

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding

This example shows how to configure selective stateless packet-based services for packet-based to flow-based forwarding.

- Requirements on page 1818
- Overview on page 1818
- Configuration on page 1819
- Verification on page 1824

Requirements

Before you begin:

- Understand how to configure stateless firewall filters. See the *Junos OS Routing Protocols Library for Security Devices*.
- Establish basic connectivity. .

Overview

In this example, you configure the IP addresses for the interfaces on each of the devices. For device R0 as 198.51.100.9/24; for R1 the are198.51.100.10/24 and 203.0.113.5/24; and for R2 it is 203.0.113.9/24. On device R1, you set an internal service interface lt-0/0/0 between routing instances and configure a peer relationship between two virtual devices. You then create a security zone called HOST, assign all interfaces to it, and configure it to allow all supported applications and protocols.

Then you configure policies and specify that all packets are permitted. You configure a virtual device routing instance Internet-VR and assign interfaces for flow-based forwarding. You enable OSPF on devices R0, R1, and R2. On Device R2, you configure the filter bypass-flow-filter with the term bypass-flow-term that contains the packet-mode action modifier. Because you have not specified any match conditions, this filter applies to all traffic that traverses the interfaces on which it is applied.

Finally, on device R1 you apply the firewall filter bypass-flow-filter to internal interfaces ge-0/0/2.0 and lt-0/0/0.0. You do not apply the filter to the interfaces associated with the Internet-VR routing instance. As a result, all traffic that traverses the LAN interfaces associated with the master routing instance uses packet-based forwarding and all traffic that traverses the Internet-VR routing instance uses flow-based forwarding.

Figure 98 shows the network topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 98: Selective Stateless Packet-Based Services for Packet-Based Forwarding

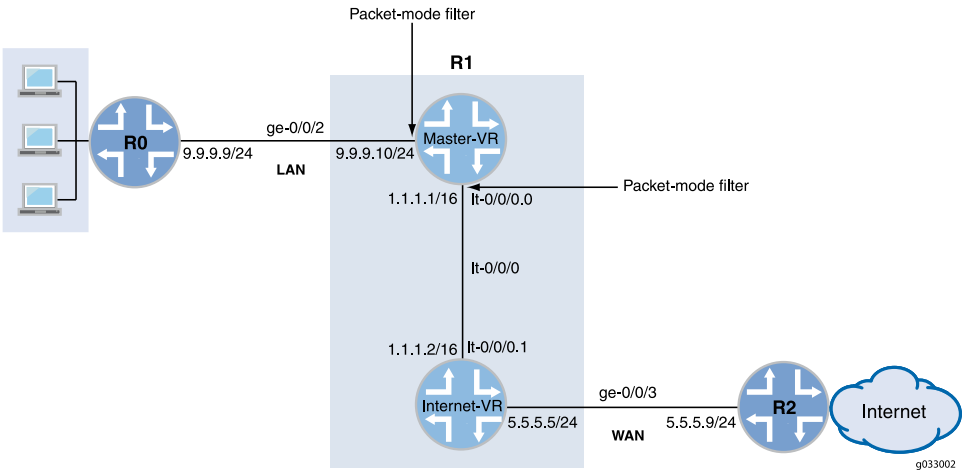

The interface facing the private LAN does not need any security services, but the interface facing the WAN needs security. In this example, you decide to configure both packet-based and flow-based forwarding for secure and not so secure traffic by configuring two routing instances—one handling the packet-based forwarding and the other handling the flow-based forwarding.

Configuration

CLI Quick Configuration

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

```
{device R0}
set interfaces description “Connect to Master VR” ge-0/0/2 unit 0 family inet address
  198.51.100.9/24
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0

{device R1}
set interfaces description “Connect to R0” ge-0/0/2 unit 0 family inet address
  198.51.100.10/24
set interfaces description “Connect to R2” ge-0/0/3 unit 0 family inet address
  203.0.113.5/24
set interfaces lt-0/0/0 unit 0 encapsulation frame-relay dlci 100 peer-unit 1 family inet
  address 192.0.2.1/16
set interfaces lt-0/0/0 unit 1 encapsulation frame-relay dlci 100 peer-unit 0 family inet
  address 192.0.2.2/16
set security zones security-zone HOST host-inbound-traffic system-services any-service
set security zones security-zone HOST host-inbound-traffic protocols all
set security zones security-zone HOST interfaces all
set security policies default-policy permit-all
set routing-instances Internet-VR instance-type virtual-router interface lt-0/0/0.1
set routing-instances Internet-VR instance-type virtual-router interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lt-0/0/0.0
set routing-instances Internet-VR protocols ospf area 0.0.0.0 interface lt-0/0/0.1
```

Copyright © 2016, Juniper Networks, Inc.

```
set routing-instances Internet-VR protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set firewall family inet filter bypass-flow-filter term bypass-flow-term then accept
set firewall family inet filter bypass-flow-filter term bypass-flow-term then packet-mode
set interfaces ge-0/0/2 unit 0 family inet bypass-flow-filter
set interfaces lt-0/0/0 unit 0 family inet bypass-flow-filter

{device R2}
set interfaces description “Connect to Internet-VR” ge-0/0/3 unit 0 family inet address
  203.0.113.9/24
set protocols ospf area 0.0.0.0 interface ge-0/0/3
```

Step-by-Step Procedure

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

To configure selective stateless packet-based services for end-to-end packet-based forwarding:

1. Configure the IP addresses for the interfaces.

```
{device R0}
[edit]
user@host# set interfaces description “Connect to Master VR” ge-0/0/2 unit 0
  family inet address 198.51.100.9/24

{device R1}
[edit]
user@host# set interfaces description “Connect to R0” ge-0/0/2 unit 0 family inet
  address 198.51.100.10/24
user@host# set interfaces description “Connect to R2” ge-0/0/3 unit 0 family inet
  address 203.0.113.5/24

{device R2}
[edit]
user@host# set interfaces description “Connect to Internet-VR” ge-0/0/3 unit 0
  family inet address 203.0.113.9/24
```

2. Set an internal service interface between routing instances.

```
{device R1}
[edit]
user@host# set interfaces lt-0/0/0 unit 0 encapsulation frame-relay dlci 100
  peer-unit 1 family inet address 192.0.2.1/16
user@host# set interfaces lt-0/0/0 unit 1 encapsulation frame-relay dlci 100
  peer-unit 0 family inet address 192.0.2.2/16
```

3. Configure security zones.

```
{device R1}
[edit]
user@host# set security zones security-zone HOST host-inbound-traffic
  system-services any-service
user@host# set security zones security-zone HOST host-inbound-traffic protocols
  all
user@host# set security zones security-zone HOST interfaces all
```

4. Configure policies.

```
{device R1}
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security policies default-policy permit-all
```

5. Configure a virtual device routing instance.

```
{device R1}
[edit]
user@host# set routing-instances Internet-VR instance-type virtual-router interface
  lt-0/0/0.1
user@host# set routing-instances Internet-VR instance-type virtual-router interface
  ge-0/0/3.0
```

6. Enable OSPF on all interfaces in the network.

```
{device R0}
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
```

```
{device R1 for Master-VR}
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
user@host# set protocols ospf area 0.0.0.0 interface lt-0/0/0.0
```

```
{device R1 for Internet-VR}
[edit]
user@host# set routing-instances Internet-VR protocols ospf area 0.0.0.0 interface
  lt-0/0/0.1
user@host# set routing-instances Internet-VR protocols ospf area 0.0.0.0 interface
  ge-0/0/3.0
```

```
{device R2}
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/3
```

7. Create a firewall filter and define a term for packet-based forwarding traffic.

```
{device R1}
[edit]
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term
  then accept
user@host# set firewall family inet filter bypass-flow-filter term bypass-flow-term
  then packet-mode
```

8. Apply the firewall filter to relevant interfaces.

```
{device R1}
[edit]
user@host# set interfaces ge-0/0/2 unit 0 family inet bypass-flow-filter
user@host# set interfaces lt-0/0/0 unit 0 family inet bypass-flow-filter
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show security**, **show routing-instances**, and **show firewall** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
{device R0}
[edit]
user@host# show interfaces
ge-0/0/2 {
```

Copyright © 2016, Juniper Networks, Inc.

```
  description “Connect to Master-VR”
  unit 0 {
    family inet {
      address 198.51.100.9/24
    }
  }
}
```

```
{device R0}
[edit]
user@host# show protocols
ospf {
  area 0.0.0.0/0 {
    interface ge-0/0/2.0;
  }
}
```

```
{device R2}
[edit]
user@host# show interfaces
ge-0/0/3 {
  description “Connect to Internet-VR”
  unit 0 {
    family inet {
      address 203.0.113.9/24;
    }
  }
}
```

```
{device R2}
[edit]
user@host# show protocols
ospf {
  area 0.0.0.0/0 {
    interface ge-0/0/3.0;
  }
}
```

```
{device R1}
[edit]
user@host# show interfaces
ge-0/0/2 {
  description “Connect to R0”
  unit 0 {
    family inet {
      filter {
        input bypass-flow-filter;
        }
      address 198.51.100.10/24;
    }
  }
}
lt-0/0/0 {
  unit 0 {
    encapsulation frame-relay;
    dlci 100;
    peer-unit 1;
```


Copyright © 2016, Juniper Networks, Inc.

```
      family inet {
        filter {
          input bypass-flow-filter
        }
        address 192.0.2.1/16;
      }
    }
    unit 1{
      encapsulation frame-relay;
      dlci 100;
      peer-unit 0;
      family inet {
        address 192.0.2.2/16 ;
      }
    }
}
{device R1}
[edit]
user@host# show protocols
ospf {
  area 0.0.0.0/0 {
    interface ge-0/0/2.0;
    interface lt-0/0/0.0;
  }
}
{device R1}
[edit]
user@host# show firewall
filter bypass-flow-filter {
  term bypass-flow-term {
    then {
      packet-mode;
      accept;
    }
  }
}

{device R1}
[edit]
user@host# show routing-instances
Internet-VR {
  instance-type virtual-router;
  interface lt-0/0/0.1;
  interface ge-0/0/3.0;
  protocols {
    ospf {
      area 0.0.0.0 {
        interface ge-0/0/3.0;
        lt-0/0/0.1;
      }
    }
  }
}

{device R1}
[edit]
user@host# show security
```

```
zones {
  security-zone HOST {
    host-inbound-traffic {
      system-services {
        any-service;
      }
      protocols {
        all;
      }
    }
    interfaces {
      all;
    }
  }
}
policies {
  default-policy {
    permit-all;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Packet-Based to Flow-Based Configuration on page 1824
- Verifying Session Establishment on LAN Traffic on page 1824
- Verifying Session Establishment on Internet Traffic on page 1825

Verifying the Packet-Based to Flow-Based Configuration

Purpose Verify that the selective stateless packet-based services are configured for packet-based to flow-based forwarding.

Action From configuration mode, enter the **show interfaces**, **show protocols**, **show security**, **show routing-instances**, and **show firewall** commands.

Verify that the output shows the intended configuration of the firewall filter, routing instances, interfaces, and policies.

Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the **insert** command.

Verifying Session Establishment on LAN Traffic

Purpose Verify that the sessions are established when traffic is transmitted on interfaces within the LAN.

Copyright © 2016, Juniper Networks, Inc.

Action To verify that sessions are established, perform the following tasks:

1. On device **R1**, from operational mode enter the **clear security flow session all** command to clear all existing security flow sessions.
2. On device **R0**, from operational mode enter the **ping** command to transmit traffic to device **Master-VR**.
3. On device **R1**, with traffic transmitting from devices **R0** through **R1**, from operational mode enter the **show security flow session** command.

NOTE: To verify established sessions, ensure that you enter the **show security flow session** command while the **ping** command is sending and receiving packets.

{device R0}
user@host> **ping 192.0.2.1**

```
PING 192.0.2.1 (192.0.2.1): 56 data bytes
64 bytes from 192.0.2.1: icmp_seq=0 ttl=63 time=2.208 ms
64 bytes from 192.0.2.1: icmp_seq=1 ttl=63 time=2.568 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=63 time=2.573 ms
64 bytes from 192.0.2.1: icmp_seq=3 ttl=63 time=2.310 ms
64 bytes from 192.0.2.1: icmp_seq=4 ttl=63 time=1.566 ms
64 bytes from 192.0.2.1: icmp_seq=5 ttl=63 time=1.569 ms
...
```

{device R1}
user@host> **show security flow session**

```
0 sessions displayed
```

The output shows traffic transmitting from **R0** to **Master-VR** and no sessions are established. In this example, you applied the **bypass-flow-filter** with the **packet-mode** action modifier on interfaces **ge-0/0/0** and **lt-0/0/0.0** for your company's LAN traffic. This output verifies that the traffic between the two interfaces is correctly bypassing flow-based forwarding and hence no sessions are established.

Verifying Session Establishment on Internet Traffic

Purpose Verify that sessions are established when traffic is transmitted to the Internet.

Action To verify that traffic to the Internet is using flow-based forwarding and sessions are established, perform the following tasks:

1. On device **R1**, from operational mode enter the **clear security flow session all** command to clear all existing security flow sessions.
2. On device **R0**, from operational mode enter the **ping** command to transmit traffic to device **R2**.

Copyright © 2016, Juniper Networks, Inc.

3. On device **R1**, with traffic transmitting from **R0** to **R2** through **R1**, from operational mode enter the **show security flow session** command.

```
root@host> show security flow session
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

NOTE: To verify established sessions, ensure that you enter the **show security flow session** command while the **ping** command is sending and receiving packets.

{device R0}
user@host> **ping 192.0.2.1 -c 10**

```
PING 60.0.0.1 (60.0.0.1) 56(84) bytes of data.
64 bytes from 192.0.2.1: icmp_seq=1 ttl=64 time=1.98 ms
64 bytes from 192.0.2.1: icmp_seq=2 ttl=64 time=1.94 ms
64 bytes from 192.0.2.1: icmp_seq=3 ttl=64 time=1.92 ms
64 bytes from 192.0.2.1: icmp_seq=4 ttl=64 time=1.89 ms

...
```

{device R1}
user@host> **show security flow session**

```
Session ID: 189900, Policy name: default-policy/2, Timeout: 2
  In: 198.51.100.1/0 --> 192.0.2.1/5924;icmp, If: lt-0/0/0.1
  Out: 192.0.2.1/5924 --> 198.51.100.1/0;icmp, If: ge-0/0/3.0

Session ID: 189901, Policy name: default-policy/2, Timeout: 2
  In: 198.51.100.1/1 --> 192.0.2.1/5924;icmp, If: lt-0/0/0.1
  Out: 192.0.2.1/5924 --> 198.51.100.1/1;icmp, If: ge-0/0/3.0

Session ID: 189902, Policy name: default-policy/2, Timeout: 4
  In: 198.51.100.1/2 --> 192.0.2.1/5924;icmp, If: lt-0/0/0.1
  Out: 192.0.2.1/5924 --> 198.51.100.1/2;icmp, If: ge-0/0/3.0

 3 sessions displayed
```

The output shows traffic transmitting from devices **R0** to **R2** and established sessions. In this example, you did not apply the **bypass-flow-filter** with the **packet-mode** action modifier on routing instance **Internet-VR** for your company's Internet traffic. This output verifies that the traffic to the Internet is correctly using flow-based forwarding and hence sessions are established.

Note that sessions are established only when traffic is flowing between **lt-0/0/0.1** and **ge-0/0/3** and not when traffic is flowing between **ge-0/0/2** and **lt-0/0/0.0**.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Selective Stateless Packet-Based Services on page 1804
- Selective Stateless Packet-Based Services Configuration Overview on page 1806
- Example: Configuring Selective Stateless Packet-Based Services for End-to-End Packet-Based Forwarding on page 1808

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 28

Configuration Statements and Operational Commands

- Configuration Statements on page 1831
- Operational Commands on page 1895

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 93

Configuration Statements

- Chassis Configuration Statement Hierarchy on page 1833
- Security Configuration Statement Hierarchy on page 1836
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-process] Hierarchy Level on page 1839
- aging on page 1839
- all-tcp on page 1840
- allow-dns-reply on page 1840
- allow-embedded-icmp on page 1841
- application-services (Security Forwarding Process) on page 1842
- apply-to-half-close-state on page 1843
- ethernet-switching on page 1844
- destination-header on page 1845
- destination-port (Security Forwarding Options) on page 1846
- destination-prefix (Security Forwarding Options) on page 1850
- early-ageout on page 1850
- error on page 1851
- fin-invalidate-session on page 1852
- force-ip-reassembly on page 1853
- forwarding-process on page 1853
- fpc error on page 1854
- fru-poweron-sequence on page 1856
- gre-in on page 1857
- gre-out on page 1858
- high-watermark on page 1858
- hop-by-hop-header on page 1859
- icmpv6-malformed on page 1860
- inline-tap on page 1860
- interface-in (Security Forwarding Options) on page 1861

Copyright © 2016, Juniper Networks, Inc.

- interface-out (Security Forwarding Options) on page 1861
- ipv4-template (Services) on page 1861
- ipv6-extension-header on page 1862
- ipv6-extension-header-limit on page 1863
- ipv6-malformed-header on page 1863
- ipv6-template (Services) on page 1864
- low-latency on page 1864
- low-watermark on page 1865
- maximize-idp-sessions on page 1865
- mirror-filter (Security Forwarding Options) on page 1866
- mode (Security Forwarding Options) on page 1867
- no-sequence-check on page 1868
- no-tcp-reset on page 1868
- output (Security Forwarding Options) on page 1869
- packet-filter on page 1870
- packet-ordering-mode (Application Services) on page 1871
- pending-sess-queue-length on page 1872
- propagate-settings on page 1872
- protocol (Security Forwarding Options) on page 1873
- resource-manager on page 1874
- route-change-timeout on page 1874
- rst-invalidate-session on page 1875
- rst-sequence-check on page 1875
- sampling on page 1876
- services-offload on page 1877
- np-cache (Flexible PIC Concentrator) on page 1878
- source-port (Security Forwarding Options) on page 1879
- source-prefix (Security Forwarding Options) on page 1883
- syn-flood-protection-mode on page 1883
- tcp-initial-timeout on page 1884
- tcp-mss (Security Flow) on page 1885
- tcp-session on page 1886
- time-wait-state on page 1887
- traceoptions (Security) on page 1888
- traceoptions (Security Flow) on page 1890
- transport (Security Log) on page 1893
- weight (Security) on page 1894

Copyright © 2016, Juniper Networks, Inc.

Chassis Configuration Statement Hierarchy

Use the statements in the **chassis** configuration hierarchy to configure alarms, aggregated devices, clusters, the Routing Engine, and other chassis properties.

```
chassis {
    aggregated-devices {
        ethernet {
            device-count number;
            lacp {
                link-protection {
                    non-revertive;
                }
                system-priority number;
            }
        }
        sonet {
            device-count number;
        }
    }
    alarm {
        ds1 {
            ais (ignore | red | yellow);
            ylw (ignore | red | yellow);
        }
        ethernet {
            link-down (ignore | red | yellow);
        }
        integrated-services {
            failure (ignore | red | yellow);
        }
        management-ethernet {
            link-down (ignore | red | yellow);
        }
        serial {
            cts-absent (ignore | red | yellow);
            dcd-absent (ignore | red | yellow);
            dsr-absent (ignore | red | yellow);
            loss-of-rx-clock (ignore | red | yellow);
            loss-of-tx-clock (ignore | red | yellow);
        }
        services {
            hw-down (ignore | red | yellow);
            linkdown (ignore | red | yellow);
            pic-hold-reset (ignore | red | yellow);
            pic-reset (ignore | red | yellow);
            rx-errors (ignore | red | yellow);
            sw-down (ignore | red | yellow);
            tx-errors (ignore | red | yellow);
        }
        t3 {
            ais (ignore | red | yellow);
            exz (ignore | red | yellow);
            ferf (ignore | red | yellow);
```

Copyright © 2016, Juniper Networks, Inc.

```
    idle (ignore | red | yellow);
    lcv (ignore | red | yellow);
    lof (ignore | red | yellow);
    los (ignore | red | yellow);
    pll (ignore | red | yellow);
    ylw (ignore | red | yellow);
  }
}
cluster {
  configuration-synchronize {
    no-secondary-bootup-auto;
  }
  control-link-recovery;
  heartbeat-interval milliseconds;
  heartbeat-threshold number;
  network-management {
    cluster-master;
  }
  redundancy-group  group-number {
    gratuitous-arp-count number;
    hold-down-interval number;
    interface-monitor interface-name {
      weight number;
    }
    ip-monitoring {
      family {
        inet {
          ipv4-address {
            interface {
              logical-interface-name;
              secondary-ip-address ip-address;
            }
            weight number;
          }
        }
      }
      global-threshold number;
      global-weight number;
      retry-count number;
      retry-interval seconds;
    }
    node (0 | 1 ) {
      priority number;
    }
    preempt;
  }
  reth-count number;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (world-readable | no-world-readable);
      size maximum-file-size;
    }
    flag flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
      level {
        (alert | all | critical | debug | emergency | error | info | notice | warning);
      }
      no-remote-trace;
    }
  }
  config-button {
    no-clear;
    no-rescue;
  }
  craft-lockout;
  fpc slot-number {
    offline;
    pic slot-number {
      aggregate-ports;
      framing {
        (e1 | e3 | sdh | sonet | t1 | t3);
      }
      ingress-policer-overhead bytes
      max-queues-per-interface (4 | 8);
      mlfr-uni-nni-bundles number;
      no-multi-rate;
      np-cache;
      port slot-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
        speed (oc12-stm4 | oc3-stm1 | oc48-stm16);
      }
      q-pic-large-buffer (large-scale | small-scale);
      services-offload {
        low-latency;
        per-session-statistics;
      }
      shdsl {
        pic-mode (1-port-atm | 2-port-atm | 4-port-atm | efm);
      }
      sparse-dlcis;
      traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode (egress-only | ingress-and-egress);
      }
      tunnel-queuing;
    }
    services-offload;
  }
  ioc-npc-connectivity {
    ioc slot-number {
      npc (npc-slot-number | none);
    }
  }
  maximum-ecmp (16 | 32 | 64);
  network-services (ethernet | IP);
  routing-engine {
    bios {
      no-auto-upgrade;
    }
```

```
        on-disk-failure {
            disk-failure-action (halt | reboot);
        }
        usb-wwan {
            port 1;
        }
    }
    usb {
        storage {
            disable;
        }
    }
}
```

Related Documentation

- cluster (Chassis) on page 3563
- *ip-monitoring*

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839

Copyright © 2016, Juniper Networks, Inc.

- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

[edit security flow] Hierarchy Level

```
security {
  flow {
    aging {
      early-ageout seconds;
      high-watermark percent;
      low-watermark percent;
    }
    allow-dns-reply;
    allow-embedded-icmp;
    ethernet-switching {
      block-non-ip-all;
      bpdu-vlan-flooding;
      bypass-non-ip-unicast;
      no-packet-flooding {
        no-trace-route;
      }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
```

```
        load distribution {
            session-affinity ipsec;
        }
        pending-sess-queue-length (high | moderate | normal);
        route-change-timeout seconds;
        syn-flood-protection-mode (syn-cookie | syn-proxy);
        tcp-mss {
            all-tcp mss value;
            gre-in {
                mss value;
            }
            gre-out {
                mss value;
            }
            ipsec-vpn {
                mss value;
            }
        }
        tcp-session {
            fin-invalidate-session;
            no-sequence-check;
            no-syn-check;
            no-syn-check-in-tunnel;
            rst-invalidate-session;
            rst-sequence-check;
            strict-syn-check;
            tcp-initial-timeout seconds;
            time-wait-state {
                (session-ageout | session-timeout seconds);
            }
        }
        traceoptions {
            file {
                filename;
                files number;
                match regular-expression;
                (no-world-readable | world-readable);
                size maximum-file-size;
            }
            flag flag;
            no-remote-trace;
            packet-filter filter-name {
                destination-port port-identifier;
                destination-prefix address;
                interface interface-name;
                protocol protocol-identifier;
                source-port port-identifier;
                source-prefix address;
            }
            rate-limit messages-per-second;
        }
    }
}
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Juniper Networks Devices Processing Overview on page 1665
- *Network Monitoring and Troubleshooting Guide for Security Devices*

[edit security forwarding-process] Hierarchy Level

```
security {
  forwarding-process {
    application-services {
      enable-gtpu-distribution;
      maximize-alg-sessions;
      maximize-idp-sessions {
        inline-tap;
        weight (equal | firewall | idp);
      }
      packet-ordering-mode {
        (hardware | software);
      }
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Juniper Networks Devices Processing Overview on page 1665

aging

Syntax

```
aging {
  early-ageout seconds;
  high-watermark percent;
  low-watermark percent;
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5.

Description Direct the device to begin aggressively aging out sessions when the percentage of entries in the session table exceeds the high-watermark setting and then stops when the percentage of sessions falls below the low-watermark setting.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

all-tcp

Syntax all-tcp mss *value*;

Hierarchy Level [edit security flow tcp-mss]

Release Information Statement introduced in Junos OS Release 8.5.

Description Set the TCP maximum segment size (MSS) value to enable MSS override for all TCP packets in network traffic.

Options **mss *value***—TCP MSS value.
Range: 64 through 65,535 bytes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- tcp-mss (Security Flow) on page 1885
- Juniper Networks Devices Processing Overview on page 1665

allow-dns-reply

Syntax allow-dns-reply;

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5.

Description Allow an incoming Domain Name Service (DNS) reply packet without a matched request. By default, if an incoming UDP first-packet has dst-port 53, the device checks the DNS message packet header to verify that the query bit (QR) is 0, which denotes a query message. If the QR bit is 1, which denotes a response message, the device drops the packet, does not create a session, and increments the illegal packet flow counter for the interface. Using the **allow-dns-reply** statement directs the device to skip the check.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

allow-embedded-icmp

Syntax allow-embedded-icmp;

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Allow an embedded ICMP packet to pass through the device even when there is no session match. Once enabled, all packets encapsulated in ICMP pass through and no policy affects this behavior. This feature is useful when you have asymmetric routing in your network and you want to use traceroute and other ICMP applications on your device.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

application-services (Security Forwarding Process)

Syntax

```
application-services {
  enable-gtpu-distribution;
  maximize-alg-sessions;
  maximize-idp-sessions {
    inline-tap;
    weight (equal | firewall | idp);
  }
  packet-ordering-mode {
    (hardware | software);
  }
}
```

Hierarchy Level [edit security forwarding-process]

Release Information Statement introduced in Junos OS Release 9.6. Statement updated in Junos OS Release 10.4. Statement updated in Junos OS Release 15.1X49-D40 with the **enable-gtpu-distribution** option.

Description You can configure the device to switch from an integrated firewall mode to maximize intrusion detection and prevention (IDP) mode to increase the capacity of IDP processing with the **maximize-idp-sessions** option. When you maximize IDP, you are decoupling IDP processes from firewall processes, allowing the device to support the same number of firewall and IDP sessions.

You can configure maximum Application Layer Gateway (ALG) sessions by using the **maximize-alg-sessions** option. By default, the session capacity number for Real-Time Streaming Protocol (RTSP), FTP, and Trivial File Transfer Protocol (TFTP) ALG sessions is 10,000 per flow Services Processing Unit (SPU). You must reboot the device (and its peer in chassis cluster mode) for the configuration to take effect. The **maximize-alg-sessions** option now enables you to increase defaults as follows:

- RTSP, FTP, and TFTP ALG session capacity: 25,000 per flow SPU
- TCP proxy connection capacity: 40,000 per flow SPU

NOTE: **Flow session capacity is reduced to half per flow SPU; therefore the aforementioned capacity numbers will not change on central point flow.**

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

apply-to-half-close-state

Syntax apply-to-half-close-state;

Hierarchy Level [edit security flow tcp-session time-wait-state]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Configure the TCP session timeout in a half-closed state. This enables the system to apply the configured session timeout on receiving only one FIN packet (either client-to-server or server-to-client). When this statement is not configured, the default behavior takes effect—applying the configured TCP session timeout on receiving both the FIN packets. The default session timeout remains 150 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

ethernet-switching

Syntax

```
ethernet-switching {
  block-non-ip-all;
  bpdu-vlan-flooding;
  bypass-non-ip-unicast;
  no-packet-flooding {
    no-trace-route;
  }
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 9.5.

Description Changes default Layer 2 forwarding behavior.

Options

- **block-non-ip-all**—Block all Layer 2 non-IP and non-ARP traffic, including multicast and broadcast traffic.
- **bpdu-vlan-flooding**—Set 802.1D bridge protocol data unit (BPDU) flooding based on VLAN.
- **bypass-non-ip-unicast**—Allow all Layer 2 non-IP traffic to pass through the device.
- **no-packet-flooding**—Stop IP flooding and send ARP or ICMP requests to discover the destination MAC address for a unicast packet.
  - **no-trace-route**—Do not send ICMP requests to discover the destination MAC address for a unicast packet. Only ARP requests are sent. This option only allows the device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnetwork as the ingress IP address.

NOTE: The **block-non-ip-all** and **bypass-non-ip-unicast** options cannot be configured at the same time.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

destination-header

Syntax

```
destination-header {
  ILNP-nonce-option;
  home-address-option;
  line-identification-option;
  tunnel-encapsulation-limit-option;
  user-defined-option-type low | <to high>;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* ip ipv6-extension-header]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Define the IPv6 destination header screen option.

Options

ILNP-nonce-option—Enable the Identifier-Locator Network Protocol nonce screen option.

home-address-option—Enable the home address screen option.

line-identification-option—Enable the line identification screen option.

tunnel-encapsulation-limit-option—Enable the tunnel encapsulation limit screen option.

user-defined-header-type *low* | <to *high*>—Define the type of header range.

Range: 1 through 255.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding IPv6 Support for Screens on page 836
- ipv6-extension-header on page 1862
- hop-by-hop-header on page 1859

Copyright © 2016, Juniper Networks, Inc.

destination-port (Security Forwarding Options)

Syntax destination-port *port-number*;

Hierarchy Level **[edit security forwarding-options mirror-filter *filter-name*]**

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number to be matched for mirroring. You can specify a numeric value or one of the text synonyms listed in Table 180.

Copyright © 2016, Juniper Networks, Inc.

Table 180: Ports Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760

Copyright © 2016, Juniper Networks, Inc.

Table 180: Ports Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162

Copyright © 2016, Juniper Networks, Inc.

Table 180: Ports Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

Copyright © 2016, Juniper Networks, Inc.

destination-prefix (Security Forwarding Options)

Syntax destination-prefix *destination-prefix*;

Hierarchy Level [edit security forwarding-options mirror-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the destination IP prefix or address to be matched for mirroring.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

early-ageout

Syntax early-ageout *seconds*;

Hierarchy Level [edit security flow aging]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the value before the device aggressively ages out a session from its session table.

Options ***seconds*** —Amount of time that elapses before the device aggressively ages out a session.
Range: 1 through 65,535 seconds
Default: 20 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

error

Syntax

```
error {
  (fatal | major | minor) {
  threshold threshold number;
  action (alarm | disable-pfe | offline-pic | log | get-state |offline | reset);
  }
}
```

Hierarchy Level

[edit chassis]

[edit chassis fpc slot-number]

Release Information

Statement introduced in Junos OS Release 15.1X49-D40.

Description

Configure the threshold at which FPC errors will take the action you configure to be performed by the device.

Some devices include an internal framework for detecting and correcting FPC errors that can have the potential to affect services. You can classify FPC errors according to severity, set an automatic recovery action for each severity, and set a threshold (the number of times the error must occur before the action is triggered).

Options

You can configure the threshold for the following severity levels:

NOTE: You cannot change the severity level of an error.

- **fatal**—Fatal error on the FPC. An error that results in the blockage of a considerable amount of traffic across modules is a fatal error.
- **major**—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error.
- **minor**—Minor error on the FPC. An error that results in the loss of a single packet but is fully recoverable is a minor error.
- **threshold *threshold-value***—Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors exceeds the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence exceeds the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold. For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence.

NOTE: You can set the threshold value to 0 for errors with a severity level of minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with a severity level of major or fatal.

Copyright © 2016, Juniper Networks, Inc.

Default: The error count for fatal and major actions is 1. The default error count for minor actions is 10.

Range: 0—429,496,729

The available detection and recovery actions are as follows:

- **alarm**—Raise an alarm.
- **disable-pfe**—Disable the Packet Forwarding Engine interfaces on the FPC.
- **get-state**—Get the current state of the FPC.
- **log**—Generate a log for the event.
- **offline**—Take the FPC offline.
- **offline-pic**—Take the PIC (installed in the FPC) offline.
- **reset**—Reset the FPC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Related Documentation

- fpc error on page 1854

fin-invalidate-session

Syntax

```
fin-invalidate-session;
```

Hierarchy Level

[edit security flow tcp-session]

Release Information

Statement introduced in Junos OS Release 10.4 R13.

Description

Invalidates a TCP session immediately on reception of a FIN/ACK packet. New incoming SYN packets will need to establish a new TCP session.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

force-ip-reassembly

Syntax force-ip-reassembly;

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 11.2.

Description Reassemble all IP fragmented packets before forwarding.

This option is disabled by default. You can disable this option by deleting this flag from the CLI.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

forwarding-process

Syntax

```
forwarding-process {
  application-services {
    enable-gtpu-distribution;
    maximize-alg-sessions;
    maximize-idp-sessions {
      inline-tap;
      weight (equal | firewall | idp);
    }
    packet-ordering-mode {
      (hardware |software);
    }
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.6.

Description If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

fpc error

Syntax

```
fpc slot number {
  error {
    (fatal | major | minor) {
    threshold threshold-value;
    action (alarm | disable-pfe | offline-pic | log | get-state |offline | reset);
    }
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Configure the threshold at which FPC errors will take the action you configure to be performed by the device.

Some devices include an internal framework for detecting and correcting FPC errors that can have the potential to affect services. For each FPC on the device, you can classify errors according to severity, set an automatic recovery action for each severity, and set a threshold (the number of times the error must occur before the action is triggered).

Options You can configure the threshold for the following severity levels:

NOTE: You cannot change the severity level of an error.

- **fatal**—Fatal error on the FPC. An error that results in the blockage of a considerable amount of traffic across modules is a fatal error.
- **major**—Major error on the FPC. An error that results in continuing loss of packet traffic but does not affect other modules is a major error.
- **minor**—Minor error on the FPC. An error that results in the loss of a single packet but is fully recoverable is a minor error.
- **threshold *threshold-value***—Configure the threshold value at which to take action. If the severity level of the error is fatal, the action is carried out only once when the total number of errors exceeds the threshold value. If the severity level of the error is major, the action is carried out once after the occurrence exceeds the threshold. If the severity level is minor, the action is carried out as many times as the value specified by the threshold. For example, when the severity level is minor, and you have configured the threshold value as 10, the action is carried out after the tenth occurrence.

NOTE: You can set the threshold value to 0 for errors with a severity level of minor. This implies that no action is taken for that error. You cannot set the threshold value to 0 for errors with a severity level of major or fatal.

Copyright © 2016, Juniper Networks, Inc.

Default: The error count for fatal and major actions is 1. The default error count for minor actions is 10.

Range: 0—429,496,729

The available detection and recovery actions are as follows:

- **alarm**—Raise an alarm.
- **disable-pfe**—Disable the Packet Forwarding Engine interfaces on the FPC.
- **get-state**—Get the current state of the FPC.
- **log**—Generate a log for the event.
- **offline**—Take the FPC offline.
- **offline-pic**—Take the PIC (installed in the FPC) offline.
- **reset**—Reset the FPC.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- error on page 1851

Copyright © 2016, Juniper Networks, Inc.

fru-poweron-sequence

Syntax fru-poweron-sequence *fru-poweron-sequence*;

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Configure the power-on sequence for FPCs installed in the chassis.

Options **fru-poweron-sequence**—Power-on sequence for the FPCs in the chassis. The numbers indicate the slot number of the FPCs.

NOTE: If the power-on sequence is not configured by including the **fru-poweron-sequence** statement, Junos OS uses the ascending order of the the FPC slot numbers as the sequence for powering on the FPCs.

NOTE: The FPC online sequence is not dependent on the FPC power-on sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

gre-in

Syntax

```
gre-in {
  mss value;
}
```

Hierarchy Level [edit security flow tcp-mss]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable and specify the TCP maximum segment size (TCP MSS) for Generic Routing Encapsulation (GRE) packets that are coming out from an IPsec VPN tunnel. If the device receives a GRE-encapsulated TCP packet with the SYN bit and TCP MSS option set and the TCP MSS option specified in the packet exceeds the TCP MSS specified by the device, the device modifies the TCP MSS value accordingly. By default, a TCP MSS for GRE packets is not set.

Options

mss *value* —TCP MSS for GRE packets. Value is optional.

Range: 64 through 63535 bytes

Default: 1320 bytes, if no value is specified

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

gre-out

Syntax

```
gre-out {
  mss value;
}
```

Hierarchy Level [edit security flow tcp-mss]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable and specify the TCP maximum segment size (TCP MSS) for Generic Routing Encapsulation (GRE) packets that are going into an IPsec VPN tunnel. If the device receives a GRE-encapsulated TCP packet with the SYN bit and TCP MSS option set and the TCP MSS option specified in the packet exceeds the TCP MSS specified by the device, the device modifies the TCP MSS value accordingly. By default, a TCP MSS for GRE packets is not set.

Options mss *value*—TCP MSS for GRE packets. Value is optional.

Range: 64 through 65,535 bytes

Default: 1320 bytes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

high-watermark

Syntax high-watermark *percent*;

Hierarchy Level [edit security flow aging]

Release Information Statement introduced in Junos OS Release 8.5.

Description Sets the point at which the aggressive aging-out process begins.

Options *percent* —Percentage of session-table capacity at which aggressive aging-out starts.

Range: 1 through 100 percent

Default: 100 percent

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

hop-by-hop-header

Syntax

```
hop-by-hop-header {
  CALIPSO-option;
  RPL-option;
  SFM-DPD-option;
  jumbo-payload-option;
  quick-start-option;
  router-alert-option;
  user-defined-option-type low | <to high>;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* ip ipv6-extension-header]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Define the IPv6 hop-by-hop screen option.

Options

CALIPSO-option—Enable the Common Architecture Label IPv6 Security Option.

RPL-option—Enable the Routing Protocol for Low-Power and Lossy Networks screen option.

SMF-DPD-option—Enable the Simplified Muliticast Forwarding IPv6 Duplicate Packet Detection screen option.

jumbo-payload-option—Enable the IPv6 jumbo payload screen option.

quick-start-option—Enable the IPv6 quick start screen option.

router-alert-option—Enable the IPv6 router alert screen option.

user-defined-header-type *low* | <to *high*>—Define the type of header range.

Range: 1 through 255.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding IPv6 Support for Screens on page 836
- ipv6-extension-header on page 1862
- destination-header on page 1845

Copyright © 2016, Juniper Networks, Inc.

icmpv6-malformed

Syntax icmpv6-malformed;

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Enable the ICMPv6 malformed intrusion detection service (IDS) option.

Options This statement has no options.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding IPv6 Support for Screens on page 836
- ipv6-extension-header on page 1862

inline-tap

Syntax inline-tap;

Hierarchy Level [edit security forwarding-process application-services maximize-idp-sessions]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable IDP inline tap mode. The inline tap feature provides passive, inline detection of application layer threats for traffic matching security policies which have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

interface-in (Security Forwarding Options)

Syntax interface-in *interface-name*;

Hierarchy Level [edit security forwarding-options mirror-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the incoming logical interface to be matched for mirroring.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

interface-out (Security Forwarding Options)

Syntax interface-out *interface-name*;

Hierarchy Level [edit security forwarding-options mirror-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the outgoing logical interface to be matched for mirroring.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

ipv4-template (Services)

Syntax ipv4-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify that the flow monitoring version 9 template is used only for IPv4 records.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

ipv6-extension-header

Syntax

```
ipv6-extension-header {
  AH-header;
  ESP-header
  HIP-header;
  destination-header {
    ILNP-nonce-option;
    home-address-option;
    line-identification-option;
    tunnel-encapsulation-limit-option;
    user-defined-option-type low | <to high>;
  }
  fragment-header;
  hop-by-hop-header {
    CALIPSO-option;
    RPL-option;
    SFM-DPD-option;
    jumbo-payload-option;
    quick-start-option;
    router-alert-option;
    user-defined-option-type low | <to high>;
  }
  mobility-header;
  no-next-header;
  routing-header;
  shim6-header
  user-defined-option-type low | <to high>;
}
```

Hierarchy Level [edit security screen ids-option *screen-name* ip]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Define the IPv6 extension header for the intrusion detection service (IDS).

Options

AH-header—Enable the IPv6 Authentication Header screen option.

ESP-header—Enable the IPv6 Encapsulating Security Payload header screen option.

HIP-header—Enable the IPv6 Host Identify Protocol header screen option.

fragment-header—Enable the IPv6 fragment header screen option.

mobility-header—Enable the IPv6 mobility header screen option.

no-next-header—Enable the IPv6 no next header screen option.

routing-header—Enable the IPv6 routing header screen option.

shim6-header—Enable the IPv6 shim header screen option.

user-defined-header-type *low* | <to *high*>—Define the type of header range.

Range: 0 through 255.

Copyright © 2016, Juniper Networks, Inc.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding IPv6 Support for Screens on page 836
- hop-by-hop-header on page 1859
- destination-header on page 1845

ipv6-extension-header-limit

Syntax
ipv6-extension-header-limit *limit*;

Hierarchy Level
[edit security screen ids-option *screen-name* ip]

Release Information
Statement introduced in Junos OS Release 12.1X46-D10.

Description
Define the IPv6 extension header number limit for screen options. The screen blocks packets that have more than the defined number of extension headers.

Options
limit—Set the number of IPv6 extension headers that can pass through the screen.
Range: 0 through 32.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding IPv6 Support for Screens on page 836
- ipv6-extension-header on page 1862

ipv6-malformed-header

Syntax
ipv6-malformed-header;

Hierarchy Level
[edit security screen ids-option *screen-name* ip]

Release Information
Statement introduced in Junos OS Release 12.1X46-D10.

Description
Enable the IPv6 malformed header intrusion detection service (IDS) option.

Options
This statement has no options.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding IPv6 Support for Screens on page 836
- ipv6-extension-header on page 1862

Copyright © 2016, Juniper Networks, Inc.

ipv6-template (Services)

Syntax ipv6-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify that the flow monitoring version 9 template is used only for IPv6 records.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

low-latency

Syntax low-latency

Hierarchy Level [edit chassis fpc *fpc-slot-number* pic *pic-slot-number* services-offload]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Enables the low-latency mode on the selected NP-IOC.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Example: Configuring Low Latency on page 1746

Copyright © 2016, Juniper Networks, Inc.

low-watermark

Syntax low-watermark *percent*;

Hierarchy Level [edit security flow aging]

Release Information Statement introduced in Junos OS Release 8.5.

Description Set the point at which the aggressive aging-out process ends.

Options *percent* —Percentage of session-table capacity at which aggressive aging-out ends.
Range: 0 through 100 percent
Default: 100 percent

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

maximize-idp-sessions

Syntax
```
maximize-idp-sessions {
  inline-tap;
  weight (equal | firewall | idp);
}
```

Hierarchy Level [edit security forwarding-process application-services]

Release Information Statement introduced in Junos OS Release 9.6.

Description If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity. See weight for information about the options provided.

NOTE: **The IDP session capacity is restricted to 100,000 sessions per SPU.**

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

mirror-filter (Security Forwarding Options)

Syntax

```
mirror-filter filter-name {
  destination-port port-number;
  destination-prefix destination-prefix;
  interface-in interface-name;
  interface-out interface-name;
  output {
    destination-mac mac-address;
    interface interface-name;
  }
  protocol protocol;
  source-port port-number;
  source-prefix source-prefix;
}
```

Hierarchy Level [edit security forwarding-options]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify match conditions for filtering packets for mirroring. You can specify the packet filter for the incoming and outgoing interfaces. A maximum of 15 filters are supported at the same time.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- show security forward-options mirror-filter on page 2167
- clear security forward-options mirror filter on page 1921

Copyright © 2016, Juniper Networks, Inc.

mode (Security Forwarding Options)

Syntax mode (drop | flow-based | packet-based);

Hierarchy Level [edit security forwarding-options family inet6]

Release Information Support on SRX Series devices for flow-based mode for family inet6 added in Junos OS Release 10.2.

Description Specify forwarding options for IPv6 traffic.

Options

- **drop**—Drop IPv6 packets. This is the default setting.
- **flow-based**—Perform flow-based packet forwarding.
- **packet-based**—Perform simple packet forwarding.

NOTE: **Packet-based processing is not supported on the following SRX Series devices: SRX5600, and SRX5800.**

If you change the forwarding option mode for IPv6, you might have to perform a reboot to initialize the configuration change. Table 181 summarizes device status upon configuration change.

Table 181: Device Status Upon Configuration Change

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

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

no-sequence-check

Syntax no-sequence-check;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that the device does not check sequence numbers in TCP segments during stateful inspection. By default, the device monitors the sequence numbers in TCP segments. The device detects the window scale specified by source and destination hosts in a session and adjusts a window for an acceptable range of sequence numbers according to their specified parameters. The device then monitors the sequence numbers in packets sent between these hosts. If the device detects a sequence number outside this range, it drops the packet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

no-tcp-reset

Syntax
```
no-tcp-reset {
  drop-all-tcp;
  drop-tcp-with-sync-only;
}
```

Hierarchy Level [edit system internet-options]

Release Information Statement introduced in Junos OS Release 11.1.

Description Do not send reset RST TCP packets for packets sent to non-listening ports.

Options
- **drop-all-tcp**—Drop all TCP packets.
- **drop-tcp-with-sync-only**— Drop TCP packets with a SYN bit.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Administration Guide for Security Devices*
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

output (Security Forwarding Options)

Syntax

```
output {
  destination-mac mac-address;
  interface interface-name;
}
```

Hierarchy Level [edit security forwarding-options mirror-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the MAC address or interface for mirrored traffic.

Options

destination-mac *mac-address*—Specify the MAC address for the mirrored traffic.

interface *interface-name*—Specify the logical interface for the mirrored traffic.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

Copyright © 2016, Juniper Networks, Inc.

packet-filter

Syntax

```
packet-filter packet-filter-name {
  action-profile (profile-name | default);
  destination-port (port-range | protocol-name);
  destination-prefix destination-prefix;
  interface logical-interface-name;
  protocol (protocol-number | protocol-name;
  source-port (port-range | protocol- name);
  source-prefix source-prefix;
  }
```

Hierarchy Level [edit security datapath-debug]

Release Information Command introduced in Junos OS Release 9.6 ; Support for IPv6 addresses for the **destination-prefix** and **source-prefix** options added in Junos OS Release 10.4. Support for IPv6 filter for the **interface** option added in Junos OS Release 10.4.

Description Set packet filter for taking the datapath-debug action. A maximum of four filters are supported at the same time.

Options

- **action-profile (*profile-name* | default)**—Identify the action profile to use. You can specify the name of the action profile to use or select default action profile.
- **destination-port (*port-range* | *protocol name*)**—Specify a destination port to match TCP/UDP destination port.
- **destination-prefix *destination-prefix***—Specify a destination IPv4/IPv6 address prefix.
- **interface *logical-interface-name***—Specify a logical interface name.
- **protocol (*protocol-number* | *protocol-name***—Match IP protocol type.
- **source-port (*port-range* | *protocol-name***—Match TCP/UDP source port.
- **source-prefix *source-prefix***—Specify a source IP address prefix.

Required Privilege Level security—To view this statement in the configuration
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

packet-ordering-mode (Application Services)

Syntax

```
packet-ordering-mode {
   (hardware | software);
}
```

Hierarchy Level [edit security forwarding-process application-services]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Enables or disables the packet-ordering functionality using the Packet Ordering Engine. By default, packet-ordering functionality using the Packet Ordering Engine (hardware) is enabled.

A system reboot is required when this feature is enabled or disabled, and a warning message is displayed during the commit.

NOTE: This feature is supported on SRX5800 and SRX5600 devices with next-generation SPCs.

Options

hardware— Enables packet-ordering functionality using the Packet Ordering Engine.

software— Disables packet-ordering functionality using the Packet Ordering Engine.

Required Privilege Level

- security—To view this statement in the configuration.
- security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

pending-sess-queue-length

Syntax pending-sess-queue-length (high | moderate | normal);

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the maximum queued length per pending session.

Options
- **high**— Allow the maximum number of queued sessions.
- **moderate**—Allow more queued sessions than the normal number.
- **normal**—Allow the normal number of queued session.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

propagate-settings

Syntax propagate-settings *interface-name*;

Hierarchy Level [edit system services dhcp]
[edit system services dhcp pool]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable or disable the propagation of TCP/IP settings received on the device acting as Dynamic Host Configuration Protocol (DHCP) client. The settings can be propagated to the server pool running on the device. Use the **system services dhcp** statement to set this feature globally. Use the **system services dhcp pool** statement to set the feature for the address pool and override the global setting.

Options ***logical-interface-name*** —Name of the logical interface to receive TCP/IP settings from the external network for propagation to the DHCP pool running on the device.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

protocol (Security Forwarding Options)

Syntax protocol *protocol*;

Hierarchy Level **[edit security forwarding-options mirror-filter *filter-name*]**

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the networking protocol name or number to be matched for mirroring.

Options ***protocol-name***—Networking protocol name or number. The following text values are supported. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

ah—IP Security Authentication header

egp—Exterior gateway protocol

esp—IPsec Encapsulating Security Payload

gre—Generic routing encapsulation

icmp—Internet Control Message Protocol

icmp6—Internet Control Message Protocol version 6

igmp—Internet Group Management Protocol

ipip—IP over IP

ospf—Open Shortest Path First

pim—Protocol Independent Multicast

rsvp—Resource Reservation Protocol

sctp—Stream Control Transmission Protocol

tcp—Transmission Control Protocol

udp—User Datagram Protocol

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

Copyright © 2016, Juniper Networks, Inc.

resource-manager

Syntax

```
resource-manager {
  traceoptions {
    flag flag;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before Junos OS Release 12.1.

Description Configure resource manager security options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

route-change-timeout

Syntax route-change-timeout *seconds*;

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5. Support for default value set to 6 seconds added in Junos OS Release 12.1X45-D10.

Description Specify the session timeout when a session is rerouted but there is a reroute failure (for example, the new route uses a different egress zone from the previous route).

Options ***seconds*** —Amount of time before sessions are timed out.

Range: 6 through 1800 seconds

Default: 6 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

rst-invalidate-session

Syntax rst-invalidate-session;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable the device to mark a session for immediate termination when it receives a TCP reset (RST) message. By default, this feature is disabled.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

rst-sequence-check

Syntax rst-sequence-check;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Verify that the TCP sequence number in a TCP segment with the RST bit enabled matches the previous sequence number for a packet in that session or is the next higher number incrementally. If the sequence number does not match either of these expected numbers, the device drops the packet and sends the host a TCP ACK message with the correct sequence number. By default, this check is disabled.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

sampling

Syntax

```
sampling {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Perform packet sampling based on particular input interfaces and various fields in the packet header.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the traffic sampling control process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

services-offload

Syntax

```
services-offload {
  low-latency;
  per-session-statistics;
}
```

Hierarchy Level [edit chassis fpc *fpc-slot-number* pic *pic-slot-number*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enables the Express Path mode (formerly known as *services offloading*) mode on the selected network processor.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

np-cache (Flexible PIC Concentrator)

Syntax np-cache;

Hierarchy Level [edit chassis fpc *fpc-slot-number*]

Release Information Statement introduced in Junos OS Release 15.1X49-D10.

Description Enable session cache table on IOC.

Starting with Junos OS Release 15.1X49-D10, NP cache is supported on the SRx5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), and SRX5K-MPC3-40G10G (IOC3) for SRX5400, SRX5600, and SRX5800 devices.

The security policy determines whether a session is for Express Path (formerly known as *services offloading*) mode on the selected Flexible PIC Concentrator (FPC).

NOTE: The IOC2 and the IOC3 utilize the delay sessions delete mechanism. The same sessions (sessions with the same five tuples) that are deleted and then reinstalled immediately are not cached on the IOCs.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Chassis Configuration Statement Hierarchy on page 1833
- Example: Configuring an SRX5K-MPC on an SRX5000 Line Device to Support Express Path on page 1737
- Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path on page 1739

Copyright © 2016, Juniper Networks, Inc.

source-port (Security Forwarding Options)

Syntax source-port *port-number*;

Hierarchy Level **[edit security forwarding-options mirror-filter *filter-name*]**

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port number to be matched for mirroring. You can specify a numeric value or one of the text synonyms listed in Table 182.

Copyright © 2016, Juniper Networks, Inc.

Table 182: Ports Supported by Services Interfaces

Port Name	Corresponding Port Number
afs	1483
bgp	179
biff	512
bootpc	68
bootps	67
cmd	514
cvspserver	2401
dhcp	67
domain	53
eklogin	2105
ekshell	2106
excc	512
finger	79
ftp	21
ftp-data	20
http	80
https	443
ident	113
imap	143
kerberos-sec	88
klogin	543
kpasswd	761
krb-prop	754
krbupdate	760

Copyright © 2016, Juniper Networks, Inc.

Table 182: Ports Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
kshell	544
ldap	389
ldp	646
login	513
mobileip-agent	434
mobilip-mn	435
msdp	639
netbios-dgm	138
netbios-ns	137
netbios-ssn	139
nfsd	2049
nntp	119
ntalk	518
ntp	123
pop3	110
pptp	1723
printer	515
radacct	1813
radius	1812
rip	520
rkinit	2108
smtp	25
snmp	161
snmp-trap	162

Copyright © 2016, Juniper Networks, Inc.

Table 182: Ports Supported by Services Interfaces *(continued)*

Port Name	Corresponding Port Number
snpp	444
socks	1080
ssh	22
sunrpc	111
syslog	514
tacacs	49
tacacs-ds	65
talk	517
telnet	23
tftp	69
timed	525
who	513
xdmcp	177

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

Copyright © 2016, Juniper Networks, Inc.

source-prefix (Security Forwarding Options)

Syntax source-prefix *source-prefix*;

Hierarchy Level [edit security forwarding-options mirror-filter *filter-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the source IP prefix or address to be matched for mirroring.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

syn-flood-protection-mode

Syntax syn-flood-protection-mode (syn-cookie | syn-proxy);

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5; support for IPv6 addresses added in Junos OS Release 10.4.

Description Enable SYN cookie or SYN proxy defenses against SYN attacks. SYN flood protection mode is enabled globally on the device and is activated when the configured **syn-flood attack-threshold value** is exceeded.

Options
- **syn-cookie**—Uses a cryptographic hash to generate a unique Initial Sequence Number (ISN). This is enabled by default.
- **syn-proxy**—Uses a proxy to handle the SYN attack.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

tcp-initial-timeout

Syntax tcp-initial-timeout *seconds*;

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the length of time (in seconds) that the device keeps an initial TCP session in the session table before dropping it, or until the device receives a FIN (no more data) or RST (reset) packet.

Options ***seconds***—Number of seconds that the device keeps an initial TCP session in the session table before dropping it.

Range: 4 through 300 seconds

Default: 20 seconds

NOTE: The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

tcp-mss (Security Flow)

Syntax

```
tcp-mss {
  all-tcp mss value;
  gre-in {
    mss value;
  }
  gre-out {
    mss value;
  }
  ipsec-vpn {
    mss value;
  }
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure TCP maximum segment size (TCP MSS) for the following packet types:

- All TCP packets for network traffic.
- GRE packets entering the IPsec VPN tunnel.
- GRE packets exiting the IPsec VPN tunnel.
- TCP packets entering the IPsec VPN tunnel.

If all the four TCP MSS options are configured simultaneously, then the order of preference is as follows:

- If TCP packet enters an IPsec VPN tunnel, then an ipsec-vpn mss value has high priority over all-tcp mss value, hence ipsec-vpn mss value is set.
- If TCP packet enters GRE , then gre-in mss value overrides all-tcp mss value, hence gre-in mss value is set.
- If TCP packet exits GRE, then all-tcp mss value overrides gre-in mss value, hence all-tcp mss value is set.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- all-tcp on page 1840
- gre-in on page 1857
- gre-out on page 1858
- ipsec-vpn (Security Flow) on page 6900
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

tcp-session

Syntax

```
tcp-session {
  no-sequence-check;
  no-syn-check;
  no-syn-check-in-tunnel;
  rst-invalidate-session;
  rst-sequence-check;
  strict-syn-check;
  tcp-initial-timeout seconds;
  time-wait-state {
    (session-ageout | session-timeout seconds);
  }
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure TCP session attributes:

- TCP sequence number checking.
- TCP SYN bit checking.
- Reset (RST) checking.
- Initial TCP session timeout—The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.
- Strict TCP SYN checking.
- TCP session timeout for time-wait state.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

time-wait-state

Syntax

```
time-wait-state {
    (session-ageout | session-timeout seconds);
}
```

Hierarchy Level [edit security flow tcp-session]

Release Information Statement introduced in Junos OS Release 11.1.

Description Defines the length of time (in seconds) that the device keeps the defined TCP session in the session table. The default is 150 seconds.

Options

- **session-ageout**—Set a TCP session to age out, using the service based timeout value.
- **session-timeout *seconds***—Set the session timeout value allowed before the device ages out a session from its session table.

  Range: 2 through 600 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
  rate-limit messages-per-second;
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure security tracing options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace all security events
  - **compilation**—Trace security compilation events
  - **configuration**—Trace security configuration events
  - **routing-socket**—Trace routing socket events
- **no-remote-trace**—Set remote tracing as disabled.
- **rate-limit *messages-per-second***—Limit the incoming rate of trace messages.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Flow)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
  packet-filter filter-name {
    destination-port port-identifier;
    destination-prefix address;
    interface interface-name;
    protocol protocol-identifier;
    source-port port-identifier;
    source-prefix address;
  }
  rate-limit messages-per-second;
  trace-level (brief | detail | error);
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 12.1X46-D10 with the **trace-level** option and additional flags.

Description Configure flow tracing options.

Options **file**—Configure the trace file options.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

files ***number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10 files

match ***regular-expression***—Refine the output to include lines that contain the regular expression.

size ***maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size,

Copyright © 2016, Juniper Networks, Inc.

trace-file.0 is renamed *trace-file*.1 and *trace-file* is renamed *trace-file*.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: *x* **K** to specify KB, *x* **m** to specify MB, or *x* **g** to specify GB

Range: 0 KB through 1 GB

Default: 128 KB

Copyright © 2016, Juniper Networks, Inc.

world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

all—Trace with all flags enabled

basic-datapath—Trace basic packet flow activity

fragmentation—Trace IP fragmentation and reassembly events

high-availability—Trace flow high-availability information

host-traffic—Trace flow host traffic information

multicast—Trace multicast flow information

route—Trace route lookup information

session—Trace session creation and deletion events

session-scan—Trace session scan information

tcp-basic—Trace TCP packet flow information

tunnel—Trace tunnel information

no-remote-trace—Set remote tracing as disabled.

packet-filter ***filter-name***—Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port ***port-identifier***—Match TCP/UDP destination port

destination-prefix ***address***—Destination IP address prefix

interface ***interface-name***—Logical interface

protocol ***protocol-identifier***—Match IP protocol type

source-port ***port-identifier***—Match TCP/UDP source port

source-prefix ***address***—Source IP address prefix

rate-limit ***messages-per-second***—Limit the incoming rate of trace messages.

Copyright © 2016, Juniper Networks, Inc.

trace-level—Set the level for trace logging. This option is available only when the flag is set.

brief—Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.

detail—Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.

error—Trace error information, such as system failure, unknown message type, and packet drop.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

transport (Security Log)

Syntax

```
transport {
  protocol (udp | tcp | tls);
  tls-profile tls-profile-name;
  tcp-connections tcp-connections;
}
```

Hierarchy Level
[edit security log]

Release Information
Statement introduced in Junos OS Release 12.1X46-D25.

Description
Configure security log transport options.

Options
protocol—Specify the type of transport protocol to be used to log the data.

- **UDP**—Set the transport protocol to UDP.
- **TCP**—Set the transport protocol to TCP.
- **TLS**—Set the transport protocol to TLS.

Default: UDP.

tls-profile *tls-profile-name*—Specify the TLS profile name.

tcp-connections *tcp-connections*—Specify the number of TCP connections per SPU.

Range: 1 through 5.

Default: 1.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding AppTrack on page 573

Copyright © 2016, Juniper Networks, Inc.

weight (Security)

Syntax weight (equal | firewall | idp);

Hierarchy Level [edit security forwarding-process application-services maximize-idp-sessions]

Release Information Statement introduced in Junos OS Release 9.6.

Description If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity.

Devices ship with an implicit default session capacity setting. This default value gives more weight to firewall sessions. You can manually override the default by using the **maximize-idp-sessions** command. The command allows you to choose between these weight values: **equal**, **firewall**, and **idp**. The following table displays the available session capacity weight and approximate throughput for each.

Table 183: Session Capacity and Resulting Throughput

Weight Value	Firewall Capacity	IDP Capacity	Firewall Throughput	IDP Throughput
Default	1,000,000	256,000	10 Gbps	2.4 Gbps
equal	1,000,000	1,000,000	8.5 Gbps	2 Gbps
firewall	1,000,000	1,000,000	10 Gbps	2.4 Gbps
idp	1,000,000	1,000,000	5.5 Gbps	1.4 Gbps

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 94

Operational Commands

- clear firewall
- clear monitor security flow filter
- clear security flow ip-action
- clear security flow session all
- clear security flow session application
- clear security flow session destination-port
- clear security flow session family
- clear security flow session interface
- clear security flow session protocol
- clear security flow session resource-manager
- clear security flow session services-offload
- clear security flow session session-identifier
- clear security flow session source-port
- clear security flow session source-prefix
- clear security forward-options mirror filter
- monitor security flow file
- monitor security flow filter
- monitor security flow start
- monitor security flow stop
- show chassis environment (Security)
- show chassis fpc (View)
- show chassis fpc errors
- show chassis hardware (View)
- show chassis pic (Security)
- show chassis power
- show chassis power sequence
- show firewall (View)
- show interfaces (View Aggregated Ethernet)

Copyright © 2016, Juniper Networks, Inc.

- show interfaces (SRX Series)
- show interfaces diagnostics optics
- show interfaces flow-statistics
- show interfaces statistics (View)
- show interfaces swfabx
- show monitor security flow
- show security flow cp-session
- show security flow cp-session destination-port
- show security flow cp-session destination-prefix
- show security flow cp-session family
- show security flow cp-session protocol
- show security flow cp-session source-port
- show security flow cp-session source-prefix
- show security flow gate
- show security flow ip-action
- show security flow gate brief node
- show security flow gate destination-port
- show security flow gate destination-prefix
- show security flow gate protocol
- show security flow gate summary node
- show security flow session
- show security flow session brief node
- show security flow session destination-port
- show security flow session destination-prefix
- show security flow session extensive node
- show security flow session family
- show security flow session interface
- show security flow session nat
- show security flow session policy-id
- show security flow session protocol
- show security flow session resource-manager
- show security flow session services-offload
- show security flow session session-identifier
- show security flow session source-port
- show security flow session source-prefix
- show security flow session summary family
- show security flow session summary node

Copyright © 2016, Juniper Networks, Inc.

- show security flow session summary services-offload
- show security flow session tunnel
- show security flow statistics
- show security flow status
- show security forward-options mirror-filter
- show security monitoring
- show security policies
- show security policies hit-count
- show security resource-manager group active
- show security resource-manager resource active
- show security resource-manager settings
- show security resource-manager summary
- show security screen ids-option
- show security screen statistics
- show security softwires
- show security zones
- show security zones type

Copyright © 2016, Juniper Networks, Inc.

clear firewall

Syntax

clear firewall
<all>
<counter *counter-name*>
<filter *filter-name*>

Release Information Command introduced in Junos OS Release 10.0.

Description Clear statistics about configured firewall filters.

NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for GRES.

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

Options

all—Clear the packet and byte counts for all filters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

Required Privilege Level clear

Related Documentation

- show firewall (View) on page 1974

List of Sample Output clear firewall all on page 1898

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear firewall all

```
user@host> clear firewall all
```


Copyright © 2016, Juniper Networks, Inc.

clear monitor security flow filter

Syntax clear monitor security flow filter <*filter-name*>

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Specify the security flow filters to be deleted. Once deleted, the filters are removed from the Packet Forwarding Engine and the Routing Engine.

NOTE: Specifying the filter name is optional. If no filter is specified, all filters are deleted.

Options This command has no options.

Required Privilege Level clear

Related Documentation

- Monitoring Security Flow Sessions Overview on page 1781
- monitor security flow start on page 1926
- monitor security flow filter on page 1924

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

Copyright © 2016, Juniper Networks, Inc.

clear security flow ip-action

Syntax **clear security flow ip-action** [*filter*]

Release Information Command introduced in Junos OS Release 10.4. Logical systems option introduced in Junos OS Release 11.2.

Description Clear IP-action entries, based on filtered options, for IP sessions running on the device.

Options *filter*—Filter the display based on the specified criteria.

The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

all | [*filter*]—All active sessions on the device.

destination-port ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535.

destination-prefix ***destination-prefix***—Destination IP prefix or address.

family (inet | inet6) [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

logical-system ***logical-system-name*** **| all** [*filter*]—Specified logical system or all logical systems.

protocol ***protocol-name*** **|** ***protocol-number*** [*filter*]—Protocol name or number and filtered options.

- **ah** or **51**
- **egp** or **8**
- **esp** or **50**
- **gre** or **47**
- **icmp** or **1**
- **icmp6** or **58**
- **igmp** or **2**
- **ipip** or **4**
- **ospf** or **89**
- **pim** or **103**
- **rsvp** or **46**
- **sctp** or **132**
- **tcp** or **6**
- **udp** or **17**

Copyright © 2016, Juniper Networks, Inc.

root-logical-system [*filter*]—Default logical system information and filtered options.

source-port ***source-port***—Source port number of the traffic. Range is 1 through 65,535.

source-prefix ***source-prefix***—Source IP prefix or address of the traffic.

Required Privilege Level clear

Related Documentation
- show security flow ip-action on page 2051

List of Sample Output
clear security flow ip-action all on page 1901
clear security flow ip-action destination-prefix on page 1901
clear security flow ip-action family inet on page 1901
clear security flow ip-action protocol udp on page 1901

Output Fields When you enter this command, the system responds with the status of your request.

Sample Output

clear security flow ip-action all

```
user@host>clear security flow ip-action all
1008 ip-action entries cleared
```

clear security flow ip-action destination-prefix

```
user@host>clear security flow ip-action destination-prefix 192.0.2.5/24
87 ip-action entries cleared
```

clear security flow ip-action family inet

```
user@host>clear security flow ip-action family inet
2479 ip-action entries cleared
```

clear security flow ip-action protocol udp

```
user@host>clear security flow ip-action protocol udp
270 ip-action entries cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session all

Syntax clear security flow session all
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5 ; **node** options added in Junos OS Release 9.0.

Description Clear all currently active security sessions on the device.

Options

- all—Clear information about all active sessions.
- **node**—(Optional) For chassis cluster configurations, clear all security sessions on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- show security flow session on page 453

List of Sample Output

clear security flow session all on page 1902
clear security flow session all node 0 on page 1902

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session all

```
user@host> clear security flow session all
node0:
--------------------------------------------------------------------------
1 active sessions cleared
node1:
-------- ------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session all node 0

```
user@host> clear security flow session all node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

 Copyright © 2016, Juniper Networks, Inc.

clear security flow session application

Syntax clear security flow session application
application-name
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear currently active sessions for application types or application sets.

Options

- *application-name* —Name of the specified application type or application set.
  - **dns**—Domain Name System
  - **ftp**—File Transfer Protocol
  - **ignore**—Ignore application type
  - **mgcp-ca**—Media Gateway Control Protocol with Call Agent
  - **mgcp-ua**—MGCP with User Agent
  - **ms-rpc**—Microsoft RPC
  - **pptp**—Point-to-Point Tunneling Protocol
  - **q931**—ISDN connection control protocol
  - **ras**—RAS
  - **realaudio**—RealAudio
  - **rsh**—UNIX remote shell services
  - **rtsp**—Real-Time Streaming Protocol
  - **sccp**—Skinny Client Control Protocol
  - **sip**—Session Initiation Protocol
  - **sqlnet-v2**—Oracle SQLNET
  - **sun-rpc**—Sun Microsystems RPC
  - **talk**—TALK program
  - **tftp**—Trivial File Transfer Protocol
- **node**—(Optional) For chassis cluster configurations, clear sessions for applications on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.

Copyright © 2016, Juniper Networks, Inc.

- **local** —Clear the local node.
- **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation
- show security flow session application on page 459

List of Sample Output
clear security flow session application dns on page 1904
clear security flow session application dns node 0 on page 1904

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session application dns

```
user@host> clear security flow session application dns
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session application dns node 0

```
user@host> clear security flow session application dns node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security flow session destination-port

Syntax clear security flow session destination-port
destination-port-number
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5 ; **node** options added in Junos OS Release 9.0.

Description Clear each session that uses the specified destination port

Options

- ***destination-port-number*** —Number of the destination port.
- **node**—(Optional) For chassis cluster configurations, clear security sessions on the port on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- show security flow session destination-port on page 2089

List of Sample Output

clear security flow session destination-port 1 on page 1905
clear security flow session destination-port 1 node 0 on page 1905

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session destination-port 1

```
user@host> clear security flow session destination-port 1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session destination-port 1 node 0

```
user@host> clear security flow session destination-port 1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session family

Syntax clear security flow session family (inet | inet6)

Release Information Command introduced in Junos OS Release 10.2.

Description Clear sessions that match the specified protocol family.

Options
- **inet**—Clear IPv4 sessions.
- **inet6**—Clear IPv6 sessions.

Required Privilege Level clear

Related Documentation
- show security flow session family on page 2103

List of Sample Output
clear security flow session family inet on page 1907
clear security flow session family inet6 on page 1907

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session family inet

```
user@host> clear security flow session family inet
1 active sessions cleared
```

clear security flow session family inet6

```
user@host> clear security flow session family inet6
1 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session interface

Syntax

clear security flow session interface
interface-name
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Clear sessions that use the specified interface.

Options

- *interface-name* —Name of a specific incoming or outgoing interface.
- **node**—(Optional) For chassis cluster configurations, clear security sessions on the interface on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level

clear

Related Documentation

- show security flow session interface on page 2108

List of Sample Output

clear security flow session interface ge-0/0/0.0 on page 1908
clear security flow session interface ge/0/0.0 node 0 on page 1908

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session interface ge-0/0/0.0

```
user@host> clear security flow session interface ge-0/0/0.0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session interface ge/0/0.0 node 0

```
user@host> clear security flow session interface ge-0/0/0.0 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

clear security flow session protocol

Syntax clear security flow session protocol *protocol-name* | *protocol-number*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5 ; **node** options added in Junos OS Release 9.0.

Description Clear each session that uses the specified IP protocol.

Options

- ***protocol-name*** — (Optional) Networking protocol name. The following text values are supported.
  - **ah**—IP Security Authentication Header
  - **egp**—Exterior gateway protocol
  - **esp**—IPsec Encapsulating Security Payload
  - **gre**—Generic routing encapsulation
  - **icmp**—Internet Control Message Protocol
  - **igmp**—Internet Group Management Protocol
  - **ipip**—IP over IP
  - **ospf**—Open Shortest Path First
  - **pim**—Protocol Independent Multicast
  - **rsvp**—Resource Reservation Protocol
  - **sctp**—Stream Control Transmission Protocol
  - **tcp**—Transmission Control Protocol
  - **udp**—User Datagram Protocol
- ***protocol-number*** —(Optional) Numeric protocol value. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

  Range: 0 through 255
- **node**—(Optional) For chassis cluster configurations, clear security on a specific node (device) in the cluster for the user with this identification number.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- show security flow session protocol on page 2118

List of Sample Output
clear security flow session protocol pim on page 1911
clear security flow session protocol 0 on page 1911

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session protocol pim

```
user@host> clear security flow session protocol pim
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session protocol 0

```
user@host> clear security flow session protocol 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session resource-manager

Syntax clear security flow session resource-manager
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Clear resource-manager sessions.

Options

- **node**—(Optional) For chassis cluster configurations, clear the resource manager sessions on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- show security flow session resource-manager on page 463

List of Sample Output

clear security flow session resource-manager on page 1912
clear security flow session resource-manager node 0 on page 1912

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session resource-manager

```
user@host> clear security flow session resource-manager
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session resource-manager node 0

```
user@host> clear security flow session resource-manager node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security flow session services-offload

Syntax **clear security flow session services-offload** [*filter*]

Release Information Command introduced in Junos OS Release 11.4.
Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 devices.

Description Clear services-offload security sessions, based on filtered options, on the device. This command also clears a services-offload security session from both the network processor and the Services Processing Unit (SPU) on which the specified session was installed.

Options *filter*—Filter the display based on the specified criteria.

The following filters clear those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

application *application-name*—Application name.

destination-port *destination-port*—Destination port number. Range is from 1 through 65,535.

destination-prefix *destination-prefix*—Destination IP prefix or address.

family (inet | inet6)—IPv4 traffic or IPv6-NAT-PT traffic.

interface *interface-name*—Incoming or outgoing interface name.

logical-system *logical-system-name* | all—Specified logical system name or all logical systems.

protocol *protocol-name* | *protocol-number* —Protocol name or number.

- **ah** or **51**
- **egp** or **8**
- **esp** or **50**
- **gre** or **47**
- **icmp** or **1**
- **icmp6** or **58**
- **igmp** or **2**
- **ipip** or **4**
- **ospf** or **89**
- **pim** or **103**
- **rsvp** or **46**

Copyright © 2016, Juniper Networks, Inc.

- **sctp** or **132**
- **tcp** or **6**
- **udp** or **17**

root-logical-system [*filter*]—Root logical system information and filtered options.

source-port ***source-port***—Source port number of the traffic. Range is from 1 through 65,535.

source-prefix ***source-prefix***—Source IP prefix or address of the traffic.

Required Privilege Level clear

Related Documentation
- show security flow session services-offload on page 2127

List of Sample Output
clear security flow session services-offload on page 1914
clear security flow session services-offload application on page 1914
clear security flow session services-offload destination-port on page 1914
clear security flow session services-offload destination-prefix on page 1914
clear security flow session services-offload family on page 1914
clear security flow session services-offload interface on page 1915
clear security flow session services-offload logical-system on page 1915
clear security flow session services-offload protocol on page 1915
clear security flow session services-offload root-logical-system on page 1915
clear security flow session services-offload source-port on page 1915
clear security flow session services-offload source-prefix on page 1915

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session services-offload

```
user@host>clear security flow session services-offload
0 active sessions cleared
```

clear security flow session services-offload application

```
user@host>clear security flow session services-offload dns
0 active sessions cleared
```

clear security flow session services-offload destination-port

```
user@host>clear security flow session services-offload destination-port 1
0 active sessions cleared
```

clear security flow session services-offload destination-prefix

```
user@host>clear security flow session services-offload destination-prefix 100.0.0.1
0 active sessions cleared
```

clear security flow session services-offload family

```
user@host>clear security flow session services-offload family inet
```


Copyright © 2016, Juniper Networks, Inc.

```
1 active sessions cleared
```

clear security flow session services-offload interface

```
user@host>clear security flow session services-offload interface ge-0/0/0.0
0 active sessions cleared
```

clear security flow session services-offload logical-system

```
user@host>clear security flow session services-offload logical-system all
0 active sessions cleared
```

clear security flow session services-offload protocol

```
user@host>clear security flow session services-offload protocol pim
0 active sessions cleared
```

clear security flow session services-offload root-logical-system

```
user@host>clear security flow session services-offload root-logical-system application dns
0 active sessions cleared
```

clear security flow session services-offload source-port

```
user@host>clear security flow session services-offload source-port 1
0 active sessions cleared
```

clear security flow session services-offload source-prefix

```
user@host>clear security flow session services-offload source-prefix 10.0.0.1
0 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session session-identifier

Syntax

clear security flow session session-identifier
session-identifier
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Clear the session with the specific identifier.

Options

- ***session-identifier*** —Number from 1 through 4,294,967,295 that identifies the security session.
- **node**—(Optional) For chassis cluster configurations, clear the specified session on a specific node (device) in the cluster.
  - ***node-id***—Identification number of the node. It can be 0 or 1.
  - **all**—Clear all nodes.
  - **local**—Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level

clear

Related Documentation

- show security flow session session-identifier on page 2132

List of Sample Output

clear security flow session session-identifier 1 on page 1916
clear security flow session session-identifier 1 node 0 on page 1916

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session session-identifier 1

```
user@host> clear security flow session session-identifier 1
0 active sessions cleared
```

Sample Output

clear security flow session session-identifier 1 node 0

```
user@host> clear security flow session session-identifier 1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security flow session source-port

Syntax

clear security flow session source-port
source-port-number
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Clear each session that uses the specified source port.

Options

- *source-port-number* —Number that identifies the source port.

  Range: 1 through 65,535

- **node**—(Optional) For chassis cluster configurations, clear sessions on the specified source port on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level

clear

Related Documentation

- show security flow session source-port on page 2136

List of Sample Output

clear security flow session source-port 1 on page 1917
clear security flow session source-port 1 node 0 on page 1917

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session source-port 1

```
user@host> clear security flow session source-port 1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

Sample Output

clear security flow session source-port 1 node 0

```
user@host> clear security flow session source-port 1 node 0
```

Copyright © 2016, Juniper Networks, Inc.

```
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security flow session source-prefix

Syntax
clear security flow session source-prefix
source-prefix-number
<node (*node-id* | all | local | primary)>

Release Information
Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0. Support for IPv6 addresses added in Junos OS Release 10.2.

Description
Clear sessions that match the source prefix.

Options
- *source-prefix-number*—Source IP prefix or address.
- **node**—(Optional) For chassis cluster configurations, clear security sessions matching the source prefix on a specific node (device) in the cluster.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - **all**—Clear all nodes.
  - **local**—Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level
clear

Related Documentation
- show security flow session source-prefix on page 2140

List of Sample Output
clear security flow session source-prefix 10.0.0.1 on page 1919
clear security flow session source-prefix 10::10 on page 1919
clear security flow session source-prefix 10.0.0.1 node 0 on page 1920

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session source-prefix 10.0.0.1

```
user@host> clear security flow session source-prefix 10.0.0.1
node0:
--------------------------------------------------------------------------
0 active sessions cleared
node1:
--------------------------------------------------------------------------
0 active sessions cleared
```

clear security flow session source-prefix 10::10

```
user@host> clear security flow session source-prefix 10::10
1 active sessions cleared
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

clear security flow session source-prefix 10.0.0.1 node 0

```
user@host> clear security flow session source-prefix 10.0.0.1 node 0
node0:
--------------------------------------------------------------------------
0 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security forward-options mirror filter

Syntax clear security forward-options mirror-filter (all | *filter-name*)

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Clear statistics about configured mirror filters.

Options **all**—Clear statistics for all configured mirror filters.

filter-name—Clear statistics for the specified mirror filter.

Required Privilege Level clear

Related Documentation

- mirror-filter (Security Forwarding Options) on page 1866
- show security forward-options mirror-filter on page 2167

Copyright © 2016, Juniper Networks, Inc.

monitor security flow file

Syntax

monitor security flow file
<*file-name*>
<files *number*>
<match *regular-expression*>
<size *maximum-file-size*>
<(world-readable | no-world-readable)>

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

Description

Configure options for the security flow monitoring output.

Options

filename—Name of the file to receive the output of the monitoring operation. All output is saved in the **/var/log/** directory.

files *number*—Maximum number of output files. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

match *regular-expression*—Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—Maximum size of each output file. When an output file named ***output*** reaches this size, it is renamed ***output*.0**. When the output file again reaches its maximum size, ***output*.0** is renamed ***output*.1** and **output** file is renamed ***output*.0**. This renaming scheme continues until the maximum number of output files is reached. Then the oldest output file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of output files with the **files** option.

Range: 10 KB through 1 GB

Default: 128 KB

(world-readable | no-world-readable)—By default, the output files can be accessed only by the user who configures the monitoring operation. The **world-readable** option enables all users to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Required Privilege Level

trace

Related Documentation

- Monitoring Security Flow Sessions Overview on page 1781
- monitor security flow filter on page 1924
- monitor security flow start on page 1926
- show monitor security flow on page 2029

Copyright © 2016, Juniper Networks, Inc.

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

monitor security flow filter

Syntax

monitor security flow filter *filter-name*
<destination-port (*port-range* | *protocol-name*)>
<destination-prefix *destination-prefix*>
<interface *interface-name*>
<logical-system *logical-system-name*>
<protocol (*protocol name* | *protocol number*)>
<root-logical-system>
<source-port (*port-range* | *protocol-name*)>
<source-prefix *source-prefix*>

Release Information

Command introduced in Junos OS Release 12.1X46-D10.

Description

Set security flow filters to define flow sessions that you want to monitor. A maximum of 64 filters is supported at a time.

Defining the filters themselves does not trigger monitoring. You have to explicitly use the **monitor security flow start** command to enable monitoring. Once monitoring starts, any traffic that matches the specified filters is saved in an output file in the **/var/log/** directory.

NOTE: Unlike filters defined in the configuration mode, these filters defined using operational mode commands are cleared when you reboot your system.

Options

filter *filter-name*—Specify a name for the filter. The filter name can contain letters, numbers, underscores (_) and hyphens (-) and can be up to 64 characters long.

destination-port (*port-range* | *protocol-name*)—Specify the TCP or UDP destination port to match. You can also specify a range of TCP or UDP destination ports and monitor all traffic in this group.

destination-prefix *destination-prefix*—Specify the destination IPv4 or IPv6 address prefix to match.

interface *interface-name*—Specify the logical interface name to match.

logical-system *logical-system-name*—Specify the logical system name to match.

protocol (*protocol name* | *protocol number*)—Specify the IP protocol type to match.

root-logical-system—(Default) Specify the root logical system to match.

source-port (*port-range* | *protocol-name*)—Specify the TCP or UDP source port to match. You can also specify a range of TCP or UDP source ports and monitor all traffic in this group.

source-prefix *source-prefix*—Specify the source IP address prefix to match.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation

- Monitoring Security Flow Sessions Overview on page 1781
- monitor security flow file on page 1922
- monitor security flow start on page 1926
- monitor security flow stop on page 1927

Copyright © 2016, Juniper Networks, Inc.

monitor security flow start

Syntax monitor security flow start

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Start the monitoring of security flow session. Once monitoring starts, any traffic that matches the specified filters is saved in an output file in the **var/log/** directory. At least one filter must be defined for the monitoring to start.

Use the **monitor security flow stop** command to stop the monitoring of flow sessions.

Options This command has no options.

Required Privilege Level trace

Related Documentation

- Monitoring Security Flow Sessions Overview on page 1781
- show monitor security flow on page 2029
- monitor security flow filter on page 1924
- monitor security flow stop on page 1927

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

monitor security flow stop

Syntax monitor security flow stop

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Stop monitoring the security flow session. Use the **monitor security flow start** command to start the monitoring of flow sessions.

Options This command has no options.

Required Privilege Level trace

Related Documentation

- Monitoring Security Flow Sessions Overview on page 1781
- monitor security flow start on page 1926

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

show chassis environment (Security)

Syntax show chassis environment

Release Information Command introduced in Junos OS Release 9.2.

Description Display environmental information about the services gateway chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

Options

none—Display environmental information about the device.

cb ***slot-number***—Display chassis environmental information for the Control Board.

fpc ***fpc-slot***—Display chassis environmental information for a specified Flexible PIC Concentrator.

fpm—Display chassis environmental information for the craft interface (FPM).

pem ***slot-number***—Display chassis environmental information for the specified Power Entry Module.

routing-engine ***slot-number***—Display chassis environmental information for the specified Routing Engine.

Required Privilege Level view

Related Documentation

- show chassis hardware (View) on page 1942

List of Sample Output show chassis environment on page 1928

Output Fields Table 184 lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 184: show chassis environment Output Fields

Field Name	Field Description
Temp	Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F).
Fan	Fan status: **OK**, **Testing** (during initial power-on), **Failed**, or **Absent**.

Sample Output

show chassis environment

```
user@host> show chassis environment
user@host> show chassis environment
Class Item                          Status     Measurement
Temp  PEM 0                         OK         40 degrees C / 104 degrees F
      PEM 1                         OK         40 degrees C / 104 degrees F
      PEM 2                         OK         40 degrees C / 104 degrees F
```


Copyright © 2016, Juniper Networks, Inc.

```
PEM 3                          OK          45 degrees C / 113 degrees F
Routing Engine 0               OK          31 degrees C / 87 degrees F
Routing Engine 0 CPU           OK          27 degrees C / 80 degrees F
Routing Engine 1               Absent
Routing Engine 1 CPU           Absent
CB 0 Intake                    OK          28 degrees C / 82 degrees F
CB 0 Exhaust A                 OK          27 degrees C / 80 degrees F
CB 0 Exhaust B                 OK          29 degrees C / 84 degrees F
CB 0 ACBC                      OK          29 degrees C / 84 degrees F
CB 0 SF A                      OK          36 degrees C / 96 degrees F
CB 0 SF B                      OK          31 degrees C / 87 degrees F
CB 1 Intake                    OK          27 degrees C / 80 degrees F
CB 1 Exhaust A                 OK          26 degrees C / 78 degrees F
CB 1 Exhaust B                 OK          29 degrees C / 84 degrees F
CB 1 ACBC                      OK          27 degrees C / 80 degrees F
CB 1 SF A                      OK          36 degrees C / 96 degrees F
CB 1 SF B                      OK          31 degrees C / 87 degrees F
CB 2 Intake                    Absent
CB 2 Exhaust A                 Absent
CB 2 Exhaust B                 Absent
CB 2 ACBC                      Absent
CB 2 XF A                      Absent
CB 2 XF B                      Absent
FPC 0 Intake                   OK          47 degrees C / 116 degrees F
FPC 0 Exhaust A                OK          44 degrees C / 111 degrees F
FPC 0 Exhaust B                OK          52 degrees C / 125 degrees F
FPC 0 xlp0 TSen                OK          51 degrees C / 123 degrees F
FPC 0 xlp0 Chip                OK          46 degrees C / 114 degrees F
FPC 0 xlp1 TSen                OK          51 degrees C / 123 degrees F
FPC 0 xlp1 Chip                OK          47 degrees C / 116 degrees F
FPC 0 xlp2 TSen                OK          44 degrees C / 111 degrees F
FPC 0 xlp2 Chip                OK          42 degrees C / 107 degrees F
FPC 0 xlp3 TSen                OK          48 degrees C / 118 degrees F
FPC 0 xlp3 Chip                OK          43 degrees C / 109 degrees F
FPC 1 Intake                   OK          41 degrees C / 105 degrees F
FPC 1 Exhaust A                OK          41 degrees C / 105 degrees F
FPC 1 Exhaust B                OK          51 degrees C / 123 degrees F
FPC 1 LU TSen                  OK          46 degrees C / 114 degrees F
FPC 1 LU Chip                  OK          45 degrees C / 113 degrees F
FPC 1 XM TSen                  OK          46 degrees C / 114 degrees F
FPC 1 XM Chip                  OK          52 degrees C / 125 degrees F
FPC 1 xlp0 TSen                OK          49 degrees C / 120 degrees F
FPC 1 xlp0 Chip                OK          42 degrees C / 107 degrees F
FPC 1 xlp1 TSen                OK          49 degrees C / 120 degrees F
FPC 1 xlp1 Chip                OK          44 degrees C / 111 degrees F
FPC 1 xlp2 TSen                OK          38 degrees C / 100 degrees F
FPC 1 xlp2 Chip                OK          39 degrees C / 102 degrees F
FPC 1 xlp3 TSen                OK          44 degrees C / 111 degrees F
FPC 1 xlp3 Chip                OK          42 degrees C / 107 degrees F
FPC 2 Intake                   OK          29 degrees C / 84 degrees F
FPC 2 Exhaust A                OK          34 degrees C / 93 degrees F
FPC 2 Exhaust B                OK          40 degrees C / 104 degrees F
FPC 2 I3 0 TSensor             OK          42 degrees C / 107 degrees F
FPC 2 I3 0 Chip                OK          41 degrees C / 105 degrees F
FPC 2 I3 1 TSensor             OK          40 degrees C / 104 degrees F
FPC 2 I3 1 Chip                OK          39 degrees C / 102 degrees F
FPC 2 I3 2 TSensor             OK          38 degrees C / 100 degrees F
FPC 2 I3 2 Chip                OK          37 degrees C / 98 degrees F
FPC 2 I3 3 TSensor             OK          35 degrees C / 95 degrees F
FPC 2 I3 3 Chip                OK          35 degrees C / 95 degrees F
FPC 2 IA 0 TSensor             OK          45 degrees C / 113 degrees F
```

Copyright © 2016, Juniper Networks, Inc.

```
                FPC 2 IA 0 Chip             OK          42 degrees C / 107 degrees F
                FPC 2 IA 1 TSensor          OK          41 degrees C / 105 degrees F
                FPC 2 IA 1 Chip             OK          43 degrees C / 109 degrees F
                FPC 9 Intake                OK          29 degrees C / 84 degrees F
                FPC 9 Exhaust A             OK          41 degrees C / 105 degrees F
                FPC 9 Exhaust B             OK          48 degrees C / 118 degrees F
                FPC 9 LU TSen               OK          48 degrees C / 118 degrees F
                FPC 9 LU Chip               OK          47 degrees C / 116 degrees F
                FPC 9 XM TSen               OK          48 degrees C / 118 degrees F
                FPC 9 XM Chip               OK          54 degrees C / 129 degrees F
                FPC 9 xlp0 TSen             OK          45 degrees C / 113 degrees F
                FPC 9 xlp0 Chip             OK          42 degrees C / 107 degrees F
                FPC 9 xlp1 TSen             OK          49 degrees C / 120 degrees F
                FPC 9 xlp1 Chip             OK          46 degrees C / 114 degrees F
                FPC 9 xlp2 TSen             OK          37 degrees C / 98 degrees F
                FPC 9 xlp2 Chip             OK          40 degrees C / 104 degrees F
                FPC 9 xlp3 TSen             OK          45 degrees C / 113 degrees F
                FPC 9 xlp3 Chip             OK          41 degrees C / 105 degrees F
                FPC 10 Intake               OK          32 degrees C / 89 degrees F
                FPC 10 Exhaust A            OK          44 degrees C / 111 degrees F
                FPC 10 Exhaust B            OK          53 degrees C / 127 degrees F
                FPC 10 LU 0 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 LU 0 Chip            OK          52 degrees C / 125 degrees F
                FPC 10 LU 1 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 LU 1 Chip            OK          44 degrees C / 111 degrees F
                FPC 10 LU 2 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 LU 2 Chip            OK          50 degrees C / 122 degrees F
                FPC 10 LU 3 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 LU 3 Chip            OK          58 degrees C / 136 degrees F
                FPC 10 XM 0 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 XM 0 Chip            OK          53 degrees C / 127 degrees F
                FPC 10 XF 0 TSen            OK          43 degrees C / 109 degrees F
                FPC 10 XF 0 Chip            OK          64 degrees C / 147 degrees F
                FPC 10 PLX Switch TSen      OK          43 degrees C / 109 degrees F
                FPC 10 PLX Switch Chip      OK          44 degrees C / 111 degrees F
                FPC 11 Intake               OK          32 degrees C / 89 degrees F
                FPC 11 Exhaust A            OK          41 degrees C / 105 degrees F
                FPC 11 Exhaust B            OK          56 degrees C / 132 degrees F
                FPC 11 LU 0 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 LU 0 Chip            OK          50 degrees C / 122 degrees F
                FPC 11 LU 1 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 LU 1 Chip            OK          47 degrees C / 116 degrees F
                FPC 11 LU 2 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 LU 2 Chip            OK          52 degrees C / 125 degrees F
                FPC 11 LU 3 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 LU 3 Chip            OK          60 degrees C / 140 degrees F
                FPC 11 XM 0 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 XM 0 Chip            OK          56 degrees C / 132 degrees F
                FPC 11 XF 0 TSen            OK          45 degrees C / 113 degrees F
                FPC 11 XF 0 Chip            OK          65 degrees C / 149 degrees F
                FPC 11 PLX Switch TSen      OK          45 degrees C / 113 degrees F
                FPC 11 PLX Switch Chip      OK          46 degrees C / 114 degrees F
Fans            Top Fan Tray Temp           OK          34 degrees C / 93 degrees F
                Top Tray Fan 1              OK          Spinning at normal speed
                Top Tray Fan 2              OK          Spinning at normal speed
                Top Tray Fan 3              OK          Spinning at normal speed
                Top Tray Fan 4              OK          Spinning at normal speed
                Top Tray Fan 5              OK          Spinning at normal speed
                Top Tray Fan 6              OK          Spinning at normal speed
                Top Tray Fan 7              OK          Spinning at normal speed
                Top Tray Fan 8              OK          Spinning at normal speed
```


Copyright © 2016, Juniper Networks, Inc.

```
Top Tray Fan 9                  OK          Spinning at normal speed
Top Tray Fan 10                 OK          Spinning at normal speed
Top Tray Fan 11                 OK          Spinning at normal speed
Top Tray Fan 12                 OK          Spinning at normal speed
Bottom Fan Tray Temp            OK          31 degrees C / 87 degrees F
Bottom Tray Fan 1               OK          Spinning at normal speed
Bottom Tray Fan 2               OK          Spinning at normal speed
Bottom Tray Fan 3               OK          Spinning at normal speed
Bottom Tray Fan 4               OK          Spinning at normal speed
Bottom Tray Fan 5               OK          Spinning at normal speed
Bottom Tray Fan 6               OK          Spinning at normal speed
Bottom Tray Fan 7               OK          Spinning at normal speed
Bottom Tray Fan 8               OK          Spinning at normal speed
Bottom Tray Fan 9               OK          Spinning at normal speed
Bottom Tray Fan 10              OK          Spinning at normal speed
Bottom Tray Fan 11              OK          Spinning at normal speed
Bottom Tray Fan 12              OK          Spinning at normal speed
  OK
```

Copyright © 2016, Juniper Networks, Inc.

show chassis fpc (View)

Syntax **show chassis fpc**
<**detail** < ***fpc-slot*** >| <node (***node-id*** | local | primary)>> |
<node (***node-id*** | local | primary)> |
<**pic-status** < ***fpc-slot*** >| <node (***node-id*** | local | primary)>>

Release Information Command modified in Junos OS Release 9.2.
Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.

NOTE: On SRX5K-MPC3-40G10G (IOC3), all four PICs cannot be powered on. A maximum of two PICs can be powered on at the same time. By default, PIC0 and PIC1 are online.

Use the **set chassis fpc <slot> pic <pic> power off** command to choose the PICs you want to power on.
When you use the **set chassis fpc <slot> pic <pic> power off** command to power off PIC0 and PIC1, PIC2 and PIC3 are automatically turned on.
When you switch from one set of PICs to another set of PICs using the **set chassis fpc <slot> pic <pic> power off** command again, ensure that there is 60 seconds duration between the two actions, otherwise core files are seen during the configuration.
The Table 185 summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 185: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary

CLI Configuration	PIC Selection
Default (i.e. no CLI configuration)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1, PIC-2 and PIC-3 powered OFF	Online: PIC-0 Offline: PIC-1, PIC-2, PIC-3
PIC-0, PIC-2 and PIC-3 powered OFF	Online: PIC-1 Offline: PIC-0, PIC-2, PIC-3
PIC-0, PIC-1 and PIC-3 powered OFF	Online: PIC-2 Offline: PIC-0, PIC-1, PIC-3
PIC-0, PIC-1 and PIC-2 powered OFF	Online: PIC-3 Offline: PIC-0, PIC-1, PIC-2

Copyright © 2016, Juniper Networks, Inc.

Table 185: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary *(continued)*

CLI Configuration	PIC Selection
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1 and PIC-2 powered OFF	Online: PIC-0, PIC-3 Offline: PIC-1, PIC-2
PIC-0 and PIC-3 powered OFF	Online: PIC-2, PIC-1 Offline: PIC-0, PIC-3
PIC-0 and PIC-1 powered OFF	Online: PIC-2, PIC-3 Offline: PIC-0, PIC-1
All other combinations of PICs being powered OFF (Invalid)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3 Default PICs will be selected for the invalid combinations. Also, a system log message will be displayed to indicate the invalid combination PIC selection.

Description Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options

- none—Display status information for all FPCs.
- **detail**—(Optional) Display detailed FPC status information.
- ***fpc-slot*** —(Optional) Display information about the FPC in this slot.
- **node**—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Copyright © 2016, Juniper Networks, Inc.

- **pic-status**—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see ***fpc-slot***).

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

List of Sample Output
show chassis fpc on page 1935
show chassis fpc (SRX5600 and SRX5800 devices) on page 1935
show chassis fpc (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3) on page 1935
show chassis fpc detail 2 on page 1936
show chassis fpc pic-status (SRX5600 and SRX5800 devices) on page 1936
show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2) on page 1936
show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC) on page 1937
show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured) on page 1937
show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP) on page 1938
show chassis fpc pic-status(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured) on page 1938
show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices) on page 1938
show chassis fpc pic-status for HA(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3) on page 1939

Output Fields Table 186 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 186: show chassis fpc Output Fields

Field Name	Field Description
Slot or Slot State	Slot number and state. The state can be one of the following conditions: • **Dead**—Held in reset because of errors. • **Diag**—Slot is being ignored while the device is running diagnostics. • **Dormant**—Held in reset. • **Empty**—No FPC is present. • **Online**—FPC is online and running. • **Present**—FPC is detected by the device, but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either **Hardware Not Supported** or **Hardware Not In Right Slot**. FPC is coming up but not yet online. • **Probed**—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • **Probe-wait**—Waiting to be probed.
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.

Copyright © 2016, Juniper Networks, Inc.

Table 186: show chassis fpc Output Fields *(continued)*

Field Name	Field Description
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
Start Time	Time when the Routing Engine detected that the FPC was running.
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.
PIC type	(**pic-status** output only) Type of FPC.

Sample Output

show chassis fpc

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)     Memory    Utilization (%)
     Slot State       (C)    Total  Interrupt        DRAM (MB) Heap  Buffer
      0  Online          -------------------- CPU less FPC --------------------
      1  Online          ---------------------  Not Usable --------------------
      2  Online          -------------------- CPU less FPC --------------------
```

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)       Memory    Utilization (%)
Slot State            (C)   Total  Interrupt         DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Empty
  3  Online           37       3          0             1024        7     42
  4  Empty
  5  Empty
  6  Online           30       8          0             1024        23    30
  7  Empty
  8  Empty
  9  Empty
 10  Empty
 11  Empty
```

show chassis fpc

(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory

                                    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB)

                                  Heap     Buffer
  0  Online           36     20          0       20     19     19    1024

                                    4         26
  1  Online           35      8          0        8      8      8    2048

                                   12         14
  2  Online           40     21          0       20     20     20    3584

                                    5         13
```

Sample Output

show chassis fpc detail 2

```
user@host> show chassis fpc detail 2
Slot 2 information:
  State                                 Online
  Temperature                            37
  Total CPU DRAM                       1024 MB
  Total RLDRAM                          0 MB
  Total DDR DRAM                        0 MB
  Start time:                             2012-07-18 07:18:50 PDT
  Uptime:                                 4 days, 21 hours, 51 minutes, 59 seconds

  Max Power Consumption                 0 Watts
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
Slot 3   Online       SRX5k SPC
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
Slot 6   Online       SRX5k DPC 4x 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

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

Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 3   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

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

Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 6   Offline      SRX5k SPC II
Slot 9   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Online       SRX5k IOC II
  PIC 0  Online       1x 100GE CFP
  PIC 2  Online       2x 40GE QSFP+
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

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

Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
```

Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

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

node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Offline      SRX5k SPC II
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 9   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP
  PIC 1  Online       10x 1GE(LAN) SFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Offline      SRX5k IOC II
```

show chassis fpc pic-status (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow

node1:
--------------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status for HA (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline

node1:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline
```

Copyright © 2016, Juniper Networks, Inc.

show chassis fpc errors

Syntax show chassis fpc errors;

Release Information Command introduced in Junos OS Release 15.1X49-D40.

Description Display chassis error information including FPC number, severity of error, number of error occurred, cleared, threshold, and corresponding action.

Error Severity Level	Default Threshold	Default Action
Fatal	1	Restart the FPC.
Major	1	Get the current state of the FPC and raise an alarm.
Minor	10	Write a log for the event.

Required Privilege Level view

Related Documentation
- fpc error on page 1854

List of Sample Output show chassis fpc errors on page 1941

Output Fields Table 187 lists the output fields for the **show chassis fpc errors** command. Output fields are listed in the approximate order in which they appear.

Table 187: show chassis fpc errors Output Fields

Field Name	Field Description
FPC	The FPC number.
Level	The severity of the error. It can be configured as follows: • **fatal**—Fatal error on the FPC. • **major**—Major error on the FPC. • **minor**—Minor error on the FPC.
Occurred	Number of error instances that have occurred.
Cleared	Number of error instances that have been cleared.
Threshold	Configured threshold value. The associated detection and recovery actions are triggered when this threshold value is exceeded.

Copyright © 2016, Juniper Networks, Inc.

Table 187: show chassis fpc errors Output Fields *(continued)*

Field Name	Field Description
Action	The detection and recovery actions that are triggered when the threshold value is exceeded. • Restart the FPC. • Get the current state of the FPC and raise an alarm. • Write a log for the event.

Sample Output

show chassis fpc errors

```
user@host> show chassis fpc errors
FPC  Level Occurred Cleared Threshold Action-Taken Action
0   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
1   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
2   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
4   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
5   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
6   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
7   Minor      0       0     10          LOG|
    Major      0       0      1        GET STATE|ALARM|
    Fatal      0       0      1         RESET
    Pfe-State: pfe-0 -ENABLED | pfe-1 -ENABLED | pfe-2 -ENABLED | pfe-3 -ENABLED
 | pfe-4 -ENABLED | pfe-5 -ENABLED | pfe-6 -ENABLED | pfe-7 -ENABLED |
```

Copyright © 2016, Juniper Networks, Inc.

show chassis hardware (View)

Syntax **show chassis hardware**
<clei-models | detail | extensive | models | node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include **node** option.

Description Display chassis hardware information.

Options

- **clei-models**—(Optional) Display Common Language Equipment Identifier Code (CLEI) barcode and model number for orderable field-replaceable units (FRUs).
- **detail | extensive**—(Optional) Display the specified level of output.
- **models**—(Optional) Display model numbers and part numbers for orderable FRUs.
- **node**—(Optional) For chassis cluster configurations, display chassis hardware information on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Interface Naming Conventions on page 2446

Output Fields Table 188 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 188: show chassis hardware Output Fields

Field Name	Field Description
Item	Chassis component—Information about the backplane; power supplies; fan trays; Routing Engine; each Physical Interface Module (PIM)—reported as FPC and PIC—and each fan, blower, and impeller.
Version	Revision level of the chassis component.
Part Number	Part number for the chassis component.
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the device chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the device chassis.
Assb ID or Assembly ID	Identification number that describes the FRU hardware.

Copyright © 2016, Juniper Networks, Inc.

Table 188: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
FRU model number	Model number of FRU hardware component.
CLEI code	Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.
EEPROM Version	ID EEPROM version used by hardware component: 0x01 (version 1) or 0x02 (version 2).

Copyright © 2016, Juniper Networks, Inc.

Table 188: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
Description	Brief description of the hardware item: • Type of power supply. • Switch Control Board (SCB) Starting with Junos OS Release 12.1X47-D15, the SRX5K-SCBE (SCB2) is introduced. ◦ There are three SCB slots in SRX5800 devices. The third slot can be used for an SCB or an FPC. When an SRX5K-SCB was used , the third SCB slot was used as an FPC. SCB redundancy is provided in chassis cluster mode. ◦ With an SCB2, a third SCB is supported. If a third SCB is plugged in, it provides intra-chassis fabric redundancy. ◦ The Ethernet switch in the SCB2 provides the Ethernet connectivity among all the FPCs and the Routing Engine. The Routing Engine uses this connectivity to distribute forwarding and routing tables to the FPCs. The FPCs use this connectivity to send exception packets to the Routing Engine. ◦ Fabric connects all FPCs in the data plane. The Fabric Manager executes on the Routing Engine and controls the fabric system in the chassis. Packet Forwarding Engines on the FPC and fabric planes on the SCB are connected through HSL2 channels. ◦ SCB2 supports HSL2 with both 3.11 Gbps and 6.22 Gbps (SerDes) link speed and various HSL2 modes. When an FPC is brought online, the link speed and HSL2 mode are determined by the type of FPC. Starting with Junos OS Release 15.1X49-D10, the SRX5K-SCB3 (SCB3) with enhanced midplanes is introduced. ◦ All existing SCB software that is supported by SCB2 is supported on SCB3. ◦ SRX5K-RE-1800X4 (RE2). Mixed Routing Engine use is not supported. ◦ SCB3 works with the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), SRX5K-MPC3-40G10G (IOC3), and SRX5K-SPC-4-15-320 (SPC2) with current midplanes and the new enhanced midplanes. ◦ Mixed SCB use is not supported. If an SCB2 and an SCB3 are used, the system will only power on the master Routing Engine's SCB and will power off the other SCBs. Only the SCB in slot 0 will be powered on and a system log is generated. ◦ SCB3 supports up to 400 Gbps per slot with old midplanes and up to 500 Gbps per slot with new midplanes. ◦ SCB3 supports fabric intra-chassis redundancy. ◦ SCB3 supports the same chassis cluster function as the SRX5K-SCB (SCB1) and the SRX5K-SCBE (SCB2), except for in-service software upgrade (ISSU) and in-service hardware upgrade (ISHU). ◦ SCB3 has a second external Ethernet port. ◦ Fabric bandwidth increasing mode is not supported.

Copyright © 2016, Juniper Networks, Inc.

Table 188: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
	• Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • IOCs Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced. ◦ IOC3 has two types of IOC3 MPCs, which have different built-in MICs: the 24x10GE + 6x40GE MPC and the 2x100GE + 4x10GE MPC. ◦ IOC3 supports SCB3 and SRX5000 line backplane and enhanced backplane. ◦ IOC3 can only work with SRX5000 line SCB2 and SCB3. If an SRX5000 line SCB is detected, IOC3 will be offline, an FPC misconfiguration alarm will be raised, and a system log message is generated. ◦ IOC3 interoperates with SCB2 and SCB3. ◦ IOC3 interoperates with the SRX5K-SPC-4-15-320 (SPC2) and the SRX5K-MPC (IOC2). ◦ The maximum power consumption for one IOC3 is 645W. An enhanced power module must be used. ◦ The IOC3 does not support the following command to set a PIC to go offline or online: **request chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot> <offline | online>** . ◦ IOC3 supports 240 Gbps of throughput with the enhanced SRX5000 line backplane. ◦ Chassis cluster functions the same as for the SRX5000 line IOC2. ◦ IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. ◦ IOC3 supports ISSU and ISHU in chassis cluster mode. ◦ IOC3 supports intra-FPC and and Inter-FPC Express Path (previously known as *services offloading*) with IPv4. ◦ NAT of IPv4 and IPv6 in normal mode and IPv4 for Express Path mode. ◦ All four PICs on the 24x10GE + 6x40GE cannot be powered on. A maximum of two PICs can be powered on at the same time. Use the **set chassis fpc <slot> pic <pic> power off** command to choose the PICs you want to power on. NOTE: Fabric bandwidth increasing mode is not supported on IOC3. • SRX Clustering Module (SCM) • Fan tray • For hosts, the Routing Engine type. ◦ Starting with Junos OS Release 12.1X47-D15, the SRX5K-RE-1800X4 (RE2) Routing Engine is introduced. ◦ The RE2 has an Intel Quad core Xeon processor, 16 GB of DRAM, and a 128-GB solid-state drive (SSD). The number 1800 refers to the speed of the processor (1.8 GHz). The maximum required power for this Routing Engine is 90W. NOTE: The RE2 provides significantly better performance than the previously used Routing Engine, even with a single core.

Copyright © 2016, Juniper Networks, Inc.

show chassis hardware

show chassis hardware

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                CM0715AK0021      SRX1500
Midplane         REV 08   750-058562   ACMA4255          SRX1500
CB 0             REV 08   711-053838   ACMA7529          CPU Board SRX700E
Routing Engine 0          BUILTIN      BUILTIN           SRX Routing Engine
FPC 0            REV 07   711-053832   ACMA3311          FEB
  PIC 0                   BUILTIN      BUILTIN           12x1G-T-4x1G-SFP-4x10G
    Xcvr 12      REV 01   740-014132   61521013          SFP-T
    Xcvr 13      REV 02   740-013111   A281604           SFP-T
    Xcvr 14      REV 02   740-011613   NRN30NV           SFP-SX
    Xcvr 15      REV 02   740-011613   NRN2PWV           SFP-SX
    Xcvr 16      REV 01   740-021308   AJA17B5           SFP+-10G-SR
    Xcvr 17      REV 01   740-021308   MSP056B           SFP+-10G-SR
    Xcvr 18      REV 01   740-031980   AS920WJ           SFP+-10G-SR
    Xcvr 19      REV 01   740-031980   AS92W5N           SFP+-10G-SR
Power Supply 0   REV 01   740-055217   1EDP4250OJZ       PS 400W 90-264V AC in
Fan Tray 0                                              SRX1500 0, Front to Back
 Airflow - AFO
Fan Tray 1                                              SRX1500 1, Front to Back
 Airflow - AFO
Fan Tray 2                                              SRX1500 2, Front to Back
 Airflow - AFO
Fan Tray 3                                              SRX1500 3, Front to Back
 Airflow - AFO
```

show chassis hardware (SRX5600 and SRX5800 devices for SRX5K-MPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN12170EAAGA      SRX 5800
Midplane         REV 01   710-041799   ACAX3849          SRX 5800 Backplane
FPM Board        REV 01   710-024632   CAAX7297          Front Panel Display
PDM              Rev 03   740-013110   QCS170250DU       Power Distribution Modu
le
PEM 0            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC i
n
PEM 1            Rev 03   740-034724   QCS17020203C      PS 4.1kW; 200-240V AC i
n
PEM 2            Rev 04   740-034724   QCS17100200A      PS 4.1kW; 200-240V AC i
n
PEM 3            Rev 03   740-034724   QCS17080200M      PS 4.1kW; 200-240V AC i
n
Routing Engine 0 REV 11   740-023530   9012047437        SRX5k RE-13-20
CB 0             REV 09   710-024802   CAAX7202          SRX5k SCB
CB 1             REV 09   710-024802   CAAX7157          SRX5k SCB
FPC 0            REV 07   750-044175   CAAD0791          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 07   750-044175   CAAD0751          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 2           REV 28  750-020751   CAAW1817         SRX5k DPC 4X 10GE
  CPU           REV 04  710-024633   CAAZ5269         SRX5k DPC PMB
  PIC 0                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
    Xcvr 0      REV 02  740-014289   T10A00404        XFP-10G-SR
  PIC 1                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 2                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 3                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
FPC 6           REV 02  750-044175   ZY2552           SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
FPC 9           REV 10  750-044175   CAAP5932         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 10          REV 22  750-043157   ZH8192           SRX5k IOC II   CPU
     REV 08   711-043360   YX3879           SRX5k MPC PMB
  MIC 0         REV 01  750-049488   YZ2084           10x 10GE SFP+
    PIC 0               BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0    REV 01  740-031980   AMB0HG3          SFP+-10G-SR
      Xcvr 1    REV 01  740-031980   AM20B6F          SFP+-10G-SR
  MIC 1         REV 19  750-049486   CAAH3504         1x 100GE CFP
    PIC 2               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   X000D375         CFP-100G-SR10
FPC 11          REV 07.04.07 750-043157 CAAJ8771      SRX5k IOC II   CPU
     REV 08   711-043360   CAAJ3881           SRX5k MPC PMB
  MIC 0         REV 19  750-049486   CAAH0979         1x 100GE CFP
    PIC 0               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   UP1020Z          CFP-100G-SR10
  MIC 1         REV 08  750-049487   CAAM1160         2x 40GE QSFP+
    PIC 2               BUILTIN      BUILTIN          2x 40GE QSFP+
      Xcvr 0    REV 01  740-032986   QB151094         QSFP+-40G-SR4
      Xcvr 1    REV 01  740-032986   QB160509         QSFP+-40G-SR4
Fan Tray 0      REV 04  740-035409   ACAE0875         Enhanced Fan Tray
Fan Tray 1      REV 04  740-035409   ACAE0876         Enhanced Fan Tray
```

show chassis hardware (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN108DA5AAGA      SRX 5800
Midplane         REV 02   710-013698   TR0037            SRX 5600 Midplane
FPM Board        REV 02   710-014974   JY4635            Front Panel Display
PDM             Rev 02   740-013110   QCS10465005       Power Distribution Module
PEM 0            Rev 03   740-023514   QCS11154040       PS 1.7kW; 200-240VAC in
PEM 2            Rev 02   740-023514   QCS10504014       PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 05   740-015113   1000681023        RE-S-1300
CB 0             REV 05   710-013385   JY4775            SRX5k SCB
FPC 1            REV 17   750-020751   WZ6349            SRX5k DPC 4X 10GE
  CPU            REV 02   710-024633   WZ0718            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0                NON-JNPR     C724XM088         XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 02   740-011571   C831XJ08S         XFP-10G-SR
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
FPC 3            REV 22   750-043157   ZH8189            SRX5k IOC II
```

Copyright © 2016, Juniper Networks, Inc.

```
  CPU            REV 06   711-043360   YX3912           SRX5k MPC PMB
  MIC 0          REV 01   750-055732   CACF9115         20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 2     REV 02   740-013111   B358549          SFP-T
      Xcvr 9     REV 02   740-011613   PNB1FQS          SFP-SX
    PIC 1                 BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 9     REV 02   740-011613   PNB1FFF          SFP-SX
FPC 5            REV 01   750-027945   JW9665           SRX5k FIOC
  CPU
FPC 8            REV 08   750-023996   XA7234           SRX5k SPC
  CPU            REV 02   710-024633   XA1599           SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN          SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
Fan Tray 0       REV 03   740-014971   TP0902           Fan Tray
Fan Tray 1       REV 01   740-014971   TP0121           Fan Tray
```

show chassis hardware (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4 (RE2))

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                               JN122A040AGA      SRX5800
Midplane         REV 01   710-041799   ACRA7817          SRX5800 Backplane
FPM Board        REV 01   760-058099   CACA2100          Front Panel Display
PDM              Rev 03   740-013110   QCS1739517Z       Power Distribution Modu

                                       le
PEM 0            Rev 05   740-034724   QCS17460203K      PS 4.1kW; 200-240V AC i

                                       n
PEM 1            Rev 04   740-034724   QCS172302017      PS 4.1kW; 200-240V AC i

                                       n
Routing Engine 0 REV 01   740-056658   9013040855        SRX5k RE-1800X4
Routing Engine 1
CB 0             REV 01   750-056587   CACG1424          SRX5k SCB II
CB 1             REV 01   750-056587   CACC9307          SRX5k SCB II
CB 2             REV 01   750-056587   CAAZ1128          SRX5k SCB II
FPC 0            REV 10   750-056758   CACS2667          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4092          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACV0038          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 10   750-043157   CACB6877          SRX5k IOC II
  CPU            REV 04   711-043360   CACH6074          SRX5k MPC PMB
```


Copyright © 2016, Juniper Networks, Inc.

```
  MIC 0          REV 19   750-049486   CAAH3504          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020Z           CFP-100G-SR10
  MIC 1          REV 04   750-049488   CACB6429          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AP21RJ5           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   AP21RLJ           SFP+-10G-SR
      Xcvr 2     REV 01   740-030658   AD1148A0AYC       SFP+-10G-USR
      Xcvr 3     REV 01   740-031980   B11E02718         SFP+-10G-SR
FPC 4            REV 10   750-056758   CACW0706          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 7            REV 10   750-056758   CACS2725          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 8            REV 11   750-043157   CABN4955          SRX5k IOC II
  CPU            REV 04   711-043360   CACT9926          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0979          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP2077V           CFP-100G-SR10
FPC 9            REV 10   750-056758   CACW0755          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 10           REV 07   750-044175   CAAD0747          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
Fan Tray 0       REV 04   740-035409   ACAE2294          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE2099          Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1235BC7AGA      SRX5800
Midplane         REV 01   710-024803   ACRC3244          SRX5800 Backplane
FPM Board        REV 01   710-024632   CACA2108          Front Panel Display
PDM              Rev 03   740-013110   QCS1739519B       Power Distribution Module
PEM 0            Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01   740-056658   9009153221        SRX5k RE-1800X4
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
FPC 0            REV 18   750-054877   CACH4004          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4082          SRX5k SPC II
```

Copyright © 2016, Juniper Networks, Inc.

```
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 2             REV 10    750-056758   CACW0713         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 3             REV 11    750-043157   CACA8792         SRX5k IOC II
  CPU               REV 04    711-043360   CACA8809         SRX5k MPC PMB
  MIC 0             REV 19    750-049486   CAAH3485         1x 100GE CFP
    PIC 0                     BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0        REV 01    740-035329   UNM0G3C          CFP-100G-SR10
  MIC 1             REV 04    750-049488   CABX0782         10x 10GE SFP+
    PIC 2                     BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0        REV 01    740-031980   AMB0HX3          SFP+-10G-SR
      Xcvr 1        REV 01    740-031980   ANT0E6V          SFP+-10G-SR
      Xcvr 2        REV 01    740-031980   ANR0ZVY          SFP+-10G-SR
      Xcvr 3        REV 01    740-031980   AP308ZU          SFP+-10G-SR
FPC 4             REV 10    750-044175   CAAS8024         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 7             REV 10    750-056758   CACS5126         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 8             REV 11    750-043157   CACA8798         SRX5k IOC II
  CPU               REV 04    711-043360   CACA8826         SRX5k MPC PMB
  MIC 0             REV 19    750-049486   CAAH0996         1x 100GE CFP
    PIC 0                     BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0        REV 01    740-035329   UP30A6N          CFP-100G-SR10
FPC 9             REV 07    750-044175   CAAD0745         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 10            REV 18    750-054877   CACD2570         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
Fan Tray 0        REV 04    740-035409   ACAE2122         Enhanced Fan Tray
Fan Tray 1        REV 04    740-035409   ACAE2254         Enhanced Fan Tray
```

show chassis hardware (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item            Version  Part number  Serial number     Description
Chassis                               JN1250870AGB      SRX5600
Midplane        REV 01   760-063936   ACRE2578          Enhanced SRX5600 Midplane
```


Copyright © 2016, Juniper Networks, Inc.

```
FPM Board        REV 02   710-017254   KD9027            Front Panel Display
PEM 0            Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V A

                                       C in
PEM 1            Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V A

                                       C in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0            REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU            REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4            REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5            REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                 Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board        test     710-017254   test                Front Panel Display
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry
Module
PEM 1            Rev 01   740-038514   QCS120311003      DC 2.6kW Power Entry
Module
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 11   750-043157   CACY1592          SRX5k IOC II
```

```
  CPU            REV 04   711-043360   CACZ8831          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACN0239          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-031980   ARN23HW           SFP+-10G-SR
      Xcvr 8     REV 01   740-031980   ARN2FVW           SFP+-10G-SR
      Xcvr 9     REV 01   740-031980   ARN2YVM           SFP+-10G-SR
FPC 5            REV 10   750-056758   CADA8736          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                Enhanced Fan Tray
```

show chassis hardware detail

show chassis hardware detail (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and (SRX5K-RE-1800X4 (RE2)

```
user@host> show chassis hardware detail
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN122A040AGA      SRX5800
Midplane         REV 01   710-041799   ACRA7817          SRX5800 Backplane
FPM Board        REV 01   760-058099   CACA2100          Front Panel Display
PDM              Rev 03   740-013110   QCS1739517Z       Power Distribution Module
PEM 0            Rev 05   740-034724   QCS17460203K      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 04   740-034724   QCS172302017      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01   740-056658   9013040855        SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200269450529 741 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1
CB 0             REV 01   750-056587   CACG1424          SRX5k SCB II
CB 1             REV 01   750-056587   CACC9307          SRX5k SCB II
CB 2             REV 01   750-056587   CAAZ1128          SRX5k SCB II
FPC 0            REV 10   750-056758   CACS2667          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4092          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACV0038          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
FPC 3            REV 10   750-043157   CACB6877         SRX5k IOC II
  CPU            REV 04   711-043360   CACH6074         SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH3504         1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020Z          CFP-100G-SR10
  MIC 1          REV 04   750-049488   CACB6429         10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AP21RJ5          SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   AP21RLJ          SFP+-10G-SR
      Xcvr 2     REV 01   740-030658   AD1148A0AYC      SFP+-10G-USR
      Xcvr 3     REV 01   740-031980   B11E02718        SFP+-10G-SR
FPC 4            REV 10   750-056758   CACW0706         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN          SPU Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
FPC 7            REV 10   750-056758   CACS2725         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN          SPU Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
FPC 8            REV 11   750-043157   CABN4955         SRX5k IOC II
  CPU            REV 04   711-043360   CACT9926         SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0979         1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP2077V          CFP-100G-SR10
FPC 9            REV 10   750-056758   CACW0755         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN          SPU Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
FPC 10           REV 07   750-044175   CAAD0747         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
Fan Tray 0       REV 04   740-035409   ACAE2294         Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE2099         Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1235BC7AGA      SRX5800
Midplane         REV 01   710-024803   ACRC3244          SRX5800 Backplane
FPM Board        REV 01   710-024632   CACA2108          Front Panel Display
PDM              Rev 03   740-013110   QCS1739519B       Power Distribution Module
PEM 0            Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01   740-056658   9009153221        SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200298450703 72 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
```

Copyright © 2016, Juniper Networks, Inc.

```
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
FPC 0            REV 18   750-054877   CACH4004          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4082          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACW0713          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 11   750-043157   CACA8792          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8809          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH3485          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UNM0G3C           CFP-100G-SR10
  MIC 1          REV 04   750-049488   CABX0782          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AMB0HX3           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   ANT0E6V           SFP+-10G-SR
      Xcvr 2     REV 01   740-031980   ANR0ZVY           SFP+-10G-SR
      Xcvr 3     REV 01   740-031980   AP308ZU           SFP+-10G-SR
FPC 4            REV 10   750-044175   CAAS8024          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 7            REV 10   750-056758   CACS5126          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 8            REV 11   750-043157   CACA8798          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8826          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0996          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP30A6N           CFP-100G-SR10
FPC 9            REV 07   750-044175   CAAD0745          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 10           REV 18   750-054877   CACD2570          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
Fan Tray 0       REV 04   740-035409   ACAE2122          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE2254          Enhanced Fan Tray
```


Copyright © 2016, Juniper Networks, Inc.

show chassis hardware detail (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 SCB3 with enhanced midplanes and (SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware detail
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1250870AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2578          Enhanced SRX5600 Midplane
FPM Board         REV 02   710-017254   KD9027              Front Panel Display
PEM 0             Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V
AC in
PEM 1             Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V
AC in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
  ad0    3831 MB  UGB30SFA4000T2       SFA4000T2 000027A0 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-043         Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0             REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0            REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU            REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4            REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5            REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                 Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
```

```
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board         test     710-017254   test                Front Panel Display
PEM 0             Rev 01   740-038514   QCS114111003        DC 2.6kW Power Entry
Module
PEM 1             Rev 01   740-038514   QCS120311O0J        DC 2.6kW Power Entry
Module
Routing Engine 0 REV 01   740-056658   9009186342          SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200313161216 109 Compact Flash
  ad1   28843 MB  UGB94BPH32H0S2-KCI   11000160387         Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000       uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087       uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0             REV 01   750-062257   CAEB8178            SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN             SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN             SPU Cp
  PIC 1                   BUILTIN      BUILTIN             SPU Flow
  PIC 2                   BUILTIN      BUILTIN             SPU Flow
  PIC 3                   BUILTIN      BUILTIN             SPU Flow
FPC 4            REV 11   750-043157   CACY1592            SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8831            SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACN0239            10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN             10x 10GE SFP+
      Xcvr 7     REV 01   740-031980   ARN23HW             SFP+-10G-SR
      Xcvr 8     REV 01   740-031980   ARN2FVW             SFP+-10G-SR
      Xcvr 9     REV 01   740-031980   ARN2YVM             SFP+-10G-SR
FPC 5            REV 10   750-056758   CADA8736            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN             SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN             SPU Flow
  PIC 1                   BUILTIN      BUILTIN             SPU Flow
  PIC 2                   BUILTIN      BUILTIN             SPU Flow
  PIC 3                   BUILTIN      BUILTIN             SPU Flow
Fan Tray                                                  Enhanced Fan Tray
```

show chassis hardware extensive node 1

show chassis hardware extensive node 1 (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4)

```
user@host> show chassis hardware extensive node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1235BC7AGA      SRX5800
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               JN1235BC7AGA
 Assembly ID:  0x051a            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SRX5800
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 1a 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 33 35 42 43 37 41 47 41 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```


Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   710-024803   ACRC3244          SRX5800 Backplane
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          710-024803        S/N:               S/N ACRC3244
 Assembly ID:  0x091a            Assembly Version:  01.01
 Date:         02-26-2014        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5800 Backplane           FRU Model Number:  SRX5800-BP-A
 Board Information Record:
  Address 0x00: ad 01 08 00 4c 96 14 d3 28 00 00 ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 38 30 33 00 00
  Address 0x20: 53 2f 4e 20 41 43 52 43 33 32 34 34 00 1a 02 07
  Address 0x30: de ff ff ff ad 01 08 00 4c 96 14 d3 28 00 00 ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 42 50 2d 41 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 01   710-024632   CACA2108          Front Panel Display
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          710-024632        S/N:               S/N CACA2108
 Assembly ID:  0x096f            Assembly Version:  01.01
 Date:         02-05-2014        Assembly Flags:    0x00
 Version:      REV 01
 ID: Front Panel Display         FRU Model Number:  SRX5800-CRAFT-A
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 6f 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 36 33 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 41 32 31 30 38 00 05 02 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 43 52 41 46 54 2d 41 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PDM              Rev 03   740-013110   QCS1739519B       Power Distribution Module

 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-013110        S/N:               QCS1739519B
 Assembly ID:  0x0416            Assembly Version:  01.03
 Date:         10-26-2013        Assembly Flags:    0x00
 Version:      Rev 03
 ID: Power Distribution Module
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 16 01 03 52 65 76 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 31 33 31 31 30 00 00
  Address 0x20: 51 43 53 31 37 33 39 35 31 39 42 00 00 1a 0a 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 0            Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
```

```
 P/N:          740-034724        S/N:               QCS17230201Z
 Assembly ID:  0x044b            Assembly Version:  01.04
 Date:         06-04-2013        Assembly Flags:    0x00
 Version:      Rev 04
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  SRX5800-PWR-4100-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4b 01 04 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 34 37 32 34 00 00
  Address 0x20: 51 43 53 31 37 32 33 30 32 30 31 5a 00 04 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 50 57 52 2d 34 31 30 30 2d
  Address 0x60: 41 43 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-034724        S/N:               QCS174502014
 Assembly ID:  0x044b            Assembly Version:  01.05
 Date:         11-06-2013        Assembly Flags:    0x00
 Version:      Rev 05
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  SRX5800-PWR-4100-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4b 01 05 52 65 76 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 34 37 32 34 00 00
  Address 0x20: 51 43 53 31 37 34 35 30 32 30 31 34 00 06 0b 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 50 57 52 2d 34 31 30 30 2d
  Address 0x60: 41 43 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 01   740-056658   9009153221        SRX5k RE-1800X4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-056658        S/N:               9009153221
 Assembly ID:  0x0c1a            Assembly Version:  01.01
 Date:         07-22-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k RE-1800X4             FRU Model Number:  SRX5K-RE-1800X4
 Board Information Record:
  Address 0x00: 54 32 30 32 37 45 43 2d 34 34 47 42 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 36 36 35 38 00 00
  Address 0x20: 39 30 30 39 31 35 33 32 32 31 00 00 00 16 07 07
  Address 0x30: dd ff ff ff 54 32 30 32 37 45 43 2d 34 34 47 42
  Address 0x40: 00 00 00 00 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 52 45 2d 31 38 30 30 58 34 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 64 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200298450703 72 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
```


Copyright © 2016, Juniper Networks, Inc.

```
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACC9541
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-07-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 43 39 35 34 31 00 07 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACG1447
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-07-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 47 31 34 34 37 00 07 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACH9058
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-06-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 48 39 30 35 38 00 06 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware extensive node 1 (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware extensive node 1
```

```
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               JN124FEC0AGB
 Assembly ID:  0x051b            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x08
 ID: SRX5600
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 1b 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 34 46 45 43 30 41 47 42 08 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midpla

                                    ne
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          760-063936        S/N:               ACRE2946
 Assembly ID:  0x0914            Assembly Version:  01.01
 Date:         03-19-2015        Assembly Flags:    0x08
 Version:      REV 01            CLEI Code:         CLEI-CODE
 ID: SRX5600 Midplane            FRU Model Number:  SRX5600X-CHAS
 Board Information Record:
  Address 0x00: ad 01 08 00 88 a2 5e 12 68 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 14 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 36 33 39 33 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 52 45 32 39 34 36 08 13 03 07
  Address 0x30: df ff ff ff ad 01 08 00 88 a2 5e 12 68 00 ff ff
  Address 0x40: ff ff ff ff 01 43 4c 45 49 2d 43 4f 44 45 20 53
  Address 0x50: 52 58 35 36 30 30 58 2d 43 48 41 53 20 20 20 20
  Address 0x60: 20 20 20 20 20 20 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ba ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        test     710-017254   test              Front Panel Display
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          710-017254        S/N:               test
 Assembly ID:  0x01ff            Assembly Version:  01.00
 Date:         06-18-2007        Assembly Flags:    0x00
 Version:      test
 ID: Front Panel Display
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 01 ff 01 00 74 65 73 74 00 00 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 32 35 34 00 00
  Address 0x20: 74 65 73 74 00 00 00 00 00 00 00 00 00 12 06 07
  Address 0x30: d7 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 ff ff ff ff ff ff ff ff ff ff ff ff
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry
Module
```


Copyright © 2016, Juniper Networks, Inc.

```
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-038514        S/N:               QCS114111003
 Assembly ID:  0x044c            Assembly Version:  01.01
 Date:         10-14-2011        Assembly Flags:    0x00
 Version:      Rev 01
 ID: DC 2.6kW Power Entry Module FRU Model Number:  SRX5600-PWR-2400-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4c 01 01 52 65 76 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 38 35 31 34 00 00
  Address 0x20: 51 43 53 31 31 34 31 31 31 30 30 33 00 0e 0a 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 36 30 30 2d 50 57 52 2d 32 34 30 30 2d
  Address 0x60: 44 43 2d 53 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 01   740-038514   QCS120311100J      DC 2.6kW Power Entry
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-038514        S/N:               QCS120311100J
 Assembly ID:  0x044c            Assembly Version:  01.01
 Date:         01-17-2012        Assembly Flags:    0x00
 Version:      Rev 01
 ID: DC 2.6kW Power Entry Module FRU Model Number:  SRX5600-PWR-2400-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4c 01 01 52 65 76 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 38 35 31 34 00 00
  Address 0x20: 51 43 53 31 32 30 33 31 31 30 30 4a 00 11 01 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 36 30 30 2d 50 57 52 2d 32 34 30 30 2d
  Address 0x60: 44 43 2d 53 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-056658        S/N:               9009186342
 Assembly ID:  0x0c1a            Assembly Version:  01.01
 Date:         02-05-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         COUCATTBAA
 ID: SRX5k RE-1800X4             FRU Model Number:  SRX5K-RE-1800X4
 Board Information Record:
  Address 0x00: 54 32 30 32 37 45 43 2d 34 34 47 42 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 36 36 35 38 00 00
  Address 0x20: 39 30 30 39 31 38 36 33 34 32 00 00 00 05 02 07
  Address 0x30: de ff ff ff 54 32 30 32 37 45 43 2d 34 34 47 42
  Address 0x40: 00 00 00 00 01 43 4f 55 43 41 54 54 42 41 41 53
  Address 0x50: 52 58 35 4b 2d 52 45 2d 31 38 30 30 58 34 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 64 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200313161216 109 Compact Flash
  ad1   28843 MB  UGB94BPH32H0S2-KCI   11000160387       Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
```

Copyright © 2016, Juniper Networks, Inc.

```
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-062257        S/N:               CAEB8178
 Assembly ID:  0x0c59            Assembly Version:  01.01
 Date:         03-19-2015        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         CLEI-CODE
 ID: SRX5k SCB3                  FRU Model Number:  SRX5K-SCB3
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 59 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 32 35 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 42 38 31 37 38 00 13 03 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4c 45 49 2d 43 4f 44 45 20 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 33 20 20 20 20 20 20 20
  Address 0x60: 20 20 20 20 20 20 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 63 ff ff ff ff ff ff ff ff ff ff ff ff
FPC 2            REV 01   750-062243   CAED0386          SRX5k IOC3 24XGE+6XLG
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-062243        S/N:               CAED0386
 Assembly ID:  0x0c57            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5k IOC3 24XGE+6XLG
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 fe 0c 57 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 32 34 33 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 44 30 33 38 36 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
  Address 0x40: f2 06 00 ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  CPU             REV 01   711-062244   CADX8554          SRX5k IOC3 PMB
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-062244        S/N:               CADX8554
 Assembly ID:  0x0c5a            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5k IOC3 PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0c 5a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 36 32 32 34 34 00 00
  Address 0x20: 53 2f 4e 20 43 41 44 58 38 35 35 34 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 2a 47 7e 50 10 05 76 5c
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0ab5            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: 12x 10GE SFP+
```


Copyright © 2016, Juniper Networks, Inc.

```
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b5 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 a9 3c 00 00 00 00 0a b6 00 00
 PIC 1                  BUILTIN      BUILTIN           12x 10GE SFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab5            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 12x 10GE SFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b5 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 e9 b4 00 00 00 00 0a b5 00 00
 PIC 2                  BUILTIN      BUILTIN           3x 40GE QSFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab6            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 3x 40GE QSFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b6 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 03 e8 c4 33 24 3a 38 00 00 00 02
 PIC 3                  BUILTIN      BUILTIN           3x 40GE QSFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab6            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 3x 40GE QSFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b6 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x70: 00 00 00 00 c0 02 ab 1c 00 00 00 00 0a b5 00 00
  WAN MEZZ       REV 01   750-062682   CAEA4788          24x 10GE SFP+ Mezz
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-062682        S/N:               CAEA4788
 Assembly ID:  0x0c76            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: 24x 10GE SFP+ Mezz
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0c 76 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 36 38 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 41 34 37 38 38 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
FPC 4            REV 11   750-043157   CACY1592          SRX5k IOC II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-043157        S/N:               CACY1592
 Assembly ID:  0x0bd1            Assembly Version:  04.11
 Date:         07-30-2014        Assembly Flags:    0x00
 Version:      REV 11            CLEI Code:         COUIBCWBAA
 ID: SRX5k IOC II                FRU Model Number:  SRX5K-MPC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b d1 04 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 31 35 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 59 31 35 39 32 00 1e 07 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
  Address 0x40: f2 06 00 ff 01 43 4f 55 49 42 43 57 42 41 41 53
  Address 0x50: 52 58 35 4b 2d 4d 50 43 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 92 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 04   711-043360   CACZ8831          SRX5k MPC PMB
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-043360        S/N:               CACZ8831
 Assembly ID:  0x0bd2            Assembly Version:  01.04
 Date:         07-28-2014        Assembly Flags:    0x00
 Version:      REV 04
 ID: SRX5k MPC PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b d2 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 33 33 36 30 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 5a 38 38 33 31 00 1c 07 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 49 fa 60 10 40 05 76 5c
  MIC 1          REV 04   750-049488   CACN0239          10x 10GE SFP+
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-049488        S/N:               CACN0239
 Assembly ID:  0x0a88            Assembly Version:  02.04
 Date:         02-26-2014        Assembly Flags:    0x00
 Version:      REV 04            CLEI Code:         COUIBCXBAA
```


Copyright © 2016, Juniper Networks, Inc.

```
  ID: 10x 10GE SFP+               FRU Model Number:  SRX-MIC-10XG-SFPP
  Board Information Record:
   Address 0x00: 34 01 03 03 ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0a 88 02 04 52 45 56 20 30 34 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 34 39 34 38 38 00 00
   Address 0x20: 53 2f 4e 20 43 41 43 4e 30 32 33 39 00 1a 02 07
   Address 0x30: de ff ff ff 34 01 03 03 ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4f 55 49 42 43 58 42 41 41 53
   Address 0x50: 52 58 2d 4d 49 43 2d 31 30 58 47 2d 53 46 50 50
   Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 9f c0 03 e4 14 55 8a 95 a8 00 00 00 02
     PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
       Xcvr 7     REV 01   740-031980   ARN23HW           SFP+-10G-SR
       Xcvr 8     REV 01   740-031980   ARN2FVW           SFP+-10G-SR
       Xcvr 9     REV 01   740-031980   ARN2YVM           SFP+-10G-SR
FPC 5              REV 10   750-056758   CADA8736          SRX5k SPC II
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-056758        S/N:               CADA8736
  Assembly ID:  0x0b4f            Assembly Version:  01.10
  Date:         09-01-2014        Assembly Flags:    0x00
  Version:      REV 10            CLEI Code:         COUCATLBAB
  ID: SRX5k SPC II                FRU Model Number:  SRX5K-SPC-4-15-320
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 4f 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 37 35 38 00 00
   Address 0x20: 53 2f 4e 20 43 41 44 41 38 37 33 36 00 01 09 07
   Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
   Address 0x40: f2 06 00 ff 01 43 4f 55 43 41 54 4c 42 41 42 53
   Address 0x50: 52 58 35 4b 2d 53 50 43 2d 34 2d 31 35 2d 33 32
   Address 0x60: 30 00 00 00 00 00 43 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 50 ff ff ff ff ff ff ff ff ff ff ff ff
   CPU                   BUILTIN      BUILTIN           SRX5k DPC PPC
   PIC 0                 BUILTIN      BUILTIN           SPU Cp
  Jedec Code:   0x0000            EEPROM Version:    0x00
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0x0a20            Assembly Version:  00.00
  Date:         00-00-0000        Assembly Flags:    0x00
  ID: SPU Cp
  Board Information Record:
   Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  I2C Hex Data:
   Address 0x00: 00 00 00 00 0a 20 00 00 00 00 00 00 00 00 00 00
   Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
   Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
   Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x70: 00 00 00 00 de ad be ef 46 10 f0 d8 40 43 43 c0
   PIC 1                 BUILTIN      BUILTIN           SPU Flow
  Jedec Code:   0x0000            EEPROM Version:    0x00
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0x0a21            Assembly Version:  00.00
  Date:         00-00-0000        Assembly Flags:    0x00
  ID: SPU Flow
  Board Information Record:
   Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  I2C Hex Data:
```

Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 13 5b d0 40 43 43 c0
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0a21            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SPU Flow
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 0c 66 40 40 43 43 c0
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0a21            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SPU Flow
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 0e db 00 40 43 43 c0
Fan Tray                                                Enhanced Fan Tray
                                FRU Model Number:  SRX5600-HC-FAN
```

show chassis hardware clei-models

show chassis hardware clei-models (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4 (RE2))

```
user@host> show chassis hardware clei-models node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   710-024803                     SRX5800-BP-A
FPM Board        REV 01   710-024632                     SRX5800-CRAFT-A
PEM 0            Rev 04   740-034724                     SRX5800-PWR-4100-AC
PEM 1            Rev 05   740-034724                     SRX5800-PWR-4100-AC
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
```


Copyright © 2016, Juniper Networks, Inc.

```
CB 1             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
CB 2             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
FPC 0            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 1            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 2            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 3            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
FPC 4            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 7            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 8            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
FPC 9            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 10           REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
Fan Tray 0       REV 04   740-035409                     SRX5800-HC-FAN
Fan Tray 1       REV 04   740-035409                     SRX5800-HC-FAN
```

show chassis hardware clei-models (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware clei-models
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   760-063936   CLEI-CODE         SRX5600X-CHAS
FPM Board        REV 02   710-017254                     CRAFT-MX480-S
PEM 0            Rev 03   740-034701                     SRX5600-PWR-2520-AC-S
PEM 1            Rev 03   740-034701                     SRX5600-PWR-2520-AC-S
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-062257   CLEI-CODE         SRX5K-SCB3
FPC 0            REV 10   750-056758   COUCATLBAB        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 2            REV 01   750-062243
FPC 4            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
FPC 5            REV 05   750-044175   PROTOXCLEI        750-044175
  CPU                     BUILTIN
Fan Tray                                                 SRX5600-HC-FAN

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   760-063936   CLEI-CODE         SRX5600X-CHAS
PEM 0            Rev 01   740-038514                     SRX5600-PWR-2400-DC-S
PEM 1            Rev 01   740-038514                     SRX5600-PWR-2400-DC-S
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-062257   CLEI-CODE         SRX5K-SCB3
FPC 0            REV 07   750-044175   COUCASFBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 4            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
```

```
FPC 5            REV 10   750-056758   COUCATLBAB       SRX5K-SPC-4-15-320
  CPU                     BUILTIN
Fan Tray                                                 SRX5600-HC-FAN
```


Copyright © 2016, Juniper Networks, Inc.

show chassis pic (Security)

Syntax show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Release Information Command introduced in Junos OS Release 9.2.

Description Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options **fpc-slot *slot-number***—Display information about the FPC in the slot.

pic-slot *slot-number*—Display information about the PIC in this particular FPC slot.

Required Privilege Level view

Related Documentation

- *Interfaces Feature Guide for Security Devices*
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output show chassis pic fpc-slot pic-slot on page 1970

Output Fields Table 189 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 189: show chassis pic Output Fields

Field Name	Field Description
Type	PIC type.
State	Status of the PIC. State is displayed only when a PIC is in the slot. • **Online**— PIC is online and running. • **Offline**—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: **LH**, **LX**, or **SX.**

Copyright © 2016, Juniper Networks, Inc.

Table 189: show chassis pic Output Fields *(continued)*

Field Name	Field Description
PIC Port Information	Port-level information for the PIC. • Port—Port number • Cable type—Type of transceiver installed. • Fiber type—Type of fiber. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal.

Sample Output

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 10 pic-slot 0
FPC slot 10, PIC slot 0 information:
  Type                            10x 10GE SFP+
  State                           Online
  PIC version                  1.1
  Uptime                          6 days, 7 hours, 29 minutes, 28 seconds

PIC port information:
                          Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number       Wavelength
  0     10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm

        Xcvr vendor
        firmware version
        0.0

PIC port information:
                          Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number       Wavelength
  1     10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm

        Xcvr vendor
        firmware version
        0.0
```


Copyright © 2016, Juniper Networks, Inc.

show chassis power

Syntax show chassis power

Release Information Command modified in Junos OS Release 12.1X44-D10.

Description Display power limits and usage information for the Power Entry Modules (PEMs).

Options This command has no options.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- show chassis power sequence on page 1973

List of Sample Output show chassis power on page 1972

Output Fields Table 190 lists the output fields for the **show chassis power** command. Output fields are listed in the approximate order in which they appear.

Table 190: show chassis power Output Fields

Field Name	Field Description

PEM number	AC or DC PEM number on the chassis. The following output fields are displayed for the PEM:
	• State—State of the PEM:
	  ◦ Online—PEM is present in the slot and online.
	  ◦ Empty—PEM is not present in the slot.
	  ◦ Present—PEM is present in the slot, but not online.
	• AC Input —State of the AC input power feed with the number of active and expected feeds (1 or 2).
	• Capacity—Actual power input capacity with maximum capacity displayed (in parentheses) in watts.
	• DC Output—DC power output, in watts, for the specified zone, at the specified amps and voltage (A @ V), and load and percentage utilization of the maximum capacity for the zone.

System	Overall power statistics for the system zone:
	• Zone number:
	  ◦ Capacity—Maximum power capacity available for the zone, in watts.
	  ◦ Allocated power—Actual capacity allocated for the zone, in watts, with remaining power displayed in parentheses.
	  ◦ Actual usage—Actual power usage for the zone, in watts.
	• Total system capacity—Cumulative power capacity of all the zones, in watts.
	• Total remaining capacity—Difference between the total system capacity and cumulative allocated power of all zones, in watts.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show chassis power

```
user@host> show chassis power

PEM 0:
  State:     Present
  DC input:  Out of range (1 feed expected, 0 feed connected)
  DC input:  48.0 V input (0 mV)
  Capacity:  0 W (maximum 4100 W)

PEM 1:
  State:     Present
  DC input:  Out of range (1 feed expected, 0 feed connected)
  DC input:  48.0 V input (0 mV)
  Capacity:  0 W (maximum 4100 W)

PEM 2:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (57500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 1881 W (zone 0, 33 A at 57 V, 45% of capacity)

PEM 3:
  State:     Online
  DC input:  OK (2 feed expected, 2 feed connected)
  DC input:  48.0 V input (56500 mV)
  Capacity:  4100 W (maximum 4100 W)
  DC output: 2464 W (zone 1, 44 A at 56 V, 60% of capacity)

System:
  Zone 0:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   2920 W (1180 W remaining)
      Actual usage:      1881 W
  Zone 1:
      Capacity:          4100 W (maximum 4100 W)
      Allocated power:   3875 W (225 W remaining)
      Actual usage:      2464 W
  Total system capacity: 8200 W (maximum 8200 W)
  Total remaining power: 1405 W
```


Copyright © 2016, Juniper Networks, Inc.

show chassis power sequence

Syntax show chassis power sequence

Release Information Command modified in Junos OS Release 12.1X44-D10.

Description Display the power-on sequence for the FPCs in the chassis. The numbers indicate the slot number of the FPCs.

Options This command has no options.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- fru-poweron-sequence on page 1856

List of Sample Output show chassis power sequence on page 1973

Output Fields Table 191 lists the output fields for the **show chassis power sequence** command. Output fields are listed in the approximate order in which they appear.

Table 191: show chassis power sequence Output Fields

Field Name	Field Description
Chassis FRU Power Sequence	Power-on sequence for the FPCs in the chassis. The numbers indicate the slot number of the FPCs.

Sample Output

show chassis power sequence

```
user@host> show chassis power sequence
Chassis FRU Power On Sequence: 0 1 2 3 4 5 6 7 8 9 10 11
```

Copyright © 2016, Juniper Networks, Inc.

show firewall (View)

Syntax

show firewall
<filter *filter-name*>
<counter *counter-name*>
<log>
<prefix-action-stats>
<terse>

Release Information

Command introduced before Junos OS Release 10.0 .

Description

Display statistics about configured firewall filters.

Options

none—Display statistics about configured firewall filters.

filter *filter-name*—Name of a configured filter.

counter *counter-name*—Name of a filter counter.

log—Display log entries for firewall filters.

prefix-action-stats—Display prefix action statistics for firewall filters.

terse—Display firewall filter names only.

Required Privilege Level

view

Related Documentation

- *firewall*

List of Sample Output

show firewall on page 1975

Output Fields

Table 192 lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 192: show firewall Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level. When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either **-i** for an input filter or **-o** for an output filter. When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either **-in** for an input filter or **-out** for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, **__ls1/filter1**).

Copyright © 2016, Juniper Networks, Inc.

Table 192: show firewall Output Fields *(continued)*

Field Name	Field Description
Counters	Display filter counter information: • **Name**—Name of a filter counter that has been configured with the **counter** firewall filter action. • **Bytes**—Number of bytes that match the filter term under which the **counter** action is specified. • **Packets**—Number of packets that matched the filter term under which the **counter** action is specified.
Policers	Display policer information: • **Name**—Name of policer. • **Bytes**—Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. • **Packets**—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Sample Output

show firewall

```
user@host> show firewall
Filter: ef_path
Counters:
Name                                            Bytes              Packets
def-count                                       0                    0
video-count                                     0                    0
voice-count                                     0                    0

Filter: __default_bpdu_filter__

Filter: deep
Counters:
Name                                            Bytes              Packets
deep2                                           302076               5031

Filter: deep-flood
Counters:
Name                                            Bytes              Packets
deep_flood_def                                  302136               5032
deep1                                           0                    0
Policers:
Name                                                        Packets
deep-pol-op-first                                             0
```

show interfaces (View Aggregated Ethernet)

Syntax

show interfaces <ae*number* | reth*number* >
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command modified in Junos OS Release 10.2.

Description

Display status information about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, status information for all interfaces is displayed.

Options

- ae*number* | reth*number*—(Optional) Display standard information about the specified aggregated Ethernet interface or redundant Ethernet interface.
- brief | detail | extensive | terse—(Optional) Display the specified level of output.
- descriptions—(Optional) Display interface description strings.
- media—(Optional) Display media-specific information.
- snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.
- statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output

show interfaces extensive (Aggregated Ethernet) on page 1982

Output Fields

Table 193 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 193: Aggregated Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels
Enabled	State of the physical interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	All levels
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up or up to down. The format is **Last flapped:** ***year-month-day hours:minutes:seconds timezone*** **(*hours:minutes:seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and frame check sequence (FCS) errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid FCS. • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions** —Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Egress queues	Total number of egress queues supported on the specified interface	**detail extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Index	Index number of the logical interface (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	SNMP interface index number of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Statistics	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface. • **Bundle**—Information about input and output bundle rates. • **Link**—(**detail** and **extensive** only) Information about specific links in the aggregate, including link state and input and output rates.	**detail extensive** none

LACP info	Link Aggregation Control Protocol (LACP) information for each aggregated interface.
	• **Role** can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **System priority**—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• **System identifier**—Actor or partner system ID, encoded as a MAC address.
	• **Port priority**—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• **Port number**—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.
	• **Port key**—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.
		detail extensive none

LACP Statistics	LACP statistics for each aggregated interface. • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.	**detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Marker Statistic	(**detail** and **extensive** only) Information about 802.3ad marker protocol statistics on the specified links. • **Marker Rx**—Number of valid marker PDUs received on this aggregation port. • **Resp Tx**—Number of marker response PDUs transmitted on this aggregation port. • **Unknown Rx**—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown protocol data unit (PDU), or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type. • **Illegal Rx**—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).	**detail extensive** none
Flow Statistics	Flow statistics for each aggregated interface.	**detail extensive** none
Flow Input statistics	Statistics for packets received by the flow module.	**detail extensive** none
Flow Output statistics	Statistics for packets sent by the flow module.	**detail extensive** none

Flow error statistics (Packets dropped due to)	Packet drop statistics for the flow module.
	• **Address spoofing**—Packet dropped when the screen module detected address spoofing.
	• **Authentication failed**—Packet dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
	• **Incoming NAT errors**—Packet dropped because the source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
	• **Invalid zone received packet**—This counter is not currently in use.
	• **Multiple user authentications**—Packet dropped if it matches more than one policy that specifies user authentication. (Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy.)
	• **Multiple incoming NAT**—Packet dropped if source NAT is specified more than once. (Sometimes packets are looped through the system more than once.)
	• **No parent for a gate**—This counter is not currently in use.
	• **No one interested in self packets**—This counter is incremented for one of the following reasons:
	  ◦ The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool.
	  ◦ No service is interested in the to-self packet
	  ◦ When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
		detail extensive none

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output

Flow error statistics (Packets dropped due to)	Packet drop statistics for the flow module (continued).
	• **No minor session**—Packet dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
	• **No more sessions**—Packet dropped because there were no more free sessions available.
	• **No NAT gate**—This counter is not currently in use.
	• **No route present**—Packet dropped because a valid route was not available to forward the packet.
	• **No SA for incoming SPI**—Packet dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
	• **No tunnel found**—Packet dropped because a valid tunnel could not be found.
	• **No session for a gate**—Packet dropped by an ALG.
	• **No zone or NULL zone binding**—Packet dropped because its incoming interface was not bound to any zone.
	• **Policy denied**—The error counter is incremented for one of the following reasons:
	  ◦ Source or destination NAT (or both) has occurred and policy says to drop the packet.
	  ◦ Policy specifies user authentication, which failed.
	  ◦ Policy was configured to deny this packet.
	• **Security association not active**—Packet dropped because an IPsec packet was received for an inactive SA.
	• **TCP sequence number out of window**—TCP packet with a sequence number failed the TCP sequence number check that was received.
	• **Syn-attack protection**—Packet dropped because of SYN attack protection or SYN cookie protection.
	• **User authentication errors**—Packet dropped because policy requires authentication; however:
	  ◦ Only Telnet, FTP, and HTTP traffic can be authenticated.
	  ◦ The corresponding authentication entry could not be found, if web-auth is specified.
	  ◦ The maximum number of authenticated sessions per user was exceeded.
		detail extensive none

protocol-family	Protocol family configured on the logical interface.	**brief**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 193: Aggregated Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about protocol family flags.	**detail extensive** none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Policer	Policer to be evaluated when packets are received or transmitted on the interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 1973, SNMP ifIndex: 501, Generation: 2176
  Link-level type: Ethernet, MTU: 1518, Speed: 3Gbps, BPDU Error: None, MAC-REWRITE
 Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:1f:12:8c:af:c0, Hardware address: 00:1f:12:8c:af:c0
  Last flapped   : 2010-04-16 14:25:36 PDT (00:02:50 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   64                   0 bps
   Output bytes  :           9816525824            463779840 bps
   Input  packets:                    1                   0 pps
   Output packets:             38345804              226455 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
```


Copyright © 2016, Juniper Networks, Inc.

```
0, Resource errors: 0
 Output errors:
   Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
 Egress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort             38270790             38270790                    0

   1 expedited-fo                   0                    0                    0

   2 assured-forw                   0                    0                    0

   3 network-cont                 526                  526                    0

  Logical interface ae0.0 (Index 69) (SNMP ifIndex 502) (Generation 692)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.11 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes           bps
    Bundle:
        Input :             1          0            64            0
        Output:      38572259     226453    9874497884    463775744
    Link:
      ge-5/0/1.0
        Input :             0          0             0            0
        Output:      12743866      75484    3262429696    154591232
      ge-5/2/0.0
        Input :             1          0            64            0
        Output:      13043256      75484    3339073116    154591232
      ge-5/2/1.0
        Input :             0          0             0            0
        Output:      12785137      75485    3272995072    154593280
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-5/0/1.0                0           0            0            0
      ge-5/2/0.0                0           0            0            0
      ge-5/2/1.0                0           0            0            0
    Security: Zone: HOST
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp dhcp
    finger ftp tftp ident-reset http https ike netconf ping reverse-telnet
reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet
    traceroute xnm-clear-text xnm-ssl lsping ntp sip
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     0
      Multicast packets :               0
      Bytes permitted by policy :       0
      Connections established :         0
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       8976842784
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 0
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol inet, MTU: 1500, Generation: 841, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.1.1/24, Local: 10.1.1.2, Broadcast: 10.1.1.255, Generation:
422
    Protocol multiservice, MTU: Unlimited, Generation: 842, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
   Logical interface ae0.32767 (Index 83) (SNMP ifIndex 503) (Generation 693)
    Flags: SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      ge-5/0/1.32767
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-5/2/0.32767
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-5/2/1.32767
        Input :             0          0             0            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port    Port  Port

                              priority          identifier  priority  number   key

      ge-5/0/1.32767    Actor      127 00:1f:12:8c:af:c0       127     833     1

      ge-5/0/1.32767  Partner      127 00:1f:12:8f:d7:c0       127     641     1

      ge-5/2/0.32767    Actor      127 00:1f:12:8c:af:c0       127     848     1

      ge-5/2/0.32767  Partner      127 00:1f:12:8f:d7:c0       127     656     1

      ge-5/2/1.32767    Actor      127 00:1f:12:8c:af:c0       127     849     1

      ge-5/2/1.32767  Partner      127 00:1f:12:8f:d7:c0       127     657     1

    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-5/0/1.32767           342         511            0            0
      ge-5/2/0.32767           344         498            0            0
      ge-5/2/1.32767           344         500            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-5/0/1.32767             0           0            0            0
      ge-5/2/0.32767             0           0            0            0
      ge-5/2/1.32767             0           0            0            0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Security: Zone: HOST
    Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp dhcp
    finger ftp tftp ident-reset http https ike netconf ping reverse-telnet
reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet
    traceroute xnm-clear-text xnm-ssl lsping ntp sip
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     0
      Multicast packets :               0
      Bytes permitted by policy :       0
      Connections established :         0
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       0
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 0
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol multiservice, MTU: Unlimited, Generation: 843, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces (SRX Series)

Syntax

```
show interfaces {
  <brief | detail | extensive | terse>
  controller interface-name
  descriptions interface-name
  destination-class (all | destination-class-name logical-interface-name)
  diagnostics optics interface-name
  far-end-interval interface-fpc/pic/port
  filters interface-name
  flow-statistics interface-name
  interval interface-name
  load-balancing (detail | interface-name)
  mac-database mac-address mac-address
  mc-ae id identifier unit number revertive-info
  media interface-name
  policers interface-name
  queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics
  redundancy (detail | interface-name)
  routing brief detail summary interface-name
  routing-instance (all | instance-name)
  snmp-index snmp-index
  source-class (all | destination-class-name logical-interface-name)
  statistics interface-name
  switch-port switch-port number
  transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all |
    interface-name)
  zone interface-name
}
```

Release Information Command modified in Junos OS Release 9.5.

Description Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

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

Options

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim*/0/*port***—E1 interface.
  - **e3-*pim*/0/*port***—E3 interface.

Copyright © 2016, Juniper Networks, Inc.

  - **fe-*pim*/0/*port***—Fast Ethernet interface.
  - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
  - **se-*pim*/0/*port***—Serial interface.
  - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
  - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
  - **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

- **brief | detail | extensive | terse**—(Optional) Display the specified level of output.
- **controller**—(Optional) Show controller information.
- **descriptions**—(Optional) Display interface description strings.
- **destination-class**—(Optional) Show statistics for destination class.
- **diagnostics**—(Optional) Show interface diagnostics information.
- **far-end-interval**—(Optional) Show far end interval statistics.
- **filters**—(Optional) Show interface filters information.
- **flow-statistics**—(Optional) Show security flow counters and errors.
- **interval**—(Optional) Show interval statistics.
- **load-balancing**—(Optional) Show load-balancing status.
- **mac-database**—(Optional) Show media access control database information.
- **mc-ae**—(Optional) Show MC-AE configured interface information.
- **media**—(Optional) Display media information.
- **policers**—(Optional) Show interface policers information.
- **queue**—(Optional) Show queue statistics for this interface.
- **redundancy**—(Optional) Show redundancy status.
- **routing**—(Optional) Show routing status.
- **routing-instance**—(Optional) Name of routing instance.
- **snmp-index**—(Optional) SNMP index of interface.
- **source-class**—(Optional) Show statistics for source class.
- **statistics**—(Optional) Display statistics and detailed output.
- **switch-port**—(Optional) Front end port number (0..15).
- **transport**—(Optional) Show interface transport information.
- **zone**—(Optional) Interface's zone.

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

List of Sample Output

show interfaces Gigabit Ethernet on page 1995
show interfaces brief (Gigabit Ethernet) on page 1996
show interfaces detail (Gigabit Ethernet) on page 1996
show interfaces extensive (Gigabit Ethernet) on page 1998
show interfaces terse on page 2001
show interfaces controller (Channelized E1 IQ with Logical E1) on page 2001
show interfaces controller (Channelized E1 IQ with Logical DS0) on page 2001
show interfaces descriptions on page 2002
show interfaces destination-class all on page 2002
show interfaces diagnostics optics on page 2002
show interfaces far-end-interval coc12-5/2/0 on page 2003
show interfaces far-end-interval coc1-5/2/1:1 on page 2003
show interfaces filters on page 2004
show interfaces flow-statistics (Gigabit Ethernet) on page 2004
show interfaces interval (Channelized OC12) on page 2005
show interfaces interval (E3) on page 2005
show interfaces interval (SONET/SDH) on page 2006
show interfaces load-balancing on page 2006
show interfaces load-balancing detail on page 2006
show interfaces mac-database (All MAC Addresses on a Port) on page 2007
show interfaces mac-database (All MAC Addresses on a Service) on page 2007
show interfaces mac-database mac-address on page 2008
show interfaces mc-ae on page 2008
show interfaces media (SONET/SDH) on page 2008
show interfaces policers on page 2009
show interfaces policers interface-name on page 2009
show interfaces queue on page 2009
show interfaces redundancy on page 2010
show interfaces redundancy (Aggregated Ethernet) on page 2010
show interfaces redundancy detail on page 2011
show interfaces routing brief on page 2011
show interfaces routing detail on page 2011
show interfaces routing-instance all on page 2012
show interfaces snmp-index on page 2012
show interfaces source-class all on page 2012
show interfaces statistics (Fast Ethernet) on page 2013
show interfaces switch-port on page 2013
show interfaces transport pm on page 2014
show security zones on page 2015

Output Fields

Table 194 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—There are two possible conditions regarding the number of oversized frames: ◦ Packet length exceeds 1518 octets, or ◦ Packet length exceeds MRU • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 194: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
```

Copyright © 2016, Juniper Networks, Inc.

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

Sample Output

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
```


Copyright © 2016, Juniper Networks, Inc.

```
Disabled,
 Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
 Device flags   : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags     : None
 CoS queues     : 8 supported, 8 maximum usable queues
 Hold-times     : Up 0 ms, Down 0 ms
 Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
 Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
 Statistics last cleared: Never
 Traffic statistics:
  Input  bytes  :                   0                   0 bps
  Output bytes  :                   0                   0 bps
  Input  packets:                   0                   0 pps
  Output packets:                   0                   0 pps
 Egress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort                   0                   0                   0

   1 expedited-fo                  0                   0                   0

   2 assured-forw                  0                   0                   0

   3 network-cont                  0                   0                   0

 Queue number:         Mapped forwarding classes
   0                   best-effort
   1                   expedited-forwarding
   2                   assured-forwarding
   3                   network-control
 Active alarms  : LINK
 Active defects : LINK
 Interface transmit statistics: Disabled

 Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
   Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Transit statistics:
    Input  bytes  :                   0                   0 bps
    Output bytes  :                   0                   0 bps
    Input  packets:                   0                   0 pps
    Output packets:                   0                   0 pps
   Security: Zone: public
   Flow Statistics :
   Flow Input statistics :
     Self packets :                    0
     ICMP packets :                    0
     VPN packets :                     0
     Multicast packets :               0
     Bytes permitted by policy :       0
     Connections established :         0
```

```
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

Sample Output

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
```


Copyright © 2016, Juniper Networks, Inc.

```
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %           bps     %           usec
    0 best-effort            95     950000000    95              0      low
none
    3 network-control         5      50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
```

```
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Local statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Transit statistics:
 Input  bytes  :                   0                   0 bps
 Output bytes  :                   0                   0 bps
 Input  packets:                   0                   0 pps
 Output packets:                   0                   0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                    0
  ICMP packets :                    0
  VPN packets :                     0
  Multicast packets :               0
  Bytes permitted by policy :       0
  Connections established :         0
Flow Output statistics:
  Multicast packets :               0
  Bytes permitted by policy :       0
Flow error statistics (Packets dropped due to):
  Address spoofing:                 0
  Authentication failed:            0
  Incoming NAT errors:              0
  Invalid zone received packet:     0
  Multiple user authentications:    0
  Multiple incoming NAT:            0
  No parent for a gate:             0
  No one interested in self packets: 0
  No minor session:                 0
  No more sessions:                 0
  No NAT gate:                      0
  No route present:                 0
  No SA for incoming SPI:           0
  No tunnel found:                  0
  No session for a gate:            0
  No zone or NULL zone binding      0
  Policy denied:                    0
  Security association not active:  0
  TCP sequence number out of window: 0
  Syn-attack protection:            0
  User authentication errors:       0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
    Generation: 150
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1              --> 0/0
                                                 10.0.0.16             --> 0/0
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

Sample Output

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                      Admin Link
ce1-1/2/6                                                       up    up
e1-1/2/6                                                        up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                      Admin Link
ce1-1/2/3                                                       up    up

ds-1/2/3:1                                                      up    up

ds-1/2/3:2                                                      up    up
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

Sample Output

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                                Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                        0                    0
                                  (                  0) (                  0)
                       silver                        0                    0
                                  (                  0) (                  0)
  Logical interface so-0/1/3.0
                                                Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                        0                    0
                                  (                  0) (                  0)
                       silver                        0                    0
                                  (                  0) (                  0)
```

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
```


Copyright © 2016, Juniper Networks, Inc.

```
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

Sample Output

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
```


Copyright © 2016, Juniper Networks, Inc.

```
    Bytes permitted by policy :        3478
    Connections established :          1
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        16994
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
```

```
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State           Last change  Member count
ams0       Up              1d 00:50     2
ams1       Up              00:00:59     2
```

show interfaces load-balancing detail

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
```

Copyright © 2016, Juniper Networks, Inc.

```
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
```


Copyright © 2016, Juniper Networks, Inc.

```
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues      : 8 supported
  Last flapped    : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  SONET alarms    : None
  SONET defects   : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

Copyright © 2016, Juniper Networks, Inc.

```
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :            118750000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :              6250000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
```

Sample Output

show interfaces redundancy

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present               sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

show interfaces redundancy (Aggregated Ethernet)

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up

ae0
ae1
```


Copyright © 2016, Juniper Networks, Inc.

```
ae2
ae3
ae4
```

show interfaces redundancy detail

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

Sample Output

show interfaces routing brief

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
```

Copyright © 2016, Juniper Networks, Inc.

```
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1    up      up     inet     10.0.0.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0  up      up     inet     10.0.0.1/32
```

Sample Output

show interfaces snmp-index

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

Sample Output

show interfaces source-class all

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                                Packets                Bytes
      Source class                  (packet-per-second)    (bits-per-second)
                             gold              1928095            161959980
                                  (                889) (            597762)
                           bronze                    0                    0
```

 Copyright © 2016, Juniper Networks, Inc.

```
                                        (                  0) (                  0)
                                silver                     0                     0
                                        (                  0) (                  0)
  Logical interface so-0/1/3.0
                                                        Packets                 Bytes
      Source class                       (packet-per-second)    (bits-per-second)
                                  gold                     0                     0
                                        (                  0) (                  0)
                                bronze                     0                     0
                                        (                  0) (                  0)
                                silver                116113               9753492
                                        (                939) (             631616)
```

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                                        Packets                 Bytes
      Destination class                  (packet-per-second)    (bits-per-second)
                              silver1                    0                     0
                                        (                  0) (                  0)
                              silver2                    0                     0
                                        (                  0) (                  0)
                              silver3                    0                     0
                                        (                  0) (                  0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

Sample Output

show interfaces switch-port

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                          Receive         Transmit
          Total bytes                       28437086         21792250
```

Copyright © 2016, Juniper Networks, Inc.

```
          Total packets                      409145          88008
          Unicast packets                      9987          83817
          Multicast packets                  145002              0
          Broadcast packets                  254156           4191
          Multiple collisions                    23             10
          FIFO/CRC/Align errors                   0              0
          MAC pause frames                        0              0
          Oversized frames                        0
          Runt frames                             0
          Jabber frames                           0
          Fragment frames                         0
          Discarded frames                        0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

Sample Output

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr         2008544300       0                NA              NA
  FEC-UncorrectedWords     0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
```


Copyright © 2016, Juniper Networks, Inc.

```
    PM                  MIN       MAX       AVG       THRESHOLD       TCA-ENABLED
     TCA-RAISED
  BER                   3.6e-5    5.8e-5    3.6e-5     10.0e-3        No
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                           CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                                 (MIN)
(MAX)    (MIN) (MAX)       (MIN) (MAX)
  Lane chromatic dispersion         0        0         0         0         0
 0        NA    NA         NA     NA
  Lane differential group delay     0        0         0         0         0
 0        NA    NA         NA     NA
    q Value                            120      120       120       120       0
   0        NA    NA         NA     NA
  SNR                               28       28        29        28        0
 0        NA    NA         NA     NA
  Tx output power(0.01dBm)          -5000    -5000     -5000     -5000     -300
 -100     No    No         No     No
  Rx input power(0.01dBm)           -3642    -3665     -3626     -3637     -1800
 -500     No    No         No     No
  Module temperature(Celsius)       46       46        46        46        -5
 75       No    No         No     No
  Tx laser bias current(0.1mA)      0        0         0         0         0
 0        NA    NA         NA     NA
  Rx laser bias current(0.1mA)      1270     1270      1270      1270      0
 0        NA    NA         NA     NA
  Carrier frequency offset(MHz)     -186     -186      -186      -186      -5000
 5000     No    No         No     No
```

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
```

Copyright © 2016, Juniper Networks, Inc.

```
Interfaces bound: 1
Interfaces:
  ge-0/0/2.0
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces diagnostics optics

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP) installed in SRX Series Services Gateways. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options *interface-name*—Name of the interface associated with the port in which the transceiver is installed: **ge-*fpc/pic/port*** .

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

List of Sample Output show interfaces diagnostics optics on page 2020

Output Fields Table 195 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 195: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is **On** or **Off**.

Copyright © 2016, Juniper Networks, Inc.

Table 195: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is **On** or **Off**.
Laser bias current high warning	Displays whether the laser bias power setting high warning is **On** or **Off**.
Laser bias current low warning	Displays whether the laser bias power setting low warning is **On** or **Off**.
Laser output power high alarm	Displays whether the laser output power high alarm is **On** or **Off**.
Laser output power low alarm	Displays whether the laser output power low alarm is **On** or **Off**.
Laser output power high warning	Displays whether the laser output power high warning is **On** or **Off**.
Laser output power low warning	Displays whether the laser output power low warning is **On** or **Off**.
Module temperature high alarm	Displays whether the module temperature high alarm is **On** or **Off**.
Module temperature low alarm	Displays whether the module temperature low alarm is **On** or **Off**.
Module temperature high warning	Displays whether the module temperature high warning is **On** or **Off**.
Module temperature low warning	Displays whether the module temperature low warning is **On** or **Off**.
Module voltage high alarm	Displays whether the module voltage high alarm is **On** or **Off**.
Module voltage low alarm	Displays whether the module voltage low alarm is **On** or **Off**.
Module voltage high warning	Displays whether the module voltage high warning is **On** or **Off**.
Module voltage low warning	Displays whether the module voltage low warning is **On** or **Off**.
Laser rx power high alarm	Displays whether the receive laser power high alarm is **On** or **Off**.
Laser rx power low alarm	Displays whether the receive laser power low alarm is **On** or **Off**.

Copyright © 2016, Juniper Networks, Inc.

Table 195: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Laser rx power high warning	Displays whether the receive laser power high warning is **On** or **Off**.
Laser rx power low warning	Displays whether the receive laser power low warning is **On** or **Off**.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.

Copyright © 2016, Juniper Networks, Inc.

Table 195: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                         :  7.408 mA
    Laser output power                         :  0.3500 mW / -4.56 dBm
    Module temperature                         :  23 degrees C / 73 degrees F
    Module voltage                             :  3.3450 V
    Receiver signal average optical power      :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Laser output power high warning            :  Off
    Laser output power low warning             :  Off
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
    Module voltage low warning                 :  Off
    Laser rx power high alarm                  :  Off
    Laser rx power low alarm                   :  On
    Laser rx power high warning                :  Off
    Laser rx power low warning                 :  On
    Laser bias current high alarm threshold    :  17.000 mA
    Laser bias current low alarm threshold     :  1.000 mA
    Laser bias current high warning threshold  :  14.000 mA
    Laser bias current low warning threshold   :  2.000 mA
    Laser output power high alarm threshold    :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold     :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold  :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold   :  0.0790 mW / -11.02 dBm
```


Copyright © 2016, Juniper Networks, Inc.

```
Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
Module temperature high warning threshold :  90 degrees C / 194 degrees F
Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
Module voltage high alarm threshold       :  3.900 V
Module voltage low alarm threshold        :  2.700 V
Module voltage high warning threshold     :  3.700 V
Module voltage low warning threshold      :  2.900 V
Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces flow-statistics

Syntax **show interfaces flow-statistics** *<interface-name>*

Release Information Command introduced in Junos OS Release 9.2.

Description Display interfaces flow statistics.

Options *Interface-name*—(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the "Interface Naming Conventions" on page 2446.

- **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
- **br-*pim*/0/*port***—Basic Rate Interface for establishing ISDN connections.
- **ce1-*pim*/0/*port***—Channelized E1 interface.
- **ct1-*pim*/0/*port***—Channelized T1 interface.
- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
- **e1-*pim*/0/*port***—E1 interface.
- **e3-*pim*/0/*port***—E3 interface.
- **fe-*pim*/0/ *port***—Fast Ethernet interface.
- **ge-*pim*/0/*port***—Gigabit Ethernet interface.
- **se-*pim*/0/*port***—Serial interface.
- **t1-*pim*/0/*port***—T1 (also called DS1) interface.
- **t3-*pim*/0/ *port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

List of Sample Output **show interfaces flow-statistics (Gigabit Ethernet) on page 2025**

Output Fields Table 196 lists the output fields for the **show interfaces flow-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 196: show interfaces flow-statistics Output Fields

Field Name	Field Description
Traffic statistics	Number of packets and bytes transmitted and received on the physical interface.

 Copyright © 2016, Juniper Networks, Inc.

Table 196: show interfaces flow-statistics Output Fields *(continued)*

Field Name	Field Description
Local statistics	Number of packets and bytes transmitted and received on the physical interface.
Transit statistics	Number of packets and bytes transiting the physical interface.
Flow input statistics	Statistics on packets received by flow module.
Flow output statistics	Statistics on packets sent by flow module.
Flow error statistics	Packet drop statistics for the flow module. For further details, see Table 197.

Table 197: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.
Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.
User authentication errors	Packet was dropped because policy requires authentication; however: • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.

Copyright © 2016, Juniper Networks, Inc.

Table 197: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Flow:	
No one interested in self packets	This counter is incremented for one of the following reasons: • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	The packet was dropped because a valid route was not available to forward the packet. For new sessions, the counter is incremented for one of the following reasons: • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because: • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found
No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.
Policy denied	The error counter is incremented for one of the following reasons: • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.
TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-

Copyright © 2016, Juniper Networks, Inc.

Table 197: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Invalid zone received packet	-
No NAT gate	-

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
      Bytes permitted by policy :       3478
      Connections established :         1
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 0
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
```

Copyright © 2016, Juniper Networks, Inc.

```
Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces statistics (View)

Syntax show interfaces statistics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level view

Related Documentation
- Understanding Interfaces on page 2441

List of Sample Output show interfaces statistics on page 2027

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
     Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
     Input packets : 2743333
     Output packets: 6790470992
     Security: Zone: untrust
     Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset
 http https ike netconf ping reverse-telnet
    reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl lsping ntp sip
     Protocol inet, MTU: 9192
       Addresses, Flags: Is-Preferred Is-Primary
         Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces swfabx

Syntax show interfaces <swfab0 | swfab1>

Release Information Command introduced in Junos OS Release 11.1.

Description List the configured interfaces for each swfab interface. The swfab interface can contain one or more members because it is an aggregated interface.

Required Privilege Level view

Related Documentation
- clear interfaces statistics swfabx on page 2763

List of Sample Output show interfaces <swfab0 | swfab1> on page 2028

Output Fields Table 198 lists the output fields for the **show interfaces <swfab0 | swfab1>** command. Output fields are listed in the approximate order in which they appear.

Table 198: show interfaces <swfab0 | swfab1> Output Fields

Field Name	Field Description
fabric-options	The fabric-options hierarchy is made to be in sync with the fab interfaces.
member-interfaces	Interfaces specified under member-interfaces are single logical aggregate-Ethernet interfaces. This interface carries internode switching traffic.

Sample Output

show interfaces <swfab0 | swfab1>

```
user@host> show interfaces <swfab0 | swfab1>
fabric-options {
       member-interfaces {
           ge-0/0/6;
           ge-0/0/7;
       }
  }
```


Copyright © 2016, Juniper Networks, Inc.

show monitor security flow

Syntax show monitor security flow

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Display information about the security flow session monitoring.

Required Privilege Level view

Related Documentation
- Monitoring Security Flow Sessions Overview on page 1781
- monitor security flow filter on page 1924
- monitor security flow start on page 1926
- clear monitor security flow filter on page 1899

List of Sample Output show monitor security flow on page 2029

Output Fields Lists the output fields for the **show monitor security flow** command. Output fields are listed in the approximate order in which they appear.

Table 199: show monitor security flow Output Fields

Field Name	Field Description
Monitor security flow session status	State of the security flow session monitoring: **active** or **inactive**.
Monitor security flow trace file	Name of the file for monitoring output.
Monitor security flow filters	• **Destination Address**—Address of the destination to be matched. • **Destination Port**—Name of the destination port to be matched. • **Interface Name**—Interface name to be matched. • **Logical System Name**—Logical system name to be matched. • **Name**—Name of the security flow filter. • **Protocol**—Name of the protocol to be matched. • **Source Address**—Address of the source to be matched. • **Source Port**—Name of the source port to be matched. • **Status**—State of the security flow filter: **active** or **inactive**.

Sample Output

show monitor security flow

```
user@host>show monitor security flow

Monitor security flow session status: Active
Monitor security flow trace file: flow
Monitor security flow filters:
  Name: server-sql
```

Copyright © 2016, Juniper Networks, Inc.

```
  Status: Active
  source: 10.2.2.1 (port *), destination: 10.20.30.40 (port 1433)
  protocol: TCP
Name: internet-access
  Status: Active
  source: * (port *), destination: * (port 80)
  protocol: TCP
```


Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session

Syntax **show security flow cp-session**
[<filter>] [summary | terse]

Release Information Command introduced in Junos OS Release 10.2. Support for connection tag added in Junos OS Release 15.1X49-D40.

Description Display central point session-related flow information.

Options
- destination-port—Destination port
- destination-prefix—Destination prefix
- family—Display session by family.
- protocol—IP protocol number
- source-port—Source port
- source-prefix—Source IP prefix or address
- summary | terse–Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow cp-session on page 2032
show security flow cp-session summary on page 2032
show security flow cp-session terse on page 2033

Output Fields Table 200 lists the output fields for the **show security flow cp-session** command. Output fields are listed in the approximate order in which they appear.

Table 200: show security flow cp-session Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Maximum sessions	Number of maximum central point sessions.
Maximum inet6 sessions	Number of maximum inet6 central point sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 200: show security flow cp-session Output Fields *(continued)*

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session

```
root> show security flow cp-session
DCP Flow Sessions on FPC0 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC0 PIC1:

Session ID: 10320276, SPU: 1, Valid
   In: 203.0.113.1/1000 --> o 203.0.113.2/2000;udp, Conn Tag: 0x0,
  Out: 0.0.0.0/0 --> 0.0.0.0/0;0, Conn Tag: 0x0,
Total sessions: 1
```

Sample Output

show security flow cp-session summary

```
root> show security flow cp-session summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 2
```


Copyright © 2016, Juniper Networks, Inc.

```
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
Maximum sessions: 7549747
Maximum inet6 sessions: 7549747
```

show security flow cp-session terse

```
root> show security flow cp-session terse
DCP Flow Sessions on FPC0 PIC1:

Session ID: 10000038, SPU: 1, Valid
  In: 203.0.113.6/1 --> 198.51.100.13/1;pim, Conn Tag: 0x0,
  Out: 198.51.100.13/1 --> 203.0.113.6/1;pim, Conn Tag: 0x0
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session destination-port

Syntax **show security flow cp-session destination-port** *destination-port-number* **[summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session-related flow information for the specified destination port.

Options

- *destination-port-number* —Number of the destination port for which to display central point session information.

  Range: 1 through 65,535

- summary | terse—Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow cp-session on page 2031
- show security flow cp-session destination-prefix on page 2036

List of Sample Output

show security flow cp-session destination-port summary on page 2035
show security flow cp-session destination-port terse on page 2035

Output Fields Table 201 lists the output fields for the **show security flow cp-session destination-port** command. Output fields are listed in the approximate order in which they appear.

Table 201: show security flow cp-session destination-port Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security flow cp-session destination-port summary

```
root> show security flow cp-session destination-port 21 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session destination-port terse

```
root> show security flow cp-session destination-port 21 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Session ID: 410003298, SPU: 41, Valid
  In: 203.0.113.10/26182 --> 198.51.100.1/21;tcp,
   Out: 198.51.100.1/21 --> 203.0.113.10/26182;tcp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC2:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session destination-prefix

Syntax **show security flow cp-session destination-prefix** *destination-IP-prefix* **[summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session-related flow information for the specified destination prefix.

Options

- *destination-IP-prefix* —Destination IP prefix or address for which to display central point session information.

  Range: 1 through 65,535.

- summary | terse—Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow cp-session on page 2031
- show security flow cp-session destination-port on page 2034

List of Sample Output

show security flow cp-session destination-prefix summary on page 2037
show security flow cp-session destination-prefix terse on page 2037

Output Fields Table 202 lists the output fields for the **show security flow cp-session destination-prefix** command. Output fields are listed in the approximate order in which they appear.

Table 202: show security flow cp-session destination-prefix Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security flow cp-session destination-prefix summary

```
root> show security flow cp-session destination-prefix 60/8 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session destination-prefix terse

```
root> show security flow cp-session destination-prefix 60/8 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002660, SPU: 42, Valid
  In: 203.0.113.10/26183 --> 192.0.2.1/21;tcp,
  Out:192.0.2.1/21 --> 203.0.113.10/26183;tcp,
Total sessions: 1


DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session family

Syntax **show security flow cp-session family *family* [summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session-related flow information for the specified family.

Options
- *family*–Display session by family.
- inet—Display IPv4 sessions.
- inet6–Display IPv6 and IPv6-NATPT sessions.
- summary | terse–Display the specified level of output.

Required Privilege Level view

Related Documentation
- show security flow cp-session on page 2031

List of Sample Output
show security flow cp-session family summary on page 2039
show security flow cp-session family terse on page 2039

Output Fields Table 203 lists the output fields for the **show security flow cp-session** family command. Output fields are listed in the approximate order in which they appear.

Table 203: show security flow cp-session family Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security flow cp-session family summary

```
root> show security flow cp-session family inet summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session family terse

```
root> show security flow cp-session family inet terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002660, SPU: 42, Valid
 Session ID: 420002660, SPU: 42, Valid
  In: 198.51.100.1/26183 --> 203.0.113.2/21;tcp,
  Out: 203.0.113.2/21 --> 198.51.100.1/26183;tcp,

Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session protocol

Syntax **show security flow cp-session protocol *protocol-name* [summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session-related flow information for the specified protocol.

Options

- *protocol-name* —Protocol to use as a central point session filter. Information about the central point session that uses this protocol is displayed.

  Possible protocols are:

  - ah—IP Security Authentication Header
  - egp—Exterior gateway protocol
  - esp—IPsec Encapsulating Security Payload
  - gre—Generic routing encapsulation
  - icmp—Internet Control Message Protocol
  - icmp6—Internet Control Message Protocol
  - igmp—Internet Group Management Protocol
  - ipip—IP over IP
  - ospf—Open Shortest Path First
  - pim—Protocol Independent Multicast
  - rsvp—Resource Reservation Protocol
  - sctp—Stream Control Transmission Protocol
  - tcp—Transmission Control Protocol
  - udp—User Datagram Protocol

- summary | terse–Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow cp-session on page 2031

List of Sample Output

show security flow cp-session protocol summary on page 2041
show security flow cp-session protocol terse on page 2042

Output Fields Table 204 lists the output fields for the **show security flow cp-session** protocol command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 204: show security flow cp-session protocol Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Sample Output

show security flow cp-session protocol summary

```
root> show security flow cp-session protocol tcp summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Sessions in other states: 0
Total sessions: 1
```

show security flow cp-session protocol terse

```
root> show security flow cp-session protocol tcp terse
Session ID: 160000015, SPU: 17, Valid
  In: 203.0.113.9/32838 --> 198.51.100.26/21;tcp,
  Out: 198.51.100.26/21 --> 203.0.113.2/32838;tcp,
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session source-port

Syntax **show security flow cp-session source-port** *source-port-number* **[summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session-related flow information for the specified source-port.

Options *source-port-number*—Number of the source port about which to display central point session information.

Range: 1 through 65,535

summary | terse—Display the specified level of output.

Required Privilege Level view

Related Documentation
- show security flow cp-session on page 2031
- show security flow cp-session source-prefix on page 2045

List of Sample Output
show security flow cp-session source-port summary on page 2044
show security flow cp-session source-port terse on page 2044

Output Fields Table 205 lists the output fields for the **show security flow cp-session source-port** command. Output fields are listed in the approximate order in which they appear.

Table 205: show security flow cp-session source-port Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security flow cp-session source-port summary

```
root> show security flow cp-session source-port 7000 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session source-port terse

```
root> show security flow cp-session source-port 7000 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002661, SPU: 42, Valid
  In: 203.0.113.64/7000 --> 192.0.2.9/8000;udp,
  Out: 192.0.2.9/8000 --> 203.0.113.64/7000;udp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow cp-session source-prefix

Syntax **show security flow cp-session source-prefix** *source-IP-prefix* **[summary | terse]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display central point session related flow information for the specified source-prefix.

Options
- *source-IP-prefix*—Source IP prefix or address for which to display central point session information.
- summary | terse—Display the specified level of output.

Required Privilege Level view

Related Documentation
- show security flow cp-session on page 2031
- show security flow cp-session source-port on page 2043

List of Sample Output
show security flow cp-session source-prefix summary on page 2046
show security flow cp-session source-prefix terse on page 2046

Output Fields Table 206 lists the output fields for the **show security flow cp-session source-prefix** command. Output fields are listed in the approximate order in which they appear.

Table 206: show security flow cp-session source-prefix Output Fields

Field Name	Field Description
Valid gates	Number of valid central point sessions.
Pending gates	Number of pending central point sessions.
Invalidated gates	Number of invalid central point sessions.
Gates in other states	Number of central point sessions in other states.
Total gates	Number of central point sessions in total.
Session ID	Number that identifies the session. Use this ID to get more information about the session.
SPU	Services Processing Unit.
In	Incoming flow (source and destination IP addresses).
Out	Reverse flow (source and destination IP addresses).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security flow cp-session source-prefix summary

```
root> show security flow cp-session source-prefix 200/8 summary
DCP Flow Sessions on FPC10 PIC0:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow cp-session source-prefix terse

```
root> show security flow cp-session source-prefix 200/8 terse
DCP Flow Sessions on FPC10 PIC0:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC1:
Total sessions: 0

DCP Flow Sessions on FPC10 PIC2:

Session ID: 420002663, SPU: 42, Valid
  In: 203.0.113.10/7000 --> 198.51.100.2/8000;udp,
  Out: 198.51.100.2/8000 --> 203.0.113.10/7000;udp,
Total sessions: 1

DCP Flow Sessions on FPC10 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate

Syntax

show security flow gate
[<filter>] [brief | summary}

Release Information

Command introduced in Junos OS Release 8.5; Filter and display options added in Junos OS Release 10.2.

Description

Display information about temporary openings known as pinholes or gates in the security firewall.

Pinholes are used by applications that commonly have both control and data sessions and must create openings in the firewall for the data sessions based on information from the parent sessions.

Options

- destination-port—Destination port
- destination-prefix—Destination IP prefix or address
- protocol—IP protocol number
- source-port—Source port
- source-prefix—Source IP prefix or address
- brief | summary—Display the specified level of output.

Required Privilege Level

view

Related Documentation

- show security flow gate brief node on page 2059
- show security flow gate destination-port on page 2065
- show security flow gate destination-prefix on page 2068
- show security flow gate protocol on page 2071
- show security flow gate summary node on page 2074

List of Sample Output

show security flow gate on page 2048
show security flow gate brief on page 2049
show security flow gate summary on page 2050

Output Fields

Table 25 lists the output fields for the **show security flow gate** command. Output fields are listed in the approximate order in which they appear.

Table 207: show security flow gate Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 207: show security flow gate Output Fields *(continued)*

Field Name	Field Description
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Maximum gates	Number of maximum gates

Sample Output

show security flow gate

```
user@host> show security flow gate
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.198.51.100.252/64515-64515
Translated: 0.0.0.0/0->10.0.31.161/25415
Protocol: udp
Application: none/0
Age: 101 seconds
Flags: 0xe001
Zone: untrust
Reference count: 1
Resource: 5-1024-8185
ole: 0.0.0.0-0.0.0.0/0-198.51.100.252/1046-1046
Translated: 198.51.100.252/36039-> 203.0.113.1/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
```


Copyright © 2016, Juniper Networks, Inc.

```
Flags: 0xe200
Zone: untrust
Reference count: 1
Resource: 5-1024-8189
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/24101-24101
Translated: 0.0.0.0/0-> 198.51.100.252/24101
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
Zone: trust
Reference count: 1
Resource: 5-1024-8188
Hole: 0.0.0.0-0.0.0.0/0-0->40.1.1.5-198.51.100.252/24100-24100
Translated: 0.0.0.0/0->198.51.100.252/24100
Protocol: udp
Application: none/0
Age: 93 seconds
Flags: 0xe001
Zone: trust
Reference count: 1
Resource: 5-1024-8191
Hole: 0.0.0.0-0.0.0.0/0-0->198.51.100.252-198.51.100.252/5060-5060
Translated: 0.0.0.0/0->198.51.100.252/5060
Protocol: udp
Application: junos-sip/63
Age: 65535 seconds
Flags: 0xe200
Zone: trust
Reference count: 1
Resource: 5-1024-8190
```

show security flow gate brief

```
root> show security flow gate brief
Flow Gates on FPC4 PIC1:

Hole: 192.0.2.1-192.0.2.1/0->192.0.2.100-192.0.2.100/38143-38143
  Translated: 192.0.2.1->192.0.2.100/38143
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65532 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate summary

```
root> show security flow gate summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```


Copyright © 2016, Juniper Networks, Inc.

show security flow ip-action

Syntax **show security flow ip-action** [<filter>] [summary family (inet | inet6)]

Release Information Command introduced in Junos OS Release 10.1. Logical systems option added in Junos OS Release 11.2 . Summary option introduced in Junos OS Release 12.1.

Description Display the current IP-action settings, based on filtered options, for IP sessions running on the device.

Options

- *filter*—Filter the display based on the specified criteria.

  The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

  all | [*filter*]—All active sessions on the device.

  destination-port ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535.

  destination-prefix ***destination-prefix***—Destination IP prefix or address.

  family (inet | inet6) [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

  logical-system ***logical-system-name*** **| all** [*filter*]—Specified logical system or all logical systems.

  protocol ***protocol-name*** **|** ***protocol-number*** [*filter*]—Protocol name or number and filtered options.

  - **ah** or **51**
  - **egp** or **8**
  - **esp** or **50**
  - **gre** or **47**
  - **icmp** or **1**
  - **icmp6** or **58**
  - **igmp** or **2**
  - **ipip** or **4**
  - **ospf** or **89**
  - **pim** or **103**
  - **rsvp** or **46**
  - **sctp** or **132**
  - **tcp** or **6**
  - **udp** or **17**

  root-logical-system [*filter*]—Default logical system information and filtered options.

Copyright © 2016, Juniper Networks, Inc.

source-port *source-port*—Source port number of the traffic. Range is 1 through 65,535.

source-prefix *source-prefix*—Source IP prefix or address of the traffic.

- **summary** —Summary information about IP-action entries.

  family—Display summary of IP-action entries by family. This option is used to filter the output.

  - **inet**—Display summary of IPv4 entries.
  - **inet6**—Display summary of IPv6 entries.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow ip-action on page 1900
- clear security flow session destination-port on page 1905
- clear security flow ip-action on page 1900

List of Sample Output

show security flow ip-action on page 2053
show security flow ip-action destination-port on page 2054
show security flow ip-action destination-prefix on page 2055
show security flow ip-action family inet protocol on page 2055
show security flow ip-action family inet logical-system all on page 2056
show security flow ip-action source-prefix on page 2057
show security flow ip-action summary on page 2058
show security flow ip-action summary family inet on page 2058
show security flow ip-action summary family inet6 on page 2058

Output Fields Table 208 lists the output fields for the **show security flow ip-action** command. Output fields are listed in the approximate order in which they appear.

Table 208: show security flow ip-action Output Fields

Field Name	Field Description
Src-Addr	Source address of outbound IP traffic.
Src-Port	Source port number of outbound IP traffic.
Dst-Addr	Destination address of inbound IP traffic.
Dst-Port/Proto	Destination port number and protocol type of inbound IP traffic.
Timeout (sec)	Configured timeouts and time remaining for an IP session.
Zone	Security zone associated with an IP session.
Action	Configured action type, for example, block, close, and notify.

Copyright © 2016, Juniper Networks, Inc.

Table 208: show security flow ip-action Output Fields *(continued)*

Field Name	Field Description
State	The active mode and passive mode describe the states of the **ip-action** entry.
IPv4 action count	The total number of IPv4 entries.
IPv6 action count	The total number of IPv6 entries.

Sample Output

show security flow ip-action

```
user@host> show security flow ip-action
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
 203.0.113.1       *       203.0.113.4         21/tcp         293/300      *
      close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *       203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300       *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *    203.0.113.4          21/tcp         292/300      *
   close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
 203.0.113.1         *    203.0.113.4          21/tcp         292/300      *
   close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
```

Copyright © 2016, Juniper Networks, Inc.

```
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action destination-port

```
user@host> show security flow ip-action destination-port 21

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
 close         Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4        21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1        *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         273/300      *
  close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```


Copyright © 2016, Juniper Networks, Inc.

show security flow ip-action destination-prefix

```
user@host>  show security flow ip-action destination-prefix 203.0.113.4/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
192.0.2.3           *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
192.0.2.3           *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet protocol

```
user@host>  show security flow ip-action family inet protocoludp
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
        close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2
```

Copyright © 2016, Juniper Networks, Inc.

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
     close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
     close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1        *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet logical-system all

```
user@host> show security flow ip-action family inet logical-system all

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
        close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
 close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *       203.0.113.4         69/udp         267/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
  close        Active            root-logical-system
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
```


Copyright © 2016, Juniper Networks, Inc.

```
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         266/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         266/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action source-prefix

```
user@host> show security flow ip-action source-prefix 192.0.2.3/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
   close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4        69/udp         244/300      *
    close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
     close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
     close        Passiveo
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
```

Copyright © 2016, Juniper Networks, Inc.

```
    close        Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary

```
user@host> show security flow ip-action summary

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action summary family inet

```
user@host> show security flow ip-action summary inet

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary family inet6

```
user@host> show security flow ip-action summary family inet6

 IPv6 action count: 1 on FPC0.PIC1
 IPv6 action count: 1 on FPC0.PIC2
 IPv6 action count: 1 on FPC0.PIC3
 IPv6 action count: 1 on FPC1.PIC0
 IPv6 action count: 1 on FPC1.PIC1
 IPv6 action count: 1 on FPC1.PIC2
 IPv6 action count: 1 on FPC1.PIC3
 IPv6 action count: Active mode 1 on all PICs
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate brief node

Syntax **show security flow gate brief node (*node-id* | all | local | primary)**

Release Information Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

Description Display information about temporary openings known as pinholes or gates in the security firewall for the specified node options in brief mode.

Options **node**—(Optional) For chassis cluster configurations, display gate information on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- show security flow gate on page 449
- show security flow gate summary node on page 2074

List of Sample Output
show security flow gate brief node 0 on page 2060
show security flow gate brief node 1 on page 2061
show security flow gate brief node all on page 2061
show security flow gate brief node local on page 2063
show security flow gate brief node primary on page 2063

Output Fields Table 209 lists the output fields for the **show security flow gate brief node** command. Output fields are listed in the approximate order in which they appear.

Table 209: show security flow gate brief node Output Fields

Field Name	Field Description
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Hole	Range of flows permitted by the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 209: show security flow gate brief node Output Fields *(continued)*

Field Name	Field Description
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.

Sample Output

show security flow gate brief node 0

```
root@antbert> show security flow gate brief node 0
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65518 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node 1

```
root@antbert> show security flow gate brief node 1
node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65514 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node all

```
root@antbert> show security flow gate brief node all
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
```

```
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65512 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65510 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node local

```
root@antbert> show security flow gate brief node local
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65504 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate brief node primary

```
root@antbert> show security flow gate brief node primary
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC4 PIC0:

Hole: 1.0.0.100-1.0.0.100/0-0->2.0.0.100-2.0.0.100/32707-32707
```

```
  Translated: 1.0.0.100/0->2.0.0.100/32707
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65500 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate destination-port

Syntax **show security flow gate destination-port *destination-port-number* [brief | summary]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display information about temporary openings known as pinholes or gates in the security firewall that for the specified destination port.

NOTE: **Destination port filter matches the gate only if the given port falls within the range of ports specified in the gate.**

Options

- *destination-port-number*—Number of the destination port for which to display gate information.

  Range: 1 through 65,535

- brief | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow gate on page 449
- show security flow gate destination-prefix on page 2068

List of Sample Output

show security flow gate destination-port brief on page 2066
show security flow gate destination-port summary on page 2067

Output Fields Table 210 lists the output fields for the **show security flow gate destination-port** command. Output fields are listed in the approximate order in which they appear.

Table 210: show security flow gate destination-port Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 210: show security flow gate destination-port Output Fields *(continued)*

Field Name	Field Description
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Maximum gates	Number of maximum gates.

Sample Output

show security flow gate destination-port brief

```
root> show security flow gate destination-port 33253 brief
Flow Gates on FPC4 PIC1:

Hole: 40.0.0.111-40.0.0.111/0-0->30.0.0.100-30.0.0.100/33253-33253
  Translated: 40.0.0.111/0->30.0.0.100/33253
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65526 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24576-86016

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:
```


Copyright © 2016, Juniper Networks, Inc.

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate destination-port summary

```
root> show security flow gate destination-port 33253 summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

Copyright © 2016, Juniper Networks, Inc.

show security flow gate destination-prefix

Syntax **show security flow gate destination-prefix *destination-IP-prefix* [brief | summary]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display information about temporary openings known as pinholes or gates in the security firewall for the specified destination prefix.

NOTE: **Destination prefix must match both the starting and ending address in the gate.**

Options

- *destination-IP-prefix*—Destination IP prefix or address for which to display gate information.
- brief | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow gate on page 449
- show security flow gate destination-port on page 2065

List of Sample Output

show security flow gate destination-prefix brief on page 2069
show security flow gate destination-prefix summary on page 2070

Output Fields Table 211 lists the output fields for the **show security flow gate destination-prefix** command. Output fields are listed in the approximate order in which they appear.

Table 211: show security flow gate destination-prefix Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 211: show security flow gate destination-prefix Output Fields *(continued)*

Field Name	Field Description
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.

Sample Output

show security flow gate destination-prefix brief

```
root> show security flow gate destination-prefix 192.0.2.1 brief
Hole: 40.0.0.111-40.0.0.111/0-0->192.0.2.1-192.0.2.1/37308-37308
  Translated: 203.0.113.1/0->192.0.2.1/37308
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65456 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24575-86015

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow gate destination-prefix summary

```
root> show security flow gate destination-prefix 192.0.2.1 summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate protocol

Syntax **show security flow gate protocol *protocol-name* [brief | summary]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display information about temporary openings known as pinholes or gates in the security firewall for the specified protocol.

Options

- *protocol-name* —Protocol to use as a gate filter. Information about gates that use this protocol is displayed.

  Possible protocols are:

  - ah—IP Security Authentication Header
  - egp—Exterior gateway protocol
  - esp—IPsec Encapsulating Security Payload
  - gre—Generic routing encapsulation
  - icmp—Internet Control Message Protocol
  - icmp6—Internet Control Message Protocol
  - igmp—Internet Group Management Protocol
  - ipip—IP over IP
  - ospf—Open Shortest Path First
  - pim—Protocol Independent Multicast
  - rsvp—Resource Reservation Protocol
  - sctp—Stream Control Transmission Protocol
  - tcp—Transmission Control Protocol
  - udp—User Datagram Protocol

- brief | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation

- show security flow gate on page 449

List of Sample Output

show security flow gate protocol brief on page 2072
show security flow gate protocol summary on page 2073

Output Fields Table 212 lists the output fields for the **show security flow gate protocol** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 212: show security flow gate protocol Output Fields

Field Name	Field Description
Hole	Range of flows permitted by the pinhole.
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.

Sample Output

show security flow gate protocol brief

```
root> root> show security flow gate protocol tcp brief
Hole: 203.0.113.1-203.0.113.1/0-0->192.0.2.1-192.0.2.1/37308-37308
  Translated: 203.0.113.1/0->192.0.2.1/37308
  Protocol: tcp
  Application: FTP ALG/79
  Age: 65414 seconds
  Flags: 0x0080
  Zone: trust
  Reference count: 1
  Resource: 1-24575-86015

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

show security flow gate protocol summary

```
root> show security flow gate protocol tcp summary
Flow Gates on FPC4 PIC1:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1

Flow Gates on FPC5 PIC0:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0

Flow Gates on FPC5 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow gate summary node

Syntax **show security flow gate summary node (*node-id* | all | local | primary)**

Release Information Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

Description Display information about temporary openings known as pinholes or gates in the security firewall for the specified node options in summary mode.

Options **node**—(Optional) For chassis cluster configurations, display gate information on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- *show security flow gate*
- show security flow gate brief node on page 2059

List of Sample Output
show security flow gate summary node 0 on page 2075
show security flow gate summary node 1 on page 2076
show security flow gate summary node all on page 2076
show security flow gate summary node local on page 2077
show security flow gate summary node primary on page 2078

Output Fields Table 213 lists the output fields for the **show security flow gate summary node** command. Output fields are listed in the approximate order in which they appear.

Table 213: show security flow gate summary node Output Fields

Field Name	Field Description
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Number of gates in total.
Hole	Range of flows permitted by the pinhole.

Copyright © 2016, Juniper Networks, Inc.

Table 213: show security flow gate summary node Output Fields *(continued)*

Field Name	Field Description
Translated	Tuples used to create the session if it matches the pinhole. • Source address and port • Destination address and port
Protocol	Application protocol, such as UDP or TCP.
Application	Name of the application.
Age	Idle timeout for the pinhole.
Flags	Internal debug flags for the pinhole.
Zone	Incoming zone.
Reference count	Number of resource manager references to the pinhole.
Resource	Resource manager information about the pinhole.

Sample Output

show security flow gate summary node 0

```
root@antbert> show security flow gate summary node 0
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node 1

```
root@antbert> show security flow gate summary node 1
node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node all

```
root@antbert> show security flow gate summary node all
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:
```


Copyright © 2016, Juniper Networks, Inc.

```
Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

node1:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node local

```
root@antbert> show security flow gate summary node local
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072
```

```
Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```

show security flow gate summary node primary

```
root@antbert> show security flow gate summary node primary
node0:
--------------------------------------------------------------------------

Flow Gates on FPC3 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072

Flow Gates on FPC4 PIC0:

Valid gates: 1
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 1
Maximum gates: 131072

Flow Gates on FPC4 PIC1:

Valid gates: 0
Pending gates: 0
Invalidated gates: 0
Gates in other states: 0
Total gates: 0
Maximum gates: 131072
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session

Syntax **show security flow session**
[*filter*] [**brief | extensive | summary**]

Release Information Command introduced in Junos OS Release 8.5. Support for filter and view options added in Junos OS Release 10.2. Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2. Policy ID filter added in Junos OS Release 12.3X48-D10. Support for connection tag added in Junos OS Release 15.1X49-D40.

Description Display information about all currently active security sessions on the device.

NOTE: For the normal flow sessions, the **show security flow session** command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. Because of this, the output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific **show** command for examples of the filtered output.

  application—Predefined application name

  application-firewall—Application firewall enabled

  application-firewall-rule-set—Application firewall enabled with the specified rule set

  application-traffic-control—Application traffic control session

  application-traffic-control-rule-set—Application traffic control rule set name and rule name

  conn-tag—Session connection tag

  destination-port—Destination port

  destination-prefix—Destination IP prefix or address

  dynamic-application—Dynamic application

  dynamic-application-group—Dynamic application

  encrypted—Encrypted traffic

  extensive—Display detailed output

Copyright © 2016, Juniper Networks, Inc.

family—Display session by family

idp—IDP enabled sessions

interface—Name of incoming or outgoing interface

logical-system (all | *logical-system-name*)—Name of a specific logical system or **all** to display all logical systems

nat—Display sessions with network address translation

policy-id—Display session information based on policy ID; the range is 1 through 4,294,967,295

protocol—IP protocol number

resource-manager—Resource manager

root-logical-system—Display root logical system as default

security-intelligence—Display security intelligence sessions

services-offload—Display services offload sessions

session-identifier—Display session with specified session identifier

source-port—Source port

source-prefix—Source IP prefix

summary—Display output summary

tunnel—Tunnel sessions

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session all on page 1902

List of Sample Output
show security flow session on page 2082
show security flow session brief on page 2083
show security flow session extensive on page 2083
show security flow session summary on page 2083

Output Fields Table 26 lists the output fields for the **show security flow session** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 214: show security flow session Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Application traffic control rule-set	AppQoS rule set for this session.
Rule	AppQoS rule for this session.
Forwarding class	The AppQoS forwarding class name for this session that distinguishes the transmission priority
DSCP code point	Differentiated Services (DiffServ) code point (DSCP) value remarked by the matching rule for this session.

Copyright © 2016, Juniper Networks, Inc.

Table 214: show security flow session Output Fields *(continued)*

Field Name	Field Description
Loss priority	One of four priority levels set by the matching rule to control discarding a packet during periods of congestion. A high loss priority means a high probability that the packet could be dropped during a period of congestion.
Rate limiter client to server	The rate-limiter profile assigned to the client-to-server traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Rate limiter server to client	The rate-limiter profile assigned to the server-to-client traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Maximum number of sessions permitted.

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
 Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/10000/2000 --> 203.0.113.11/10000/1000;udp, Conn Tag: 0x0, If:
 reth0.0, Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session brief node

Syntax **show security flow session brief node (*node-id* | all | local | primary)**

Release Information Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

Description Display information about all currently active security sessions on the device for the specified node options in brief mode.

Options **node**—(Optional) For chassis cluster configurations, display session information on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- *show security flow session*

List of Sample Output

show security flow session brief node 0 on page 2086
show security flow session brief node 1 on page 2086
show security flow session brief node all on page 2087
show security flow session brief node local on page 2087
show security flow session brief node primary on page 2088

Output Fields Table 215 lists the output fields for the **show security flow session brief node** command. Output fields are listed in the approximate order in which they appear.

Table 215: show security flow session brief node Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Policy name	Policy that permitted the traffic.
State	Session state.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).

Copyright © 2016, Juniper Networks, Inc.

Table 215: show security flow session brief node Output Fields *(continued)*

Field Name	Field Description
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Total sessions	Total number of sessions.

Sample Output

show security flow session brief node 0

```
root@host> show security flow session brief node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout:
1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes:
626, CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes:
744, CP Session ID: 10000001

Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0


Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node 1

```
root@host> show security flow session brief node 1
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Backup, Timeout:
12, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 0, Bytes: 0,
CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 0, Bytes: 0,
 CP Session ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node all

```
root@host> show security flow session brief node all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout:
1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes:
626, CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes:
744, CP Session ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout:
1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes:
626, CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes:
744, CP Session ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node local

```
root@host> show security flow session brief node local
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout:
1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes:
626, CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes:
```

```
744, CP Session ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session brief node primary

```
root@host> show security flow session brief node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000001, Policy name: default-policy-00/2, State: Active, Timeout:
1696, Valid
Resource information : FTP ALG, 1, 0
  In: 203.0.113.1/60059 --> 203.0.113.2/21;tcp, If: reth0.0, Pkts: 14, Bytes:
626, CP Session ID: 10000001
  Out: 203.0.113.2/21 --> 203.0.113.1/60059;tcp, If: reth1.0, Pkts: 13, Bytes:
744, CP Session ID: 10000001
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session destination-port

Syntax **show security flow session destination-port**
destination-port-number **[brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 8.5; Filter and view options added in Junos OS Release 10.2.

Description Display information about each session that uses the specified destination port.

Options

- ***destination-port-number*** —Number of the destination port for which to display sessions information.
- **Range:** 1 through 65,535

brief | extensive | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session destination-port on page 1905

List of Sample Output

show security flow session destination-port 23 on page 2090
show security flow session destination-port 23 brief on page 2091
show security flow session destination-port 23 extensive on page 2091
show security flow session destination-port 23 summary on page 2092

Output Fields Table 216 lists the output fields for the **show security flow session destination-port** command. Output fields are listed in the approximate order in which they appear.

Table 216: show security flow session destination-port Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 216: show security flow session destination-port Output Fields *(continued)*

Field Name	Field Description
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session destination-port 23

```
root> show security flow session destination-port 23
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 198.51.100.10/15190 --> 192.0.2.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 109,
Bytes: 5874, CP Session ID: 430000093
```


Copyright © 2016, Juniper Networks, Inc.

```
  Out: 198.51.100.2/23 --> 198.51.100.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64,
Bytes: 4015, CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-port 23 brief

```
root> show security flow session destination-port 23 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 198.51.100.10/15190 --> 198.51.100.2/23;tcp, If: ge-7/1/0.0, Pkts: 109,
Bytes: 5874, CP Session ID: 430000093
  Out: 198.51.100.2/23 --> 198.51.100.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64,
Bytes: 4015, CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-port 23 extensive

```
root> show security flow session destination-port 23 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Status: Normal
Flags: 0x40/0x0/0x2008003
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1630
Session State: Valid
Start time: 65490, Duration: 207
   In: 198.51.100.10/15190 --> 198.51.100.2/23;tcp
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 109, Bytes: 5874
    CP Session ID: 430000093
   Out: 198.51.100.2/23 --> 2198.51.100.10/15190;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 64, Bytes: 4015
    CP Session ID: 430000093
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session destination-port 23 summary

```
root> show security flow session destination-port 23 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session destination-prefix

Syntax
show security flow session destination-prefix
destination-IP-prefix **[brief | extensive | summary]**

Release Information
Command introduced in Junos OS Release 8.5. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.
Support for IPv6 addresses added in Junos OS Release 10.2.
Filter and view options added in Junos OS Release 10.2.

Description
Display information about each session that matches the specified IP destination prefix.

Options
- ***destination-IP-prefix***—Destination IP prefix or address for which to display session information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session destination-port on page 1905

List of Sample Output
show security flow session destination-prefix 60/8 on page 2094
show security flow session destination-prefix 60/8 brief on page 2095
show security flow session destination-prefix 60/8 extensive on page 2095
show security flow session destination-prefix 60/8 summary on page 2096
show security flow session destination-prefix 10::10 on page 2096

Output Fields
Table 217 lists the output fields for the **show security flow session destination-prefix** command. Output fields are listed in the approximate order in which they appear.

Table 217: show security flow session destination-prefix Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).

Copyright © 2016, Juniper Networks, Inc.

Table 217: show security flow session destination-prefix Output Fields *(continued)*

Field Name	Field Description
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session destination-prefix 60/8

```
root> show security flow session destination-prefix 60/8
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1450, Valid
```


Copyright © 2016, Juniper Networks, Inc.

```
  In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes:
 5874, CP Session ID: 430000093
  Out: 198.51.100.1/23 --> 192.0.2.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes:
 4015, CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 60/8 brief

```
root> show security flow session destination-prefix 60/8 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1450, Valid
  In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes:
 5874, CP Session ID: 430000093
  Out: 198.51.100.1/23 --> 192.0.2.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes:
 4015, CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 60/8 extensive

```
root> show security flow session destination-prefix 60/8 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000098, Status: Normal
Flags: 0x40/0x0/0x2008003
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1172
Session State: Valid
Start time: 65490, Duration: 666
   In: 192.0.2.10/15190 --> 198.51.100.1/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 109, Bytes: 5874
    CP Session ID: 430000093
   Out: 198.51.100.1/23 --> 200.0.0.10/15190;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 64, Bytes: 4015
```

Copyright © 2016, Juniper Networks, Inc.

```
    CP Session ID: 430000093
Total sessions: 1
```

show security flow session destination-prefix 60/8 summary

```
root> show security flow session destination-prefix 60/8 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session destination-prefix 10::10

```
user@host> show security flow session destination-prefix 5001::2
Session ID: 50000004, Policy name: self-traffic-policy/1, Timeout: 2
In: 10::11/42756 --> 10::10/0;icmp, If: .local..0
Out: 10::10/0 --> 10::11/42756;icmp, If: ge-0/3/0.0

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session extensive node

Syntax **show security flow session extensive node (*node-id* | all | local | primary)**

Release Information Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.

Description Display information about all currently active security sessions on the device for the specified node options in extensive mode.

Options **node**—(Optional) For chassis cluster configurations, display session information on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- show security flow session on page 453

List of Sample Output
show security flow session extensive node 0 on page 2098
show security flow session extensive node 0 on page 2099
show security flow session extensive node all on page 2099
show security flow session extensive node local on page 2101
show security flow session extensive node primary on page 2102

Output Fields Table 218 lists the output fields for the **show security flow session extensive node** command. Output fields are listed in the approximate order in which they appear.

Table 218: show security flow session extensive node Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Status	Session status.
State	Session state.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Policy that permitted the traffic.

Copyright © 2016, Juniper Networks, Inc.

Table 218: show security flow session extensive node Output Fields *(continued)*

Field Name	Field Description
Source NAT pool	The name of the source pool where NAT is used.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Start time	Time when the session was created, offset from the system start time.
Duration	Length of time for which the session is active.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, flag, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, flag, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.

Sample Output

show security flow session extensive node 0

```
root@host> show security flow session extensive node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1778
Session State: Valid
Start time: 6466, Duration: 28
   In: 10.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
```


Copyright © 2016, Juniper Networks, Inc.

```
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node 0

```
root@host> show security flow session extensive node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1778
Session State: Valid
Start time: 6466, Duration: 28
   In: 10.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node all

```
root@host> show security flow session extensive node all
```

```
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1778
Session State: Valid
Start time: 6466, Duration: 28
   In: 10.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1


Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Backup
Flags: 0x10000042/0x0/0x10103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 14298
Session State: Valid
Start time: 6248, Duration: 115
   In: 10.0.2.1/52080 --> 203.0.113.1/24;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x60002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
```


Copyright © 2016, Juniper Networks, Inc.

```
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
   Out: 203.0.113.1/24 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x60002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10000003
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

show security flow session extensive node local

```
root@host> show security flow session extensive node local
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1584
Session State: Valid
Start time: 6466, Duration: 221
   In: 10.0.2.1/52080 --> 203.0.113.1/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/21 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session extensive node primary

```
root@host> show security flow session extensive node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000003, Status: Normal, State: Active
Flags: 0x8000042/0x8000000/0x110103
Policy name: default-policy-00/2
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1554
Session State: Valid
Start time: 6466, Duration: 252
   In: 10.0.2.1/52080 --> 203.0.113.1/21;tcp,
    Interface: reth0.0,
    Session token: 0x6, Flag: 0x40002621
    Route: 0x86193c2, Gateway: 100.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 10000004
   Out: 203.0.113.1/21 --> 10.0.2.1/52080;tcp,
    Interface: reth1.0,
    Session token: 0x6, Flag: 0x40002620
    Route: 0x86033c2, Gateway: 120.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 420
    CP Session ID: 10000004
Total sessions: 1

Flow Sessions on FPC0 PIC2:
Total sessions: 0

Flow Sessions on FPC0 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session family

Syntax show security flow session family (inet | inet6)
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 10.2.

Description Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options
- **inet**—Display details summary of IPv4 sessions.
- **inet6**—Display details summary of IPv6 sessions.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session family on page 1907

List of Sample Output
show security flow session family inet on page 2104
show security flow session family inet brief on page 2105
show security flow session family inet extensive on page 2105
show security flow session family inet summary on page 2107

Output Fields Table 219 lists the output fields for the **show security flow session family** command. Output fields are listed in the approximate order in which they appear.

Table 219: show security flow session family Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.

Copyright © 2016, Juniper Networks, Inc.

Table 219: show security flow session family Output Fields *(continued)*

Field Name	Field Description
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session family inet

```
root> show security flow session family inet
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000107, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 200.0.0.10/3 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000202
  Out: 60.0.0.1/19001 --> 200.0.0.10/3;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000202
Total sessions: 1

Flow Sessions on FPC10 PIC3:
```


Copyright © 2016, Juniper Networks, Inc.

```
Session ID: 430000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000110
  Out: 60.0.0.1/19001 --> 200.0.0.10/2;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000110

Session ID: 430000117, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 200.0.0.10/4 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000111
  Out: 60.0.0.1/19001 --> 200.0.0.10/4;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000111
Total sessions: 2
```

show security flow session family inet brief

```
root> show security flow session family inet brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/1 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000206
  Out: 60.0.0.1/19769 --> 200.0.0.10/1;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000206

Session ID: 420000117, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000207
  Out: 60.0.0.1/19769 --> 200.0.0.10/2;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000207
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000119, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/3 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000112
  Out: 60.0.0.1/19769 --> 200.0.0.10/3;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000112
Total sessions: 1
```

show security flow session family inet extensive

```
root> show security flow session family inet extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000111, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76455, Duration: 0
   In: 200.0.0.10/4 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
```

```
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
   Out: 60.0.0.1/20537 --> 200.0.0.10/4;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000123, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 2
Session State: Valid
Start time: 76454, Duration: 2
   In: 200.0.0.10/3 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
   Out: 60.0.0.1/20537 --> 200.0.0.10/3;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000131, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76421, Duration: 1
   In: 200.0.0.10/5 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
   Out: 60.0.0.1/20537 --> 200.0.0.10/5;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
Total sessions: 1
```

show security flow session family inet summary

```
root> show security flow session family inet summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session interface

Syntax
show security flow session interface
interface-name **[brief | extensive | summary]**

Release Information
Command introduced in Junos OS Release 8.5; Filter and view options added in Junos OS Release 10.2.

Description
Display information about each session that uses the specified interface. The interface name can be a session's incoming or outgoing interface.

Options
- ***interface-name***—Name of the interface on the device for which to display sessions information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session interface on page 1908

List of Sample Output
show security flow session interface ge-0/0/2.0 on page 2109
show security flow session interface ge-0/0/2.0 brief on page 2110
show security flow session interface ge-0/0/2.0 extensive on page 2110
show security flow session interface ge-7/1/1.0 summary on page 2111

Output Fields
Table 220 lists the output fields for the **show security flow session interface** command. Output fields are listed in the approximate order in which they appear.

Table 220: show security flow session interface Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 220: show security flow session interface Output Fields *(continued)*

Field Name	Field Description
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session interface ge-0/0/2.0

```
root> show security flow session interface ge-7/1/1.0
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000146, Policy name: default-policy-00/2, Timeout: 58, Valid
  In: 200.0.0.10/9 --> 60.0.0.2/21562;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000247
  Out: 60.0.0.2/21562 --> 200.0.0.10/9;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0,
CP Session ID: 420000247
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC3:

Session ID: 430000146, Policy name: default-policy-00/2, Timeout: 56, Valid
  In: 200.0.0.10/8 --> 60.0.0.2/21562;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000131
  Out: 60.0.0.2/21562 --> 200.0.0.10/8;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0,
CP Session ID: 430000131
Total sessions: 1
```

show security flow session interface ge-0/0/2.0 brief

```
root> show security flow session interface ge-7/1/1.0 brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000137, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/5 --> 60.0.0.2/23354;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 410000269
  Out: 60.0.0.2/23354 --> 200.0.0.10/5;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 410000269
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000151, Policy name: default-policy-00/2, Timeout: 54, Valid
  In: 200.0.0.10/1 --> 60.0.0.2/23354;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000252
  Out: 60.0.0.2/23354 --> 200.0.0.10/1;icmp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0,
CP Session ID: 420000252
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session interface ge-0/0/2.0 extensive

```
root> show security flow session interface ge-7/1/1.0 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000151, Status: Normal
Flags: 0x40/0x0/0x2000003
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 60, Current timeout: 48
Session State: Valid
Start time: 83328, Duration: 12
   In: 200.0.0.10/1 --> 60.0.0.2/23354;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000252
   Out: 60.0.0.2/23354 --> 200.0.0.10/1;icmp,
```


Copyright © 2016, Juniper Networks, Inc.

```
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0000020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000252
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session interface ge-7/1/1.0 summary

```
root> show security flow session interface ge-7/1/1.0 summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session nat

Syntax **show security flow session nat [brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 10.2.

Description Display sessions with network address translation.

Options **brief | extensive | summary**—Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- show security flow session on page 453

List of Sample Output
show security flow session nat brief on page 2113
show security flow session nat extensive on page 2113
show security flow session nat summary on page 2114

Output Fields Table 221 lists the output fields for the **show security flow session nat** command. Output fields are listed in the approximate order in which they appear.

Table 221: show security flow session nat Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
Resource information	Information about the session particular to the resource manager, including the name of the ALG, the group ID. and the resource ID.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.

Copyright © 2016, Juniper Networks, Inc.

Table 221: show security flow session nat Output Fields *(continued)*

Field Name	Field Description
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.

Sample Output

show security flow session nat brief

```
root> show security flow session nat brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000390, Policy name: default-policy-00/2, Timeout: 1778, Valid
  In: 200.0.0.10/41043 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 414,
 CP Session ID: 420001090
  Out: 60.0.0.2/21 --> 60.0.0.1/19473;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 479,
CP Session ID: 430000964
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session nat extensive

```
root> show security flow session nat extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000390, Status: Normal
Flags: 0x2/0x0/0x2010103
Policy name: default-policy-00/2
Source NAT pool: interface, Application: junos-ftp/1
```

Copyright © 2016, Juniper Networks, Inc.

```
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1770
Session State: Valid
Start time: 151971, Duration: 55
   In: 200.0.0.10/41043 --> 60.0.0.2/21;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0002621
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 414
    CP Session ID: 420001090
   Out: 60.0.0.2/21 --> 60.0.0.1/19473;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xe0002620
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 479
    CP Session ID: 430000964
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session nat summary

```
root> show security flow session nat summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session policy-id

Syntax **show security flow session policy-id** *policy-id-number* **[brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 12.3X48-D10.

Description Display information about each session by using policy id of the session.

Options

- ***policy-id-number*** —ID of the policy that the first packet of the session matches with.

  Range: 1through 4294967295

- brief | extensive | summary–Display the specified level of output.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session protocol on page 1910

List of Sample Output

show security flow session policy-id 4 on page 2116
show security flow session policy-id 4 extensive on page 2116

Output Fields Table 222 lists the output fields for the **show security flow session policy-id** command. Output fields are listed in the approximate order in which they appear.

Table 222: show security flow session policy-id Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	For the input flow: • Source and destination addresses and protocol tuple for the input flow. • **Interface**: Input flow interface. • **Session token**: Internal token derived from the virtual routing instance. • **Flag**: Internal debugging flags. • **Route**: Internal next hop of the route to be used by the flow. • **Gateway**: Next-hop gateway of the flow. • **Tunnel**: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • **Port Sequence, FIN sequence, FIN state, Cookie**: Internal TCP state tracking information.

Copyright © 2016, Juniper Networks, Inc.

Table 222: show security flow session policy-id Output Fields *(continued)*

Field Name	Field Description
Out	For the reverse flow: • Source and destination addresses, and protocol tuple for the reverse flow. • **Interface**: Reverse flow interface. • **Session token**: Internal token derived from the virtual routing instance. • **Flag**: Internal debugging flags. • **Route**: Internal next hop of the route to be used by the flow. • **Gateway**: Next-hop gateway of the flow. • **Tunnel**: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • **Port Sequence, FIN sequence, FIN state, Cookie**: Internal TCP state tracking information.
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.

Sample Output

show security flow session policy-id 4

```
root> show security flow session policy-id 4
Flow Sessions on FPC1 PIC0:

Session ID: 20093273, Policy name: p1/4, Timeout: 1784, Valid
  In: 101.0.0.2/1 --> 111.0.0.3/1;0, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
  Out: 111.0.0.3/1 --> 201.0.0.1/22643;0, If: ge-0/0/1.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

show security flow session policy-id 4 extensive

```
root> show security flow session policy-id 4 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC2:

Session ID: 420000428, Status: Normal
Flags: 0x0/0x0/0x2008003
Policy name: p1/4
Source NAT pool: interface, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1740
Session State: Valid
Start time: 152305, Duration: 64
   In: 200.0.0.10/15192 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0001021
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 40, Bytes: 2251
    CP Session ID: 420001128
   Out: 60.0.0.2/23 --> 60.0.0.1/8078;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xe0001020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 28, Bytes: 1714
    CP Session ID: 430000965
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session protocol

Syntax **show security flow session protocol** (***protocol-name*** | ***protocol-number***)
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

Description Display information about each session that uses the specified protocol.

Options ***protocol-name*** —(Optional) Protocol to use as a sessions filter. Information about sessions that use this protocol is displayed. Possible protocols are:

- **ah**—IP Security Authentication Header
- **egp**—Exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload
- **gre**—Generic routing encapsulation
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP over IP
- **ospf**—Open Shortest Path First
- **pim**—Protocol Independent Multicast
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Transmission Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

protocol-number —(Optional) Numeric protocol value. For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers* (*for the Internet Protocol Suite*).

Range: 0 through 255

brief | extensive | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session protocol on page 1910

List of Sample Output
show security flow session protocol icmp on page 2120
show security flow session protocol icmp brief on page 2120
show security flow session protocol icmp extensive on page 2121

Copyright © 2016, Juniper Networks, Inc.

show security flow session protocol icmp summary on page 2121

Output Fields Table 223 lists the output fields for the **show security flow session protocol** command. Output fields are listed in the approximate order in which they appear.

Table 223: show security flow session protocol Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 223: show security flow session protocol Output Fields *(continued)*

Field Name	Field Description
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session protocol icmp

```
root> show security flow session protocol icmp
Flow Sessions on FPC10 PIC1:

Session ID: 410000654, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.2/15685;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 410001264
  Out: 60.0.0.2/15685 --> 200.0.0.10/2;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 410001264
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Session ID: 430000399, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/3 --> 60.0.0.2/15685;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430001053
  Out: 60.0.0.2/15685 --> 200.0.0.10/3;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 430001053
Total sessions: 1
```

show security flow session protocol icmp brief

```
root> show security flow session protocol icmp brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000658, Policy name: p1/4, Timeout: 4, Valid
  In: 200.0.0.10/4 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 410001268
  Out: 60.0.0.2/16453 --> 200.0.0.10/4;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 410001268
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000612, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/5 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420001316
  Out: 60.0.0.2/16453 --> 200.0.0.10/5;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 420001316
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC3:

Session ID: 430000405, Policy name: p1/4, Timeout: 2, Valid
  In: 200.0.0.10/6 --> 60.0.0.2/16453;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430001059
  Out: 60.0.0.2/16453 --> 200.0.0.10/6;icmp, If: ge-7/1/1.0, Pkts: 1, Bytes: 84,
 CP Session ID: 430001059
Total sessions: 1
```

show security flow session protocol icmp extensive

```
root> show security flow session protocol icmp extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000660, Status: Normal
Flags: 0x80000040/0x0/0x2800003
Policy name: p1/4
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 2
Session State: Valid
Start time: 153201, Duration: 3
   In: 200.0.0.10/8 --> 60.0.0.2/16453;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410001270
   Out: 60.0.0.2/16453 --> 200.0.0.10/8;icmp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0000020
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410001270
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session protocol icmp summary

```
root> show security flow session protocol icmp summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 1
Sessions in other states: 0
Total sessions: 3

Flow Sessions on FPC10 PIC2:
```

Copyright © 2016, Juniper Networks, Inc.

```
Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 1
Sessions in other states: 0
Total sessions: 3
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager

Syntax **show security flow session resource-manager**
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

Description Display information about sessions created by the resource manager.

Options none—Display all resource manager sessions.

brief | extensive | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session resource-manager on page 1912

List of Sample Output
show security flow session resource-manager on page 2124
show security flow session resource-manager brief on page 2125
show security flow session resource-manager extensive on page 2125
show security flow session resource-manager summary on page 2126

Output Fields Table 28 lists the output fields for the **show security flow session resource-manager** command. Output fields are listed in the approximate order in which they appear.

Table 224: show security flow session resource-manager Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
Resource information	Information about the session particular to the resource manager, including the name of the ALG, the group ID. and the resource ID.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 224: show security flow session resource-manager Output Fields *(continued)*

Field Name	Field Description
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.

Sample Output

show security flow session resource-manager

```
root> show security flow session resource-manager
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1734, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586,
 CP Session ID: 410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes:
803, CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager brief

```
root> show security flow session resource-manager brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Policy name: p1/4, Timeout: 1704, Valid
Resource information : FTP ALG, 1, 0
  In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 586,
 CP Session ID: 410001274
  Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp, If: ge-7/1/1.0, Pkts: 13, Bytes:
803, CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session resource-manager extensive

```
root> show security flow session resource-manager extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000664, Status: Normal
Flags: 0x42/0x0/0x2010103
Policy name: p1/4
Source NAT pool: Null, Application: junos-ftp/1
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1682
Session State: Valid
Start time: 160496, Duration: 153
Client: FTP ALG, Group: 1, Resource: 0
   In: 200.0.0.10/41047 --> 60.0.0.2/21;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0002621
    Route: 0x70010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 586
    CP Session ID: 410001274
   Out: 60.0.0.2/21 --> 200.0.0.10/41047;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0xc0002620
    Route: 0x80010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 13, Bytes: 803
    CP Session ID: 410001274
Total sessions: 1

Flow Sessions on FPC10 PIC2:
Total sessions: 0

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session resource-manager summary

```
root> show security flow session resource-manager summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session services-offload

Syntax **show security flow session services-offload**
[*filter*] [**brief | extensive | summary**]

Release Information Command introduced in Junos OS Release 11.4. Low-latency option introduced in Junos OS Release 12.1X44-D10.
Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 devices.

Description Display information about all currently active services-offload security sessions on the device.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter:

  application —Application name.

  application-firewall-rule-set—Application firewall enabled with the specified rule set.

  application-traffic-control-rule-set—Application traffic control enabled with the specified rule set.

  destination-port—Destination port.

  destination-prefix—Destination IP prefix or address.

  dynamic-application—Dynamic application name.

  dynamic-application-group—Dynamic application group name.

  encrypted—Show encrypted traffic.

  family—Protocol family.

  interface—Name of incoming or outgoing interface.

  logical-system—Logical system name.

  protocol—IP protocol number.

  root-logical-system—Root logical system name.

  source-port—Source port.

  source-prefix—Source IP prefix or address.

- **brief | extensive | summary**—Display the specified level of output.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session services-offload on page 1913

List of Sample Output
show security flow session services-offload on page 2129
show security flow session services-offload brief on page 2129
show security flow session services-offload extensive on page 2130
show security flow session services-offload summary on page 2130

Output Fields Table 225 lists the output fields for the **show security flow session services-offload** command. Output fields are listed in the approximate order in which they appear.

Table 225: show security flow session services-offload Output Fields

Field Name	Field Description
Session ID	Number that identifies the services-offload session. Use this ID to get more information about the session.
Policy name	Policy that permits the services-offload traffic.
Timeout	Idle timeout period after which the services-offload session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets, and bytes).
Total sessions	Total number of services-offload sessions.
Status	Services-offload session status.
Flag	Internal flag depicting the state of the services-offload session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the services-offload session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Dynamic application	Name of the dynamic application.
Maximum timeout	Maximum amount of idle time allowed for the services-offload session.
Current timeout	Number of seconds that the current services-offload session has been idle.

Copyright © 2016, Juniper Networks, Inc.

Table 225: show security flow session services-offload Output Fields *(continued)*

Field Name	Field Description
Session State	Services-offload session state.
Start time	Time when the services-offload session was created, offset from the system start time.
Duration	Duration of the services-offload session.
Valid sessions	Number of valid services-offload sessions.
Pending sessions	Number of pending services-offload sessions.
Invalidated sessions	Number of invalidated services-offload sessions.
Sessions in other states	Number of services-offload sessions in other states.
Total sessions	Total number of services-offload sessions.

Sample Output

show security flow session services-offload

```
user@host>show security flow session services-offload
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 1788, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session services-offload brief

```
user@host>show security flow session services-offload brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 1748, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session services-offload extensive

```
user@host>show security flow session services-offload extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1718
Session State: Valid
Start time: 165, Duration: 89
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session services-offload summary

```
user@host>show security flow session services-offload summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

```
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session session-identifier

Syntax **show security flow session session-identifier** ***session-identifier***

Release Information Command introduced in Junos OS Release 8.5. Output changed to support natflag2 and flag in Junos OS Release 12.3X48-D10.

Description Display detailed information for the session with this identifier.

Options ***session-identifier*** —Identifier of the session about which to display information.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session session-identifier on page 1916

List of Sample Output
show security flow session session-identifier 420000002 on page 2133
show security flow session session-identifier 2218 on page 2134
show security flow session session-identifier 33 on page 2134

Output Fields Table 226 lists the output fields for the **show security flow session session-identifier** command. Output fields are listed in the approximate order in which they appear.

Table 226: show security flow session session-identifier Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Virtual system	Virtual system to which the session belongs.
Policy name	Name and ID of the policy that the first packet of the session matched.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Start time	Time when the session was created, offset from the system start time.

Copyright © 2016, Juniper Networks, Inc.

Table 226: show security flow session session-identifier Output Fields *(continued)*

Field Name	Field Description
Duration	Length of time for which the session is active.
In	For the input flow: • Source and destination addresses and protocol tuple for the input flow. • **Interface**: Input flow interface. • **Session token**: Internal token derived from the virtual routing instance. • **Flag**: Internal debugging flags. • **Route**: Internal next hop of the route to be used by the flow. • **Gateway**: Next-hop gateway of the flow. • **Tunnel**: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • **Port Sequence, FIN sequence, FIN state, Cookie**: Internal TCP state tracking information.
Out	For the reverse flow: • Source and destination addresses, and protocol tuple for the reverse flow. • **Interface**: Reverse flow interface. • **Session token**: Internal token derived from the virtual routing instance. • **Flag**: Internal debugging flags. • **Route**: Internal next hop of the route to be used by the flow. • **Gateway**: Next-hop gateway of the flow. • **Tunnel**: If the flow is going into a tunnel, the tunnel ID. Otherwise, 0 (zero). • **Port Sequence, FIN sequence, FIN state, Cookie**: Internal TCP state tracking information.

Sample Output

show security flow session session-identifier 420000002

```
root> show security flow session session-identifier 420000002
Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 1586
Session State: Valid
Start time: 165, Duration: 220
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
```

Copyright © 2016, Juniper Networks, Inc.

```
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1
```

Sample Output

show security flow session session-identifier 2218

```
user@host> show security flow session session-identifier 2218
Flow Sessions on FPC4 PIC1:

Session ID: 2218, Status: Normal,
Flags: 0x80000040, 0xffffffff, 0xffffffff
Virtual system: Root VSYS(I), Policy name: foo/4
Maximum timeout: 60, Current timeout: 60
Start time: 0, Duration: 0
Client: MGCP ALG, Group: 2047, Resource: 8188
   In: 12.0.102.26/28072 --> 11.0.101.236/23252;udp,
    Interface: ge-0/0/2.0,
    Session token: 0xa, Flag: 0x8094740
    Route: 0xb0010, Gateway: 12.0.102.26, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0, Cookie: 0,
   Out: 11.0.101.236/23252 --> 12.0.102.26/28072;udp,
    Interface: ge-0/0/1.0,
    Session token: 0x8, Flag: 0x8094740
Route: 0xa0010, Gateway: 11.0.101.236, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0, Cookie: 0,
Total sessions: 1
```

Sample Output

show security flow session session-identifier 33

```
user@host> show security flow session session-identifier 33
Flow Sessions on FPC4 PIC1:

Session ID: 33, Status: Normal,
Flags: 0x80000040, 0xffffffff, 0xffffffff
Virtual system: Root VSYS(I), Policy name: default-policy/2
Application: junos-ftp/1
Maximum timeout: 1800, Current timeout: 1492
Start time: 31128, Duration: 121
   In: 10.10.10.1/2851 --> 192.168.0.2/21;tcp,
    Interface: t1-1/0/0.0,
Session token: 0x6, Flag: 0x80a15e0
    Route: 0x60010, Gateway: 10.10.10.0, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0, Cookie: 0,
   Out: 192.168.0.2/21 --> 10.10.10.1/2851;tcp,
    Interface: ge-0/0/1.0,
    Session token: 0x6, Flag: 0x80a15e0
    Route: 0x90010, Gateway: 192.168.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0, Cookie: 0,
Total sessions: 1
```


Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show security flow session source-port

Syntax
show security flow session source-port
source-port-number
[brief | extensive | summary]

Release Information
Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

Description
Display information about each session that uses the specified source port.

Options
- ***source-port-number*** —Number of the source port about which to display sessions information.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session source-port on page 1917

List of Sample Output
show security flow session source-port 15198 on page 2137
show security flow session source-port 15198 brief on page 2138
show security flow session source-port 15198 extensive on page 2138
show security flow session source-port 15198 summary on page 2138

Output Fields
Table 227 lists the output fields for the **show security flow session source-port** command. Output fields are listed in the approximate order in which they appear.

Table 227: show security flow session source-port Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
Resource information	Information about the session particular to the resource manager, including the name of the ALG, the group ID, and the resource ID.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).

Copyright © 2016, Juniper Networks, Inc.

Table 227: show security flow session source-port Output Fields *(continued)*

Field Name	Field Description
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.

Sample Output

show security flow session source-port 15198

```
root> show security flow session source-port 15198
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 770, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session source-port 15198 brief

```
root> show security flow session source-port 15198 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 740, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-port 15198 extensive

```
root> show security flow session source-port 15198 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 750
Session State: Valid
Start time: 165, Duration: 1056
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-port 15198 summary

```
root> show security flow session source-port 15198 summary
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session source-prefix

Syntax

show security flow session source-prefix
source-prefix-number
[brief | extensive | summary]

Release Information

Command introduced in Junos OS Release 8.5.
Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.
Filter and view options introduced in Junos OS Release 10.2.

Description

Display information about each session that uses the specified source prefix.

Options

source-prefix-number—Source IP prefix or address for which to display sessions information.

brief | extensive | summary—Display the specified level of output.

Required Privilege Level

view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session source-prefix on page 1919

List of Sample Output

show security flow session source-prefix 200.0.0.10 on page 2141
show security flow session source-prefix 200.0.0.10 brief on page 2142
show security flow session source-prefix 200.0.0.10 extensive on page 2142
show security flow session source-prefix 200.0.0.10 summary on page 2142

Output Fields

Table 228 lists the output fields for the **show security flow session source-prefix** command. Output fields are listed in the approximate order in which they appear.

Table 228: show security flow session source-prefix Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).

Copyright © 2016, Juniper Networks, Inc.

Table 228: show security flow session source-prefix Output Fields *(continued)*

Field Name	Field Description
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.

Sample Output

show security flow session source-prefix 200.0.0.10

```
root> show security flow session source-prefix 200.0.0.10
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 488, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session source-prefix 200.0.0.10 brief

```
root> show security flow session source-prefix 200.0.0.10 brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Policy name: p1/4, Timeout: 482, Valid
  In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 9, Bytes: 507,
 CP Session ID: 420000002
  Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp, If: ge-7/1/1.0, Pkts: 8, Bytes: 462,
 CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-prefix 200.0.0.10 extensive

```
root> show security flow session source-prefix 200.0.0.10 extensive
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000002, Status: Normal
Flags: 0x40/0x0/0x2408003, services-offload
Policy name: p1/4
Source NAT pool: Null, Application: junos-telnet/10
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 1800, Current timeout: 436
Session State: Valid
Start time: 165, Duration: 1370
   In: 200.0.0.10/15198 --> 60.0.0.2/23;tcp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0x42001021
    Route: 0x80010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 9, Bytes: 507
    CP Session ID: 420000002
   Out: 60.0.0.2/23 --> 200.0.0.10/15198;tcp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x42001020
    Route: 0x70010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 8, Bytes: 462
    CP Session ID: 420000002
Total sessions: 1

Flow Sessions on FPC10 PIC3:
Total sessions: 0
```

show security flow session source-prefix 200.0.0.10 summary

```
root> show security flow session source-prefix 200.0.0.10 summary
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session summary family

Syntax show security flow session summary family (inet | inet6)

Release Information Command introduced in Junos OS Release 10.2.
Support on SRX Series devices for flow-based mode for family inet6 added in Junos OS Release 10.2.
Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

Description Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options
- **inet**—Display details summary of IPv4 sessions.
- **inet6**—Display details summary of IPv6 sessions.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session all on page 1902

List of Sample Output
show security flow session summary family inet on page 2144
show security flow session summary family inet6 on page 2145

Output Fields Table 229 lists the output fields for the **show security flow session summary family** command. Output fields are listed in the approximate order in which they appear.

Table 229: show security flow session summary Output Fields

Field Name	Field Description
Valid sessions	Count of valid sessions.
Pending sessions	Count of pending sessions.
Invalidated sessions	Count of sessions the security device has determined to be invalid.
Sessions in other states	Count of sessions not in valid, pending, or invalidated state.
Total sessions	Total of the above counts.

Sample Output

show security flow session summary family inet

```
user@host> show security flow session summary family inet
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0
```

show security flow session summary family inet6

```
user@host> show security flow session summary family inet6
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session summary node

Syntax **show security flow session summary node (*node-id* | all | local | primary)**

Release Information Command introduced in Junos OS Release 8.5; node options added in Junos OS Release 9.0. Filter options added in Junos OS Release 10.2.
Support on SRX Series devices for flow-based mode for family inet6 added in Junos OS Release 10.2.
Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4.

Description Display information about all currently active security sessions on the device for the specified node options in summary mode.

Options **node**—(Optional) For chassis cluster configurations, display session information on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- show security flow session on page 453

List of Sample Output
show security flow session summary node 0 on page 2147
show security flow session summary node 1 on page 2148
show security flow session summary node all on page 2148
show security flow session summary node local on page 2150
show security flow session summary node primary on page 2150

Output Fields Table 230 lists the output fields for the **show security flow session summary node** command. Output fields are listed in the approximate order in which they appear.

Table 230: show security flow session summary node Output Fields

Field Name	Field Description
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.

Copyright © 2016, Juniper Networks, Inc.

Table 230: show security flow session summary node Output Fields *(continued)*

Field Name	Field Description
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session summary node 0

```
root@host> show security flow session summary node 0
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session summary node 1

```
root@host> show security flow session summary node 1
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node all

```
root@host> show security flow session summary node all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
```


Copyright © 2016, Juniper Networks, Inc.

```
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session summary node local

```
root@host> show security flow session summary node local
node0:
--------------------------------------------------------------------------


Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security flow session summary node primary

```
root@host> show security flow session summary node primary
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC0 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC0 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session summary services-offload

Syntax **show security flow session summary services-offload** [*filter*]

Release Information Command introduced in Junos OS Release 11.4.
Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) with Express Path (formerly known as *services offloading*) support are introduced for SRX5400, SRX5600, and SRX5800 devices.

Description Display information about all currently active services-offload security sessions on the device in summary mode.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter:

  application—Application name.

  application-firewall-rule-set—Application firewall enabled with the specified rule set.

  application-traffic-control-rule-set—Application traffic control enabled with the specified rule set.

  destination-port—Destination port.

  destination-prefix—Destination IP prefix or address.

  dynamic-application—Dynamic application name.

  dynamic-application-group—Dynamic application group name.

  family—Protocol family.

  interface—Name of incoming or outgoing interface.

  logcal-system—Logical system name.

  protocol—IP protocol number.

  root-logical-system—Root logical system name.

  source-port—Source port.

  source-prefix—Source IP prefix or address.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session services-offload on page 1913

Copyright © 2016, Juniper Networks, Inc.

List of Sample Output

show security flow session summary services-offload on page 2153
show security flow session summary services-offload application on page 2154
show security flow session summary services-offload destination-port on page 2154

Output Fields

Table 231 lists the output fields for the **show security flow session summary services-offload** command. Output fields are listed in the approximate order in which they appear.

Table 231: show security flow session summary services-offload Output Fields

Field Name	Field Description
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Services-offload-sessions	Number of services-offload sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use: • Valid • Pending • Invalidated • Sessions in other states
Maximum-sessions	Maximum number of sessions.

Sample Output

show security flow session summary services-offload

```
user@host> show security flow session summary services-offload
Flow Sessions on FPC1 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 409600

Flow Sessions on FPC2 PIC0:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 1
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC3 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200

Flow Sessions on FPC5 PIC0:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 819200
```

show security flow session summary services-offload application

```
user@host> show security flow session summary services-offload application telnet
Flow Sessions on FPC10 PIC1:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Valid sessions: 0
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 0

Flow Sessions on FPC10 PIC3:

Valid sessions: 1
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 1
```

show security flow session summary services-offload destination-port

```
user@host> show security flow session summary services-offload destination-port 23
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC3:

Session ID: 430000004, Policy name: p1/4, Timeout: 1500, Valid
  In: 200.0.0.10/15200 --> 60.0.0.2/23;tcp, If: ge-7/1/0.0, Pkts: 13, Bytes: 718,
 CP Session ID: 430000003
  Out: 60.0.0.2/23 --> 200.0.0.10/15200;tcp, If: ge-7/1/1.0, Pkts: 12, Bytes:
677, CP Session ID: 430000003
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session tunnel

Syntax **show security flow session tunnel**
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 8.5; Filter and view options introduced in Junos OS Release 10.2.

Description Display information about all tunnel sessions.

Options none—Display all tunnel sessions.

brief | extensive | summary—Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow session tunnel on page 2157
show security flow session tunnel brief on page 2158
show security flow session tunnel extensive on page 2158
show security flow session tunnel summar on page 2160

Output Fields Table 232 lists the output fields for the **show security flow session tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 232: show security flow session tunnel Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. You can use this ID to get additional information about the session.
Policy name	Policy that permitted the traffic. NA (Not Applicable) for a tunnel session.
Timeout	Idle timeout after which the session expires. NA (Not Applicable) for a tunnel session.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.

Copyright © 2016, Juniper Networks, Inc.

Table 232: show security flow session tunnel Output Fields *(continued)*

Field Name	Field Description
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Valid sessions	Number of valid sessions.
Pending sessions	Number of pending sessions.
Invalidated sessions	Number of invalidated sessions.
Sessions in other states	Number of sessions in other states.

Sample Output

show security flow session tunnel

```
root> show security flow session tunnel
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/43405 --> 60.0.0.3/494;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP
 Session ID: 420000000

Session ID: 410000002, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000

Session ID: 420000004, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000

Session ID: 430000006, Policy name: N/A, Timeout: N/A, Valid
```

Copyright © 2016, Juniper Networks, Inc.

```
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2
```

show security flow session tunnel brief

```
root> show security flow session tunnel brief
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/43405 --> 60.0.0.3/494;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP
 Session ID: 420000000

Session ID: 410000002, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000

Session ID: 420000004, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;esp, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000

Session ID: 430000006, Policy name: N/A, Timeout: N/A, Valid
  In: 60.0.0.2/0 --> 60.0.0.3/0;ah, If: ge-7/1/1.0, Pkts: 0, Bytes: 0, CP Session
 ID: 420000000
Total sessions: 2
```

show security flow session tunnel extensive

```
root> show security flow session tunnel extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000001, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3548, Duration: 797
   In: 60.0.0.2/43405 --> 60.0.0.3/494;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x80100621
    Route: 0x60010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
```


Copyright © 2016, Juniper Networks, Inc.

```
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000

Session ID: 410000002, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3548, Duration: 797
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x60010, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Session ID: 420000003, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 798
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000

Session ID: 420000004, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 798
   In: 60.0.0.2/0 --> 60.0.0.3/0;ah,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
    CP Session ID: 420000000
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000005, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 799
   In: 60.0.0.2/0 --> 60.0.0.3/0;esp,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000

Session ID: 430000006, Status: Normal
Flags: 0x10000/0x0/0x1
Policy name: N/A
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: N/A, Current timeout: N/A
Session State: Valid
Start time: 3513, Duration: 799
   In: 60.0.0.2/0 --> 60.0.0.3/0;ah,
    Interface: ge-7/1/1.0,
    Session token: 0x7, Flag: 0x621
    Route: 0x0, Gateway: 60.0.0.2, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 420000000
Total sessions: 2
```

show security flow session tunnel summar

```
root> show security flow session tunnel summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2

Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 0
Sessions in other states: 0
Total sessions: 2
```

Copyright © 2016, Juniper Networks, Inc.

show security flow statistics

Syntax **show security flow statistics**

Release Information Command introduced in Junos OS Release 10.2.
Starting with Junos OS Release 15.1X49-D10, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based datapath packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).
The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Description Display flow-related system statistics.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow statistics on page 2162
show security flow statistics (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3) on page 2163

Output Fields Table 233 lists the output fields for the **show security flow statistics** command. Output fields are listed in the approximate order in which they appear.

Table 233: show security flow statistics Output Fields

Field Name	Field Description
Current sessions	Number of current sessions.
Packets forwarded	Number of packets forwarded.
Packets dropped	Number of Packets dropped.
Fragment packets	Number of fragment packets.

Sample Output

show security flow statistics

```
root> show security flow statistics
Flow Statistics of FPC4 PIC1:
    Current sessions: 63
    Packets forwarded: 3001
    Packets dropped: 1281
    Fragment packets: 0

  Flow Statistics of FPC5 PIC0:
    Current sessions: 22
    Packets forwarded: 859
```


Copyright © 2016, Juniper Networks, Inc.

```
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC5 PIC1:
    Current sessions: 22
    Packets forwarded: 858
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics Summary:
    System total valid sessions: 107
    Packets forwarded: 4718
    Packets dropped: 1281
    Fragment packets: 0
```

show security flow statistics (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
root> show security flow statistics
Flow Statistics of FPC0 PIC1:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC0 PIC2:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC0 PIC3:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics Summary:
    System total valid sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow status

Syntax **show security flow status**

Release Information Command introduced in Junos OS Release 10.2; session distribution mode option added in Junos OS Release 12.1X44-D10; enhanced route scaling mode option added in Junos OS Release 12.1X45-D10. GTP-U distribution option added in Junos OS Release 15.1X49-D40.
Starting with Junos OS Release 15.1X49-D10, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based data path packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).
The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Description Display the flow processing modes and logging status.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow status on page 2165
show security flow status (IPsec Performance Acceleration) on page 2165
show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3) on page 2166

Output Fields Table 234 lists the output fields for the **show security flow status** command. Output fields are listed in the approximate order in which they appear.

Table 234: show security flow status Output Fields

Field Name	Field Description
Flow forwarding mode	Flow processing mode. • Inet forwarding mode • Inet6 forwarding mode • MPLS forwarding mode • ISO forwarding mode • Session distribution mode • Enhanced route scaling mode
Flow trace status	Flow logging status. • Flow tracing status • Flow tracing options

Copyright © 2016, Juniper Networks, Inc.

Table 234: show security flow status Output Fields *(continued)*

Field Name	Field Description
flow session distribution	SPU load distribution mode. • RR-based • Hash-based GTP-U distribution • Enabled
Flow packet ordering	packet-ordering mode. • Hardware • Software
Flow ipsec performance acceleration	IPsec VPN performance acceleration status.

Sample Output

show security flow status

```
root> show security flow status
Flow forwarding mode:
Inet forwarding mode: flow based
Inet6 forwarding mode: flow based
MPLS forwarding mode: drop
ISO forwarding mode: drop
Enhanced route scaling mode: Enabled (reboot needed to disable)
Flow trace status
Flow tracing status: on
Flow tracing options: all
Flow session distribution
Distribution mode: Hash-based
GTP-U distribution: Enabled
Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
```

show security flow status (IPsec Performance Acceleration)

```
root> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
Flow trace status
    Flow tracing status: off
Flow session distribution
    Distribution mode: RR-based
    GTP-U distribution: Enabled Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
Flow ipsec performance acceleration: on
```

Copyright © 2016, Juniper Networks, Inc.

show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
root> show security flow status
node0:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware

node1:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
```


Copyright © 2016, Juniper Networks, Inc.

show security forward-options mirror-filter

Syntax show security forward-options mirror-filter (all | *filter-name*)

Release Information Command introduced in Junos OS Release 12.1X46-D10.

Description Display information about the status of the mirror filters.

Options **all**—Display counters for all mirror filters.

filter-name—Name of the mirror filter for which the counter is displayed.

Required Privilege Level view

Related Documentation
- mirror-filter (Security Forwarding Options) on page 1866
- clear security forward-options mirror filter on page 1921

List of Sample Output show security forward-options mirror-filter on page 2167

Output Fields Lists the output fields for the **show security forward-options mirror-filter** command. Output fields are listed in the approximate order in which they appear in the output.

Table 235: show security forward-options mirror-filter

Field Name	Field Description
mirror-filter-name	Name of the mirror filter configured on the device.
interface-in	Display the name of the incoming logical interface to be matched for mirroring.
interface-out	Display the outgoing logical interface to be matched for mirroring.
protocol	Display the networking protocol name or number to be matched for mirroring.
source-prefix	Display the source IP prefix or address to be matched for mirroring.
destination-prefix	Display the destination IP prefix or address to be matched for mirroring.
filter-counters	Display the number of packets matched for mirroring.
output-counter	Display the number of packets sent to the sniffing device.

Sample Output

show security forward-options mirror-filter

```
user@host> show security forward-options mirror-filter m10

node0:
```

```
----------------------------------
  Security mirror status

    mirror-filter-name: m10
    interface-in: reth0.0
    interface-out: reth1.0
    protocol: 132
    source-prefix: 11.1.1.0
    destination-prefix: 41.1.1.0
    filter-counters: 143
    output-counters: 143

node1:
----------------------------------
  Security mirror status

    mirror-filter-name: m10
    interface-in: reth0.0
    interface-out: reth1.0
    protocol: 132
    source-prefix: 11.1.1.0
    destination-prefix: 41.1.1.0
    filter-counters: 0
    output-counters: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security monitoring

Syntax show security monitoring

Release Information Command introduced in Junos OS Release 10.2.

Description Displays a count of security flow and central point (CP) sessions, CPU utilization (as a percentage of maximum), and memory in use (also as a percentage of maximum) at the moment the command is run.

Required Privilege Level View

Related Documentation

- *show security monitoring fpc fpc-number*
- *show security monitoring performance session*
- *show security monitoring performance spu*

show security monitoring

user@host>show security monitoring

```
user@host> show security monitoring
-------------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem         current        maximum        current        maximum
-------------------------------------------------------------------------------
  1   0   0  11               0              0              0              0
  1   1   0   5               3        6291456              1        7549747
  1   2   0   5               2        6291456              0        7549747
  1   3   0   5               3        6291456              1        7549747
  8   0   0  65               4           6963              2           8355
  8   1   0  65               2           6963              0           8355
Total Sessions:              14       18888294              4       22665951
```

show security monitoring (vSRX)

user@host>show security monitoring

```
user@host> show security monitoring
-------------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem         current        maximum        current        maximum
-------------------------------------------------------------------------------
  0   0   0  68               2         524288            N/A            N/A
```

show security monitoring (vSRX in a Chassis Cluster)

user@host>show security monitoring

```
node0:
-------------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem         current        maximum        current        maximum
```

```
--------------------------------------------------------------------------------
  0   0   0  67              0         524288            N/A            N/A

node1:
--------------------------------------------------------------------------------
                  Flow session   Flow session     CP session     CP session
FPC PIC CPU Mem         current        maximum        current        maximum
--------------------------------------------------------------------------------
  0   0   0  67              0         524288            N/A            N/A
```


Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 2174
show security policies policy-name p1 detail on page 2175
show security policies (services-offload) on page 2176
show security policies detail on page 2176
show security policies detail (TCP Options) on page 2177
show security policies policy-name p1 (Negated Address) on page 2177
show security policies policy-name p1 detail (Negated Address) on page 2178
show security policies global on page 2178

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 236: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 236: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 236: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```


Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                  545 bps
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                  272 bps
    Output bytes       :               18144                  545 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                 272 bps
    Input  packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Output packets     :                 216                   6 pps
      Initial direction:                 108                   3 bps
      Reply direction  :                 108                   3 bps
    Session rate       :                 108                   3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

show security policies hit-count

Syntax

show security policies hit-count
<ascending | descending>
<from-zone *zone-name*>
<greater-than *count*>
<less-than *count*>
<to-zone *zone-name*>

Release Information

Command introduced in Junos OS Release 12.1.

Description

Display the utility rate of security policies according to the number of hits they receive. The number of hits can be listed without an order or sorted in either ascending or descending order, and they can be restricted to the number of hits that fall above or below a specific count or within a range. Data is shown for all zones associated with the policies or named zones.

Use this command without options to display the number of hits in random order for all security policies and for all zones.

Options

- **ascending** | **descending**—(Optional) Display the number of hits for security policies in ascending or descending order.
- **from-zone** ***zone-name***—(Optional) Display the number of hits for security policies associated with the named source zone.
- **greater-than** ***count***—(Optional) Display security policies for which the number of hits is greater than the specified number.

  Range: 0 through 4,294,967,295
- **less-than** ***count***—(Optional) Display security policies for which the number of hits is less than the specified number.

  Range: 0 through 4,294,967,295
- **to-zone** ***zone-name***—(Optional) Display the number of hits for security policies associated with the named destination zone.

Required Privilege Level

view

Related Documentation

- clear security policies hit-count on page 1343
- Security Policies Overview on page 1085

List of Sample Output

show security policies hit-count on page 2180
show security policies hit-count ascending on page 2180
show security policies hit-count descending greater-than 70 less-than 100 on page 2180
show security policies hit-count from-zone untrust to-zone trust on page 2180

Output Fields

Table 104 lists the output fields for the **show security policies hit-count** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 237: show security policies hit-count Output Fields

Field Name	Field Description
from-zone	Name of the source zone.
to-zone	Name of the destination zone.
policy	Name of the security policy.
hit-count	Number of hits for each security policy.
Number of policy	Number of security policies for which hit counts are displayed.

Sample Output

show security policies hit-count

```
user@host> show security policies hit-count
from-zone   to-zone  policy  hit-count
untrust     vrtrust  u2t1      40
untrust     trust    u2t2      20
untrust     trust    u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count ascending

```
user@host> show security policies hit-count ascending
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      20
untrust     vrtrust  u2t1      40
untrust     trust    u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count descending greater-than 70 less-than 100

```
user@host> show security policies hit-count descending greater-than 70 less-than 100
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      100
untrust     vrtrust  u2t1      90
untrust     vrtrust  u2t3      80

Number of policy: 3
```

Sample Output

show security policies hit-count from-zone untrust to-zone trust

```
user@host> show security policies hit-count from-zone untrust to-zone trust
from-zone   to-zone  policy  hit-count
untrust     trust    u2t2      20
```


Copyright © 2016, Juniper Networks, Inc.

```
untrust     trust    u2t3      80

Number of policy: 2
```

Copyright © 2016, Juniper Networks, Inc.

show security resource-manager group active

Syntax

show security resource-manager group active
<group-number >
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description

Display security information about active groups created through the resource manager.

Options

- none—Display resource manager group service information for all active groups.
- ***group-number*** —(Optional) Display resource manager group service information for a specific group identification number.
- **node**—(Optional) For chassis cluster configurations, display active resource manager group service information on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output

show security resource-manager group active on page 2183
show security resource-manager group active 2048 on page 2183
show security resource-manager group active node primary on page 2183
show security resource-manager group active node all on page 2183
show security resource-manager group active 1024 node all on page 2183

Output Fields

Table 29 lists the output fields for the **show security resource-manager group** command. Output fields are listed in the approximate order in which they appear.

Table 238: show security resource-manager group Output Fields

Field Name	Field Description
Total groups	Total number of groups in the system.
active groups	Number of active groups.
Group ID	Identification number whose group information is displayed.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security resource-manager group active

```
user@host> show security resource-manager group active
Total groups 32, active groups 0
```

Sample Output

show security resource-manager group active 2048

```
user@host> show security resource-manager group active 2048
Total groups 2048, active groups 1
Group ID 2048: state - Active
          : Virtual System      - root
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8190
               Resource ID - 8188
               Resource ID - 8187
```

Sample Output

show security resource-manager group active node primary

```
user@host> show security resource-manager group active node primary
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active node all

```
user@host> show security resource-manager group active node all
node0:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
node1:
--------------------------------------------------------------------------
Group ID 1024: Application - SIP ALG
Total groups 1024, active groups 1
```

Sample Output

show security resource-manager group active 1024 node all

```
user@host> show security resource-manager group active 1024 node all
node0:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8192
               Resource ID - 8188
               Resource ID - 8187
```

Copyright © 2016, Juniper Networks, Inc.

```
node1:
--------------------------------------------------------------------------
Group ID 1024: state - Active
          : Application         - SIP ALG
          : Group Timeout       - 65535
          : Number of resources - 3
               Resource ID - 8187
               Resource ID - 8186
               Resource ID - 8190
```


Copyright © 2016, Juniper Networks, Inc.

show security resource-manager resource active

Syntax
show security resource-manager resource **active**
<resource-id >
<node (*node-id* | all | local | primary)>

Release Information
Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description
Display security information about active resources created through the resource manager.

Options
- none—Display information for all active resources.
- ***resource-id*** —(Optional) Display information for a resource with a specific identification number.
- **node**—(Optional) For chassis cluster configurations, display active resource manager information on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level
view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security resource-manager resource active on page 2186
show security resource-manager resource active 5 on page 2186
show security resource-manager resource active node local on page 2186
show security resource-manager resource active node primary on page 2186

Output Fields
Table 30 lists the output fields for the **show security resource-manager resource** command. Output fields are listed in the approximate order in which they appear.

Table 239: show security resource-manager resource Output Fields

Field Name	Field Description
Total resources	Total number of resources in the system.
active resources	Number of active resources.
Resource ID	Identification number whose resource information is displayed.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security resource-manager resource active

```
user@host> show security resource-manager resource active
    Resource ID 7: Group ID - 2, Application - JSF_sip

    Resource ID 6: Group ID - 2, Application - JSF_sip

    Resource ID 5: Group ID - 2, Application - JSF_sip

    Resource ID 4: Group ID - 2, Application - JSF_sip

    Resource ID 3: Group ID - 2, Application - JSF_sip

    Resource ID 1: Group ID - 2, Application - JSF_sip

    Resource ID 2: Group ID - 2, Application - JSF_sip
Total Resources 4326, active resources 7
```

Sample Output

show security resource-manager resource active 5

```
user@host> show security resource-manager resource active 5
Resource ID 5: state - Active
    Application         - asl_client
    Parent group        - 2
    Policy              - 5
    From zone           - untrust
    To zone             - trust
    Resource timeout    - 0
    Number of sessions  - 0
    Number of Holes     - 1
        Source IP range        - {0.0.0.0, 0.0.0.0}
        Source port range      - {0, 0}
        Destination IP range   - {33.1.0.200, 33.1.0.200}
        Destination port range - {5060, 5060}
        Translated             - {0.0.0.0/0 -> 33.1.0.200/5060}
        Protocol               - 17
        Reference count        - 1
```

Sample Output

show security resource-manager resource active node local

```
user@host> show security resource-manager resource active node local
node0:
--------------------------------------------------------------------------
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Sample Output

show security resource-manager resource active node primary

```
user@host> show security resource-manager resource active node primary
node0:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

```
    Resource ID 8192: Group ID - 1024, Application - SIP ALG
    Resource ID 8188: Group ID - 1024, Application - SIP ALG
    Resource ID 8187: Group ID - 1024, Application - SIP ALG
Total Resources 8192, active resources 3
```

Copyright © 2016, Juniper Networks, Inc.

show security resource-manager settings

Syntax show security resource-manager settings
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5; **node** options added in Junos OS Release 9.0.

Description Display resource manager settings.

Options **node**—(Optional) For chassis cluster configurations, display resource manager settings on a specific node.

- ***node-id*** —Identification number of the node. It can be 0 or 1.
- **all**—Display information about all nodes.
- **local**—Display information about the local node.
- **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security resource-manager settings on page 2188
show security resource-manager settings node primary on page 2189
show security resource-manager settings node all on page 2189

Output Fields Table 240 lists the output fields for the **show security resource-manager settings** command. Output fields are listed in the approximate order in which they appear.

Table 240: show security resource-manager settings Output Fields

Field Name	Field Description
Client Heartbeat	Time after which idle an resource manager client is timed out.
Count	Number of active clients.
Pinhole age	Duration for which the temporary opening in the security firewall (pinhole) is open for specified traffic. If the specified traffic does not exist during this time period, the pinhole is timed out.

Sample Output

show security resource-manager settings

```
user@host> show security resource-manager settings
Client Heartbeat: timeout 600 seconds, count 5
Pinhole age: 32 seconds
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security resource-manager settings node primary

```
user@host> show security resource-manager settings node primary
node0:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
```

Sample Output

show security resource-manager settings node all

```
user@host> show security resource-manager settings node all
node0:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
node1:
--------------------------------------------------------------------------
Client heartbeat: timeout 600 seconds, count 5
Pinhole age: 120 seconds
```

Copyright © 2016, Juniper Networks, Inc.

show security resource-manager summary

Syntax show security resource-manager summary

Release Information Command introduced in Junos OS Release 11.4.

Description Display summary information about active resources, clients, groups, and sessions created through the resource manager.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output show security resource-manager summary on page 2190

Output Fields Table 31 lists the output fields for the **show security resource-manager summary** command. Output fields are listed in the approximate order in which they appear.

Table 241: show security resource-manager summary Output Fields

Field Name	Field Description
Active resource-manager clients	Number of active resource manager clients.
Active resource-manager groups	Number of active resource manager groups.
Active resource-manager resources	Number of active resource manager resources.
Active resource-manager sessions	Number of active resource manager sessions.

Sample Output

show security resource-manager summary

```
user@host> show security resource-manager summary

Active resource-manager clients   : 15
Active resource-manager groups    : 1
Active resource-manager resources : 1
Active resource-manager sessions  : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security screen ids-option

Syntax

show security screen ids-option
screen-name
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5. Support for **node** options added in Junos OS Release 9.0. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10. Support for UDP **port scan** added in Junos OS Release 12.1X47-D10.

Description

Display configuration information about the specified security screen.

Options

- *screen-name* —Name of the screen.
- **node**—(Optional) For chassis cluster configurations, display the configuration status of the security screen on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- ids-option on page 990
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

show security screen ids-option jscreen on page 2193
show security screen ids-option jscreen (IPv6) on page 2194
show security screen ids-option jscreen1 node all on page 2194

Output Fields

Table 68 lists the output fields for the **show security screen ids-option** command. Output fields are listed in the approximate order in which they appear.

Table 242: show security screen ids-option Output Fields

Field Name	Field Description
TCP address sweep threshold	Number of microseconds for which the device accepts 10 TCP packets from the same remote source to different destination addresses.
TCP port scan threshold	Number of microseconds during which the device accepts packets from the same remote source with up to 10 different port numbers.
ICMP address sweep threshold	Maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device.

Copyright © 2016, Juniper Networks, Inc.

Table 242: show security screen ids-option Output Fields *(continued)*

Field Name	Field Description
UDP flood threshold	Number of UDP packets per second allowed to ping the same destination address before the device rejects further UDP packets.
UDP port scan threshold	Number of microseconds during which the device accepts packets from the same remote source IP with up to 10 different destination port numbers.
TCP winnuke	Enable or disable the detection of TCP WinNuke attacks.
TCP SYN flood attack threshold	Number of SYN packets per second required to trigger the SYN proxy response.
TCP SYN flood alarm threshold	Number of half-complete proxy connections per second at which the device makes entries in the event alarm log.
TCP SYN flood source threshold	Number of SYN segments to be received per second before the device begins dropping connection requests.
TCP SYN flood destination threshold	Number of SYN segments received per second before the device begins dropping connection requests.
TCP SYN flood timeout	Maximum length of time before a half-completed connection is dropped from the queue.
TCP SYN flood queue size	Number of proxy connection requests that can be held in the proxy connection queue before the device begins rejecting new connection requests.
ICMP large packet	Enable or disable the detection of any ICMP frame with an IP length greater than 1024 bytes.
UDP address sweep threshold	Number of microseconds for which the device accepts 10 UDP packets from the same remote source to different destination addresses.
IPv6 extension routing	Enable or disable the IPv6 extension routing screen option.
IPv6 extension shim6	Enable or disable the IPv6 extension shim6 screen option.
IPv6 extension fragment	Enable or disable the IPv6 extension fragment screen option.
IPv6 extension AH	Enable or disable the IPv6 extension Authentication Header Protocol screen option.
IPv6 extension ESP	Enable or disable the IPv6 extension Encapsulating Security Payload screen option.
IPv6 extension mobility	Enable or disable the IPv6 extension mobility screen option.
IPv6 extension HIP	Enable or disable the IPv6 extension Host Identify Protocol screen option.
IPv6 extension no next	Enable or disable the IPv6 extension no-next screen option.
IPv6 extension user-defined	Enable or disable the IPv6 extension user-defined screen option.

Copyright © 2016, Juniper Networks, Inc.

Table 242: show security screen ids-option Output Fields *(continued)*

Field Name	Field Description
IPv6 extension HbyH jumbo	Enable or disable the IPv6 extension HbyH jumbo screen option.
IPv6 extension HbyH RPL	Enable or disable the IPv6 extension HbyH RPL screen option.
IPv6 extension HbyH router alert	Enable or disable the IPv6 extension HbyH router screen option.
IPv6 extension HbyH quick start	Enable or disable the IPv6 extension HbyH quick-start screen option.
IPv6 extension HbyH CALIPSO	Enable or disable the IPv6 extension HbyH Common Architecture Label IPv6 Security Screen option.
IPv6 extension HbyH SMF DPD	Enable or disable the IPv6 extension HbyH Simplified Multicast Forwarding IPv6 Duplicate Packet Detection screen option.
IPv6 extension HbyH user-defined	Enable or disable the IPv6 extension HbyH user-defined screen option.
IPv6 extension Dst tunnel encap limit	Enable or disable the IPv6 extension distributed (network) storage tunnel encapsulation limit screen option.
IPv6 extension Dst home address	Enable or disable the IPv6 extension DST home address screen option.
IPv6 extension Dst ILNP nonce	Enable or disable the IPv6 extension DST Identifier-Locator Network Protocol nonce screen option.
IPv6 extension Dst line-id	Enable or disable the IPv6 extension DST line-ID screen option.
IPv6 extension Dst user-defined	Enable or disable the IPv6 extension DST user-defined screen option.
IPv6 extension header limit	Threshold for the number of IPv6 extension headers that can pass through the screen.
IPv6 malformed header	Enable or disable the IPv6 malformed header screen option.
ICMPv6 malformed header	Enable or disable the ICMPv6 malformed packet screen option.

Sample Output

show security screen ids-option jscreen

```
user@host> show security screen ids-option jscreen
Screen object status:
Name                                        Value
  TCP port scan threshold                   5000
  UDP port scan threshold                   10000
  ICMP address sweep threshold              5000
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security screen ids-option jscreen (IPv6)

```
user@host> show security screen ids-option jscreen

Screen object status:

   Name                                        Value
  ICMP ping of death                          enabled
  ……

  IPv6 extension routing                      enabled
  IPv6 extension shim6                        enabled
  IPv6 extension fragment                     enabled
  IPv6 extension AH                           enabled
  IPv6 extension ESP                          enabled
  IPv6 extension mobility                     enabled
  IPv6 extension HIP                          enabled
  IPv6 extension no next                      enabled
  IPv6 extension user-defined                 enabled
  IPv6 extension HbyH jumbo                   enabled
  IPv6 extension HbyH RPL                     enabled
  IPv6 extension HbyH router alert            enabled
  IPv6 extension HbyH quick start             enabled
  IPv6 extension HbyH CALIPSO                 enabled
  IPv6 extension HbyH SMF DPD                 enabled
  IPv6 extension HbyH user-defined            enabled
  IPv6 extension Dst tunnel encap limit       enabled
  IPv6 extension Dst home address             enabled
  IPv6 extension Dst ILNP nonce               enabled
  IPv6 extension Dst line-id                  enabled
  IPv6 extension Dst user-defined             enabled
  IPv6 extension header limit                 20
  IPv6 Malformed header                       enabled
  ICMPv6 malformed packet                     enabled
```

Sample Output

show security screen ids-option jscreen1 node all

```
user@host> show security screen ids-option jscreen1 node all
node0:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  UDP flood threshold                       1000
  TCP winnuke                               enabled
  TCP SYN flood attack threshold            200
  TCP SYN flood alarm threshold             512
  TCP SYN flood source threshold            4000
  TCP SYN flood destination threshold       4000
  TCP SYN flood timeout                     20
  TCP SYN flood queue size                  1024
  ICMP large packet                         enabled
node1:
--------------------------------------------------------------------------
Screen object status:
Name                                        Value
  UDP flood threshold                       1000
```


Copyright © 2016, Juniper Networks, Inc.

```
TCP winnuke                                enabled
TCP SYN flood attack threshold             200
TCP SYN flood alarm threshold              512
TCP SYN flood source threshold             4000
TCP SYN flood destination threshold        4000
TCP SYN flood timeout                      20
TCP SYN flood queue size                   1024
ICMP large packet                          enabled
```

Copyright © 2016, Juniper Networks, Inc.

show security screen statistics

Syntax

show security screen statistics (zone *zone-name* | interface *interface-name*)
<logical-system (*logical-system-name* | all)>
<node (*node-id* | all | local | primary)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 8.5. **node** options added in Junos OS Release 9.0. **logical-system all** option added in Junos OS Release 11.2R6. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

Description

Display intrusion detection service (IDS) security screen statistics.

Options

- **zone** ***zone-name*** —Display screen statistics for this security zone.
- **interface** *interface-name*—Display screen statistics for this interface.
- **logical-system**—(Optional) Display screen statistics for configured logical systems.
  - ***logical-system-name***—Display screen statistics for the named logical system.
  - **all**—Display screen statistics for all logical systems, including the master (root) logical system.
- **node**—(Optional) For chassis cluster configurations, display screen statistics on a specific node.
  - ***node-id***—Identification number of a node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.
- **root-logical-system**—(Optional) Display screen statistics for the master logical system only.

Required Privilege Level

view

Related Documentation

- clear security screen statistics on page 1022
- clear security screen statistics interface on page 1023
- clear security screen statistics zone on page 1025
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

show security screen statistics zone scrzone on page 2199
show security screen statistics zone untrust (IPv6) on page 2199
show security screen statistics interface ge-0/0/3 on page 2200
show security screen statistics interface ge-0/0/1 (IPv6) on page 2200
show security screen statistics interface ge-0/0/1 node primary on page 2201
show security screen statistics zone trust logical-system all on page 2201

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 69 lists the output fields for the **show security screen statistics** command. Output fields are listed in the approximate order in which they appear.

Table 243: show security screen statistics Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.

Copyright © 2016, Juniper Networks, Inc.

Table 243: show security screen statistics Output Fields *(continued)*

Field Name	Field Description
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.
UDP address sweep	Number of UDP address sweeps.
IPv6 extension header	Number of packets filtered for the defined IPv6 extension headers.
IPv6 extension hop by hop option	Number of packets filtered for the defined IPv6 hop-by-hop option types.
IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).

Copyright © 2016, Juniper Networks, Inc.

Table 243: show security screen statistics Output Fields *(continued)*

ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                            Statistics
  ICMP flood                                    0
  UDP  flood                                    0
  TCP  winnuke                                  0
  ……
  IPv6   extension header                        0
  IPv6   extension hop by hop option             0
  IPv6   extension destination option            0
  IPv6   extension header limit                  0
  IPv6   malformed header                        0
```

Copyright © 2016, Juniper Networks, Inc.

```
  ICMPv6 malformed packet                    0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                                 Statistics
  ICMP flood                                        0
  UDP flood                                         0
  ……
  IPv6 extension header                             0
  IPv6 extension hop by hop option                  0
  IPv6 extension destination option                 0
  IPv6 extension header limit                       0
  IPv6 malformed header                             0
  ICMPv6 malformed packet                           0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 1
  UDP flood                                  1
  TCP winnuke                                1
  TCP port scan                              1
  ICMP address sweep                         1
  TCP sweep                                  1
  UDP sweep                                  1
  IP tear drop                               1
  TCP SYN flood                              1
  IP spoofing                                1
  ICMP ping of death                         1
  IP source route option                     1
  TCP land attack                            1
  TCP SYN fragment                           1
  TCP no flag                                1
  IP unknown protocol                        1
  IP bad options                             1
  IP record route option                     1
  IP timestamp option                        1
  IP security option                         1
  IP loose source route option               1
  IP strict source route option              1
  IP stream option                           1
  ICMP fragment                              1
  ICMP large packet                          1
  TCP SYN FIN                                1
  TCP FIN no ACK                             1
  Source session limit                       1
  TCP SYN-ACK-ACK proxy                      1
  IP block fragment                          1
  Destination session limit                  1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
```

Copyright © 2016, Juniper Networks, Inc.

```
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls1
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls2
Screen statistics:

IDS attack type                             Statistics
```


Copyright © 2016, Juniper Networks, Inc.

```
ICMP flood                          0
UDP flood                           0
TCP winnuke                         0
TCP port scan                       0
ICMP address sweep                  0
TCP sweep                           0
UDP sweep                           0
IP tear drop                        0
TCP SYN flood                       0
IP spoofing                         0
ICMP ping of death                  0
IP source route option              0
TCP land attack                     0
TCP SYN fragment                    0
TCP no flag                         0
IP unknown protocol                 0
IP bad options                      0
IP record route option              0
IP timestamp option                 0
IP security option                  0
IP loose source route option        0
IP strict source route option       0
IP stream option                    0
ICMP fragment                       0
ICMP large packet                   0
TCP SYN FIN                         0
TCP FIN no ACK                      0
Source session limit                0
TCP SYN-ACK-ACK proxy               0
IP block fragment                   0
Destination session limit           0
```

Copyright © 2016, Juniper Networks, Inc.

show security softwires

Syntax show security softwires <softwire-name *softwire-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced in Junos OS Release 10.4. The **logical-system** option introduced in Junos OS Release 12.1.

Description Display a summary of information of all the softwire concentrators and details on concentrators with specified name.

Options **softwire-name *softwire-name***—Display the details of the specified softwire concentrator.

logical-system (all | *logical-system-name*)—Display softwire information for all logical systems or for a specified logical system. This option is only available to the master administrator.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Sample Output

```
user@host> show security softwires
Software Name             SC Address         Status   Number of SI connected
  SC-CSSI-1                  3001::1          Connected      2
  SC-CSSI-str00              3100::1          Active         0
  SC-CSSI-str01              3101::1          Inactive       0
  SC-CSSI-str02              3001::1          Connected      2520
```

```
user@host> show security softwires softwire-name SC-CSSI-1
Name of softwire: SC-CSSI-1
   SC status: Connected
   SC address: 3001::1
   Zone: trust
   VR ID: 0
 SI Address               SI Status                  SPU
 3001::2                  Active                     spu-1
 3001::2                  Active                     spu-21
 SI number: 2
```

```
user@host> show security softwires logical-system ls-product-design
Software Name             SC Address         Status   Number of SI connected
  sc_1                       3000::1         Connected      1
```


Copyright © 2016, Juniper Networks, Inc.

show security zones

Syntax
show security zones
<detail | terse>
< *zone-name* >

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones.

Options
- none—Display information about all zones.
- **detail | terse**—(Optional) Display the specified level of output.
- ***zone-name*** —(Optional) Display information about the specified zone.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones on page 2206
show security zones abc on page 2206
show security zones abc detail on page 2206
show security zones terse on page 2207

Output Fields
Table 32 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 244: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                     Type
my-internal              Security
my-external              Security
dmz                      Security
```

Copyright © 2016, Juniper Networks, Inc.

show security zones type

Syntax
show security zones type
(functional | security)
<detail | terse>

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones of the specified type.

Options
- **functional**—Display functional zones.
- **security**—Display security zones.
- **detail | terse**—(Optional) Display the specified level of output.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones type functional on page 2209
show security zones type security on page 2209
show security zones type security terse on page 2209
show security zones type security detail on page 2209

Output Fields
Table 33 lists the output fields for the **show security zones type** command. Output fields are listed in the approximate order in which they appear.

Table 245: show security zones type Output Fields

Field Name	Field Description
Security zone	Zone name.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones type functional

```
user@host> show security zones type functional
Functional zone: management
  Description: management zone
  Policy configurable: No
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security

```
user@host> show security zones type security
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Sample Output

show security zones type security terse

```
user@host> show security zones type security terse
Zone                    Type
trust                   Security
untrust                 Security
junos-host              Security
```

Sample Output

show security zones type security detail

```
user@host> show security zones type security detail
Security zone: trust
  Description: trust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
Security zone: untrust
  Description: untrust zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
  ge-0/0/1.0
Security zone: junos-host
  Description: junos-host zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```


Copyright © 2016, Juniper Networks, Inc.

General Packet Radio Service Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 29

Overview

- Introduction to General Packet Radio Service on page 2215

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 95

Introduction to General Packet Radio Service

- GPRS Overview on page 2215
- Understanding Central Point Architecture Support for GTP on page 2218
- Understanding GTP Handover Messages on page 2219

GPRS Overview

General Packet Radio Service (GPRS) networks connect to several external networks including those of roaming partners, corporate customers, GPRS Roaming Exchange (GRX) providers, and the public Internet. GPRS network operators face the challenge of protecting their network while providing and controlling access to and from these external networks. Juniper Networks provides solutions to many of the security problems plaguing GPRS network operators.

In the GPRS architecture, the fundamental cause of security threats to an operator's network is the inherent lack of security in the GPRS tunneling protocol (GTP). GTP is the protocol used between GPRS support nodes (GSNs). GTP is used to establish a GTP tunnel for individual mobile stations (MSs) and between a Serving GPRS Support Node (SGSN) and a Gateway GPRS Support Node (GGSN). A GTP tunnel is a channel between GSNs through which two hosts exchange data. The SGSN receives packets from the MS and encapsulates them within a GTP header before forwarding them to the GGSN through the GTP tunnel. When the GGSN receives the packets, it decapsulates them and forwards them to the external host.

Communication between different GPRS networks is not secure because GTP does not provide any authentication, data integrity, or confidentiality protection. Implementing IP Security (IPsec) for connections between roaming partners, setting traffic rate limits, and using stateful inspection can eliminate a majority of the GTP's security risks. The GTP firewall features in Junos OS address key security issues in mobile operators' networks.

Juniper Networks security devices mitigate a wide variety of attacks on the following types of GPRS interfaces:

- Gn—The Gn interface is the connection between an SGSN and a GGSN within the same public land mobile network (PLMN).

Copyright © 2016, Juniper Networks, Inc.

- Gp—The Gp interface is the connection between two PLMNs.
- Gi—The Gi interface is the connection between a GGSN and the Internet or destination networks connected to a PLMN.

NOTE: The term *interface* has different meanings in Junos OS and in GPRS technology. In Junos OS, an interface is a doorway to a security zone that allows traffic to enter and exit the zone. In GPRS, an interface is a connection, or a reference point, between two components of a GPRS infrastructure, for example, an SGSN and a GGSN.

This topic contains the following sections:

- Gp and Gn Interfaces on page 2216
- Gi Interface on page 2216
- Operational Modes on page 2217
- GTP In-Service Software Upgrade on page 2218

Gp and Gn Interfaces

You implement a security device on the Gn interface to protect core network assets such as the SGSN and GGSN. To secure GTP tunnels on the Gn interface, you place the security device between SGSNs and GGSNs within a common PLMN.

When you implement a security device to the Gp interface, you protect a PLMN from another PLMN. To secure GTP tunnels on the Gp interface, you place the SGSNs and GGSNs of a PLMN behind the security device so that all traffic, incoming and outgoing, goes through the firewall.

Figure 99 illustrates the placement of Juniper Networks SRX Series devices used to protect PLMNs on the Gp and Gn interfaces.

Figure 99: Gp and Gn Interfaces

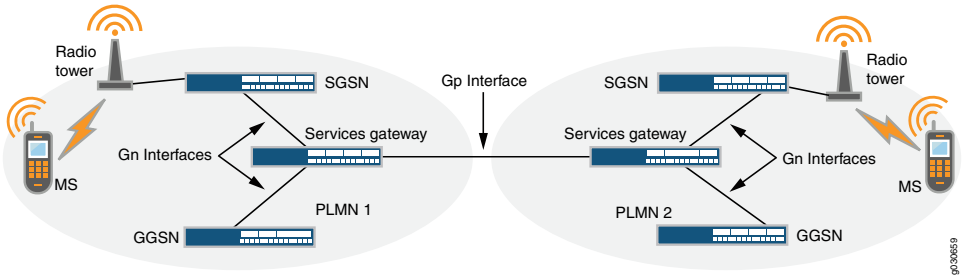

Gi Interface

When you implement a security device on the Gi interface, you can simultaneously control traffic for multiple networks, protect a PLMN against the Internet and external networks, and protect mobile users from the Internet and other networks. Junos OS provides a great

Copyright © 2016, Juniper Networks, Inc.

number of virtual routers, making it possible for you to use one virtual router per customer network and thereby allow the separation of traffic for each customer network.

The security device can securely forward packets to the Internet or destination networks using the Layer 2 Tunneling Protocol (L2TP) for IPsec virtual private network (VPN) tunnels.

NOTE: SRX Series devices do not support full L2TP.

Figure 100 illustrates the implementation of a security device to protect a PLMN on the Gi interface.

Figure 100: Gi Interface

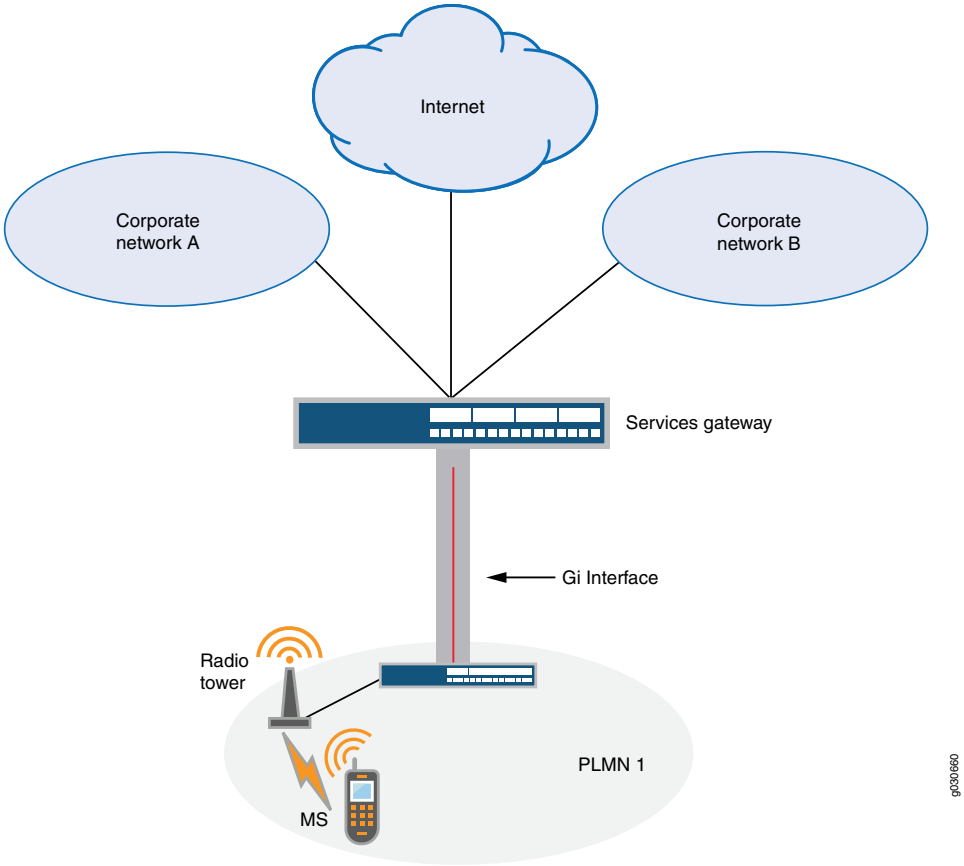

Operational Modes

Junos OS supports two interface operational modes with GTP: transparent mode and route mode. If you want the security device to participate in the routing infrastructure of your network, you can run it in route mode. This requires a certain amount of network redesign. Alternatively, you can implement the security device into your existing network in transparent mode without having to reconfigure the entire network. In transparent mode, the security device functions as a Layer 2 switch or bridge, and the IP addresses

Copyright © 2016, Juniper Networks, Inc.

of interfaces are set at 0.0.0.0, making the presence of the security device invisible, or *transparent*, to users.

Junos OS supports Network Address Translation (NAT) on interfaces and policies that do not have GTP inspection enabled.

Currently in Junos OS, route mode supports active/passive, and active/active chassis cluster. Transparent mode supports active/passive only.

GTP In-Service Software Upgrade

GTP supports unified in-service software upgrade (ISSU) between two SRX Series devices running two different Junos OS releases. Unified ISSU is performed on a chassis cluster, enabling a software upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

Related Documentation

- *Layer 2 Bridging and Transparent Mode for Security Devices*
- *Chassis Cluster Overview*
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Inspection Objects on page 2229
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235

Understanding Central Point Architecture Support for GTP

User equipment (for example, a cellphone) attaches to a Serving GPRS Support Node (SGSN) for General Packet Radio Service (GPRS) data service. The SGSN connects to a gateway GPRS support node to access the Internet. The user equipment requests the SGSN to create one or multiple GPRS tunneling protocol (GTP) tunnels to the GGSN for Internet access. In situations where the user equipment moves to a new location, the user equipment has to attach to another SGSN. The new SGSN notifies the GGSN to update the new SGSN information in the original tunnel.

The GTP Application Layer Gateway (ALG) maintains the status of the tunnels and permits tunnel update request packets only for the existing tunnels. When the user equipment moves to a new location and attaches to another SGSN, the new SGSN information must be updated in the original tunnel. Because few GTP-C messages are bidirectional, and messages can be sent either sent by the SGSN or the GGSN, correct session distribution is not guaranteed. That is, the GTP ALG stops creating a session if the first packet originates from an unknown direction. In this case, the first packet and the other pending packets are dropped.

To prevent GTP-C packets from being dropped, a new flow session is created and the GTP-C traffic is allowed to pass even if the GGSN or SGSN direction is not determined. Later, the GGSN IP is determined using the correct SPU to create the flow session; otherwise, the session is migrated to the designated SPU.

Copyright © 2016, Juniper Networks, Inc.

Starting with Junos OS Release 15.1X49-D40, the central point architecture is enhanced to address the following issues:

- To prevent GTP-C packet drop issues during the SGSN handover.
- To support the GTP-C message rate-limiting to protect the GGSN from flooding of GTP-C messages.
- To distribute GTP-U traffic handled by a GGSN and SGSN pair on all SPUs by switching to tunnel-based session distribution in which the GTP-U traffic of different tunnels is spread across different SPUs. Use the **enable-gtpu-distribution** command to enable GTP-U session distribution.

Related Documentation

- enable-gtpu-distribution on page 2362

Understanding GTP Handover Messages

During handover procedures, Serving GPRS Support Node (SGSN) context messages (request, response, and acknowledge) or forward relocation messages are sent between the new and the old mobility management entity (MME) and SGSN. For GPRS tunneling protocol (GTP) version 2, the messages should be context messages or forward relocation messages. For simplicity, these types of messages are uniformly referred as handover messages. The packet data protocol (PDP) context information is acquired from these messages. The PDP context is set up on the SRX Series device when these messages are received, and then subsequent GTP messages can be normally inspected according to the new PDP context.

Use the **set security gprs gtp profile *<profile-name>* handover-on-roaming-intf** command to enable PDP context setup by handover messages. Use the **delete security gprs gtp profile *<profile-name>* handover-on-roaming-intf** command to disable PDP context setup by handover messages.

The addresses and tunnel endpoint identifiers (TEIDs) for forwarding data traffic are also acquired from handover messages. In addition, the forward tunnel can be set up on SRX Series devices for forwarding GPRS tunneling protocol, user plane (GTP-U) stateful check.

NOTE: **Handover between different GTP versions is supported.**

Key features of GTP handover are:

- Support for GTP inter-MME/SGSN handover messages for GTPv0, v1, and v2
- Inter-MME/SGSN handover messages inspection
- GTP PDP context and forwarding tunnel setup according to the information in handover messages
- GTP-U inspection for forwarding data traffic

Copyright © 2016, Juniper Networks, Inc.

- Support for PDP context update by updating and modifying messages with different versions
- System log and counter for handover messages

Related Documentation

- Understanding GTP Message Filtering on page 2233

Copyright © 2016, Juniper Networks, Inc.

PART 30

Configuring GPRS Tunnel Protocol v1

- Configuring Policy-Based GTP on page 2223
- Configuring GTP Inspection Objects on page 2229
- Configuring GTP Message Filtering on page 2233
- Configuring GTP Information Elements on page 2247
- Configuring NAT for GTP on page 2265
- Configuring GGSN on page 2279

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 96

Configuring Policy-Based GTP

- Understanding Policy-Based GTP on page 2223
- Example: Enabling GTP Inspection in Policies on page 2224

Understanding Policy-Based GTP

By default, the public land mobile network (PLMN) that the Juniper Networks device protects is in the Trust zone. The device protects the PLMN in the Trust zone against other PLMNs in other zones. You can place all the PLMNs against which you are protecting your PLMN in the Untrust zone, or you can create user-defined zones for each PLMN. A PLMN can occupy one security zone or multiple security zones.

You must create policies to enable traffic to flow between zones and PLMNs. Policies contain rules that permit, deny, or tunnel traffic. The device performs GPRS tunneling protocol (GTP) policy filtering by checking every GTP packet against policies that regulate GTP traffic and by then forwarding, dropping, or tunneling the packet based on these policies.

By selecting the GTP service in a policy, you enable the device to permit, deny, or tunnel GTP traffic. However, this does not enable the device to inspect GTP traffic. For the device to inspect GTP traffic, you must apply a GTP configuration, also referred to as a *GTP inspection object*, to a policy.

You can apply only one GTP inspection object per policy, but you can apply a GTP inspection object to multiple policies. Using policies, you can permit or deny the establishment of GTP tunnels from certain peers such as a Serving GPRS Support Node (SGSN).

You can configure policies that specify “Any” as the source or destination zone (thereby including all hosts in the zone), and you can configure policies that specify multiple source and destination addresses.

In policies, you can enable traffic logging.

Related Documentation

- GPRS Overview on page 2215
- Understanding GTP Inspection Objects on page 2229
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

- Example: Enabling GTP Inspection in Policies on page 2224

Example: Enabling GTP Inspection in Policies

This example shows how to enable GTP inspection in policies.

- Requirements on page 2224
- Overview on page 2224
- Configuration on page 2224
- Verification on page 2227

Requirements

Before you begin, the device must be restarted after GTP is enabled. By default, GTP is disabled on the device.

Overview

In this example, you configure interfaces as ge-0/0/1 and ge-0/0/2, the addresses are 2.0.0.254/8 and 3.0.0.254/8. You then configure the security zone and specify address as 2.0.0.5/32 and 3.0.0.6/32. You enable the GTP service in the security policies to allow bidirectional traffic between two networks within the same PLMN.

Configuration

CLI Quick Configuration

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

```
set security gprs gtp profile gtp1
set interfaces ge-0/0/1 unit 0 family inet address 2.0.0.254/8
set interfaces ge-0/0/2 unit 0 family inet address 3.0.0.254/8
set security zones security-zone sgsn interfaces ge-0/0/1.0 host-inbound-traffic
    system-services all
set security zones security-zone sgsn host-inbound-traffic protocols all
set security zones security-zone ggsn interfaces ge-0/0/2.0 host-inbound-traffic
    system-services all
set security zones security-zone ggsn host-inbound-traffic protocols all
set security address-book global address local-sgsn 2.0.0.5/32
set security address-book global address remote-ggsn 3.0.0.6/32
set security policies from-zone sgsn to-zone ggsn policy sgsn_to_ggsn match
    source-address local-sgsn destination-address remote-ggsn application junos-gprs-gtp
set security policies from-zone sgsn to-zone ggsn policy sgsn_to_ggsn then permit
    application-services gprs-gtp-profile gtp1
set security policies from-zone ggsn to-zone sgsn policy ggsn_to_sgsn match
    source-address remote-ggsn destination-address local-sgsn application junos-gprs-gtp
set security policies from-zone ggsn to-zone sgsn policy ggsn_to_sgsn then permit
    application-services gprs-gtp-profile gtp1
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To configure GTP inspection in policies:

1. Enable GTP.

```
[edit]
user@host# set security gprs gtp enable
user@host# commit
user@host# exit
user@host# request system reboot
```

2. Create the GTP inspection object.

```
[edit]
user@host# set security gprs gtp profile gtp1
```

3. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 2.0.0.254/8
user@host# set ge-0/0/2 unit 0 family inet address 3.0.0.254/8
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone sgsn interfaces ge-0/0/1.0
user@host# set security-zone sgsn host-inbound-traffic system-services all
user@host# set security-zone sgsn host-inbound-traffic protocols all
user@host# set security-zone ggsn interfaces ge-0/0/2.0
user@host# set security-zone ggsn host-inbound-traffic system-services all
user@host# set security-zone ggsn host-inbound-traffic protocols all
```

5. Specify addresses.

```
[edit security address-book global]
user@host# set address local-sgsn 2.0.0.5/32
user@host# set address remote-ggsn 3.0.0.6/32
```

6. Enable the GTP service in the security policies.

```
[edit security policies]
user@host# set from-zone sgsn to-zone ggsn policy sgsn_to_ggsn match
  source-address local-sgsn destination-address remote-ggsn application
  junos-gprs-gtp
user@host# set from-zone sgsn to-zone ggsn policy sgsn_to_ggsn then permit
  application-services gprs-gtp-profile gtp1
user@host# set from-zone ggsn to-zone sgsn policy ggsn_to_sgsn match
  source-address remote-ggsn destination-address local-sgsn application
  junos-gprs-gtp
user@host# set from-zone ggsn to-zone sgsn policy ggsn_to_sgsn then permit
  application-services gprs-gtp-profile gtp1
```

Results

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

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

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security
...
gprs {
gtp {
profile gtp1;
  }
}
zones {
security-zone Trust {
host-inbound-traffic {
system-services {
all;
}
protocols {
all;
  }
}
interfaces {
ge-0/0/1.0;
  }
}
...

host-inbound-traffic {
system-services {
all;
}
protocols {
all;
  }
}
interfaces {
ge-0/0/1.0;
  }
}
host-inbound-traffic {
system-services {
all;
}
protocols {
all;
  }
}
interfaces {
ge-0/0/2.0;
   }
}
}
address-book {
global {
address local-sgsn 2.0.0.5/32;
address remote-ggsn 3.0.0.6/32;
   }
}
policies {
from-zone sgsn to-zone ggsn {
policy sgsn_to_ggsn {
match {
source-address local-sgsn;
destination-address remote-ggsn;
application junos-gprs-gtp;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
then {
permit {
application-services {
gprs-gtp-profile gtp1;
   }
   }
  }
  }
}
from-zone ggsn to-zone sgsn {
policy ggsn_to_sgsn {
match {
source-address remote-ggsn;
destination-address local-sgsn;
application junos-gprs-gtp;
}
}
then {
permit {
application-services {
gprs-gtp-profile gtp1;
     }
     }
  }
  }
default-policy {
permit-all;
}
}
...
```

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

Verification

Confirm that the configuration is working properly.

Verifying GTP Inspection in Policies

Purpose Verify that GTP inspection is enabled.

Action From operational mode, enter the **show security** command.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 97

Configuring GTP Inspection Objects

- Understanding GTP Inspection Objects on page 2229
- Example: Creating a GTP Inspection Object on page 2229
- Understanding GTP-U Inspection on page 2230

Understanding GTP Inspection Objects

For the device to perform the inspection of GPRS tunneling protocol (GTP) traffic, you must create a GTP inspection object and then apply it to a policy. GTP inspection objects provide more flexibility in that they allow you to configure multiple policies that enforce different GTP configurations. You can configure the device to control GTP traffic differently based on source and destination zones and addresses, action, and so on.

To configure GTP features, you must enter the context of a GTP configuration. To save your settings in the CLI, you must first exit the GTP configuration, then enter the **commit** command.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235
- Example: Creating a GTP Inspection Object on page 2229

Example: Creating a GTP Inspection Object

This example shows how to create a GTP inspection object.

- Requirements on page 2230
- Overview on page 2230
- Configuration on page 2230
- Verification on page 2230

Copyright © 2016, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you create a GTP inspection object named LA-NY. You preserve most of the default values, and enable the sequence number validation feature.

Configuration

Step-by-Step Procedure

To configure a GTP inspection object:

1. Create a GTP inspection object.

   ```
   [edit]
   user@host# set security gprs gtp profile la-ny
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235

Understanding GTP-U Inspection

The GPRS tunneling protocol user plane (GTP-U) inspection performs security checks on GTP-U packets. When GTP-U inspection is enabled, the invalid GTP-U packets are blocked and the GPRS support node (GSN) is protected from a GTP-U attack.

Once GTP-U inspection is enabled and depending on the device configuration, GTP-U inspection might include checks on GTP-in-GTP packets, end-user authorization, packet sequence validity, and tunnel validity. If any configured check fails, the GTP-U packet is dropped.

The following list describes the various types of GTP-U inspections that are performed on the traffic:

Copyright © 2016, Juniper Networks, Inc.

- **GTP-U tunnel check**—The GTP-U module checks that the GTP-U packet matches a GTP tunnel. If no tunnel matches the GTP-U packet, then the GTP-U packet is dropped.
- **GTP-in-GTP check**—In the SPU, the GTP module checks to ensure that the GTP-U payload is not a GTP packet. If the payload is a GTP packet, then the GTP packet is dropped.
- **End-user address check**—If the user tunnel is found for the GTP-U packet, then the GTP-U module checks for the end-user address. If the GTP-U payload address does not match the end-user address, then the GTP-U packet is dropped.

NOTE: In certain scenarios, the end-user address is not carried in GTP create messages. For example, if DHCPv4 is used for IPv4 address allocation, the IPv4 address field in the GTP create message will be set to 0.0.0.0. The user equipment and GGSN/PGW get the address from the DHCP server. In this scenario, the GTP module cannot get the address for the end-user address check. Subsequently, if this configuration is enabled, the GTP create message will be dropped.

- **Sequence number check**—The GTP-U module compares the GTP-U packet sequence number with the sequence number stored in the GTP-U tunnel. If it is not in the specified range, then the GTP-U packet is dropped. If it is in the range, then the GTP-U tunnel refreshes the sequence number and allows the GTP-U packet to pass.

NOTE: At the end of the GTP-U inspection, the GTP-U tunnel refreshes the timers and counters.

Related Documentation

- GPRS Overview on page 2215

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 98

Configuring GTP Message Filtering

- Understanding GTP Message Filtering on page 2233
- Understanding GTP Message-Length Filtering on page 2233
- Example: Setting the GTP Message Lengths on page 2234
- Understanding GTP Message-Type Filtering on page 2235
- Supported GTP Message Types on page 2235
- Example: Permitting and Denying GTP Message Types on page 2238
- Understanding GTP Message-Rate Limiting on page 2239
- Understanding GTP Control Message Path Rate Limiting on page 2239
- Example: Limiting the Message Rate and Path Rate for GTP Control Messages on page 2240
- Example: Enabling GTP Sequence Number Validation on page 2244
- Understanding GTP IP Fragmentation on page 2245

Understanding GTP Message Filtering

When the device receives a GPRS tunneling protocol (GTP) packet, it checks the packet against policies configured on the device. If the packet matches a policy, the device inspects the packet according to the GTP configuration applied to the policy. If the packet fails to meet any of the GTP configuration parameters, the device will pass or drop the packets based on the configuration of the GTP inspection object.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message-Length Filtering on page 2233
- Supported GTP Message Types on page 2235

Understanding GTP Message-Length Filtering

You can configure the device to drop packets that do not meet your specified minimum or maximum message lengths. In the GPRS tunneling protocol (GTP) header, the message

Copyright © 2016, Juniper Networks, Inc.

length field indicates the length, in octets, of the GTP payload. It does not include the length of the GTP header itself, the UDP header, or the IP header. The default minimum and maximum GTP message lengths are 0 and 65,535 bytes, respectively.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Example: Setting the GTP Message Lengths on page 2234
- Supported GTP Message Types on page 2235

Example: Setting the GTP Message Lengths

This example shows how to set the GTP message lengths.

- Requirements on page 2234
- Overview on page 2234
- Configuration on page 2234
- Verification on page 2235

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the minimum GTP message length to 8 octets and the maximum GTP message length to 1200 octets for the GTP inspection object.

Configuration

Step-by-Step Procedure

To configure the GTP message lengths:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Specify the minimum message length.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 min-message-length 8
   ```

3. Specify the maximum message length.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 max-message-length 1200
   ```

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

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


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message-Length Filtering on page 2233
- Supported GTP Message Types on page 2235

Understanding GTP Message-Type Filtering

You can configure the device to filter GPRS tunneling protocol (GTP) packets and permit or deny them based on their message type. By default, the device permits all GTP message types.

A GTP message type includes one or many messages. When you permit or deny a message type, you automatically permit or deny all messages of the specified type. For example, if you select to drop the sgsn-context message type, you thereby drop sgsn-context-request, sgsn-context-response, and sgsn-context-acknowledge messages.

You permit and deny message types based on the GTP version number. For example, you can deny message types for one version while you permit them for the other version.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Example: Permitting and Denying GTP Message Types on page 2238
- Supported GTP Message Types on page 2235

Supported GTP Message Types

Table 246 lists the GTP messages supported in GTP Releases 1997 and 1999 (including charging messages for GTP) and the message types that you can use to configure GTP message-type filtering.

Table 246: GTP Messages

Message	Message Type	Version 0	Version 1
create AA pdp context request	create-aa-pdp	b	
create AA pdp context response	create-aa-pdp	b	
create pdp context request	create-pdp	b	b

Copyright © 2016, Juniper Networks, Inc.

Table 246: GTP Messages *(continued)*

Message	Message Type	Version 0	Version 1
create pdp context response	create-pdp	b	b
data record request	data-record	b	b
data record response	data-record	b	b
delete AA pdp context request	delete-aa-pdp	b	
delete AA pdp context response	delete-aa-pdp	b	
delete pdp context request	delete-pdp	b	b
delete pdp context response	delete-pdp	b	b
echo request	echo	b	b
echo response	echo	b	b
error indication	error-indication	b	b
failure report request	failure-report	b	b
failure report response	failure-report	b	b
forward relocation request	fwd-relocation	b	b
forward relocation response	fwd-relocation	b	b
forward relocation complete	fwd-relocation	b	b
forward relocation complete acknowledge	fwd-relocation	b	b
forward SRNS context	fwd-srns-context	b	b
forward SRNS context acknowledge	fwd-srns-context	b	b
identification request	identification	b	b
identification response	identification	b	b
node alive request	node-alive	b	b
node alive response	node-alive	b	b
note MS GPRS present request	note-ms-present	b	b

Copyright © 2016, Juniper Networks, Inc.

Table 246: GTP Messages *(continued)*

Message	Message Type	Version 0	Version 1
note MS GPRS present response	note-ms-present	b	b
pdu notification request	pdu-notification	b	b
pdu notification response	pdu-notification	b	b
pdu notification reject request	pdu-notification	b	b
pdu notification reject response	pdu-notification	b	b
RAN info relay	ran-info	b	b
redirection request	redirection	b	b
redirection response	redirection	b	b
relocation cancel request	relocation-cancel	b	b
relocation cancel response	relocation-cancel	b	b
send route info request	send-route	b	b
send route info response	send-route	b	b
sgsn context request	sgsn-context	b	b
sgsn context response	sgsn-context	b	b
sgsn context acknowledge	sgsn-context	b	b
supported extension headers notification	supported-extension	b	b
g-pdu	gtp-pdu	b	b
update pdp context request	update-pdp	b	b
updated pdp context response	update-pdp	b	b
version not supported	version-not-supported	b	b

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message-Type Filtering on page 2235

Copyright © 2016, Juniper Networks, Inc.

- Example: Setting the GTP Message Lengths on page 2234

Example: Permitting and Denying GTP Message Types

This example shows how to permit and deny GTP message types.

- Requirements on page 2238
- Overview on page 2238
- Configuration on page 2238
- Verification on page 2238

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, for the gtp1 profile, you configure the device to drop the error-indication and failure-report message types for version 1.

Configuration

Step-by-Step Procedure

To permit and deny GTP message types:

1. Configure the device.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Drop the error indication.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 drop error-indication 1
   ```

3. Drop the failure report messages.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 drop failure-report 1
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message-Type Filtering on page 2235

Copyright © 2016, Juniper Networks, Inc.

- Supported GTP Message Types on page 2235

Understanding GTP Message-Rate Limiting

You can configure the device to limit the rate of network traffic going to a GPRS support node (GSN). You can set separate thresholds, in packets per second, for GGSN tunneling protocol, control (GTP-C) messages. Because GTP-C messages require processing and replies, they can potentially overwhelm a GSN. By setting a rate limit on GTP-C messages, you can protect your GSNs from possible denial-of-service (DoS) attacks such as the following:

- Border gateway bandwidth saturation—A malicious operator connected to the same GPRS Roaming Exchange (GRX) as your public land mobile network (PLMN) can direct so much network traffic at your Border Gateway that legitimate traffic is starved for bandwidth in or out of your PLMN, thus denying roaming access to or from your network.
- GTP flood—GPRS tunneling protocol (GTP) traffic can flood a GSN, forcing it to spend its CPU cycles processing illegitimate data. This can prevent subscribers from roaming and forwarding data to external networks, and it can prevent a General Packet Radio Service (GPRS) from attaching to the network.

This feature limits the rate of traffic sent to each GSN from the Juniper Networks device. The default rate is unlimited.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Example: Limiting the Message Rate and Path Rate for GTP Control Messages on page 2240
- Supported GTP Message Types on page 2235

Understanding GTP Control Message Path Rate Limiting

You can restrict the maximum packets per second for specific control messages on a path on all high-end SRX Series devices. These GPRS tunneling protocol (GTP) messages include **create-req**, **delete-req**, and other GTP messages. However, you can restrict the maximum packets per minute for an **echo-req** GTP message.

The **path-rate-limit** function controls specific GTP messages in both the forward and reverse directions. A drop threshold and an alarm threshold can be configured for each control message in the forward and reverse direction for one path. If the control messages on one path reach the alarm threshold, an alarm log is generated. If the number of control messages received reaches the drop threshold, a packet drop log is generated and all other control messages of this type received later are dropped.

To control message traffic in the forward and reverse directions, configure a policy on the device such that the direction that is consistent with the configured policy is defined

Copyright © 2016, Juniper Networks, Inc.

as forward, and the opposite direction is defined as reverse. Use the **set security gprs gtp profile <profile-name> path-rate-limit** statement to restrict the maximum packets per second for specific control messages on a path.

NOTE: You can configure both the **rate-limit** and the **path-rate-limit** options at the same time.

Related Documentation

- Example: Limiting the Message Rate and Path Rate for GTP Control Messages on page 2240

Example: Limiting the Message Rate and Path Rate for GTP Control Messages

This example shows how to limit the message rate and the path rate for GTP control messages. The **rate-limit** option limits the GTP messages per second and the **path-rate-limit** option controls specific GTP messages in both the forward and reverse directions.

- Requirements on page 2240
- Overview on page 2240
- Configuration on page 2241
- Verification on page 2243

Requirements

This example uses the following hardware and software components:

- A high-end SRX Series device
- Junos OS Release 12.1X45-D10

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you limit the rate of incoming GTP messages to 300 packets per second and you limit the path rate for GTP control messages in both the forward and reverse directions. You configure the device to limit the rate of network traffic going to a GPRS support node (GSN), and you restrict the maximum packets per second or per minute for specific control messages on a path. For **create-req**, **delete-req**, and **other** GTP messages you restrict the maximum packets per second. However, for an **echo-req** GTP message, you restrict the maximum packets per minute.

The **path-rate-limit** function controls specific GTP messages in both the forward and reverse directions. Configure the **alarm-threshold** parameter to configure the device to raise an alarm when the GTP control messages on a path have reached the configured limit. Configure the **drop-threshold** to drop traffic when the number of packets per second or per minute exceeds the configured limit.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

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

```
set security gprs gtp profile gtp1 rate-limit 300
set security gprs gtp profile gtp1 path-rate-limit message-type create-req alarm-threshold
    forward 50 reverse 50
set security gprs gtp profile gtp1 path-rate-limit message-type delete-req alarm-threshold
    forward 50 reverse 50
set security gprs gtp profile gtp1 path-rate-limit message-type echo-req alarm-threshold
    forward 50 reverse 50
set security gprs gtp profile gtp1 path-rate-limit message-type other alarm-threshold
    forward 50 reverse 50
set security gprs gtp profile gtp1 path-rate-limit message-type create-req drop-threshold
    forward 80 reverse 80
set security gprs gtp profile gtp1 path-rate-limit message-type delete-req drop-threshold
    forward 80 reverse 80
set security gprs gtp profile gtp1 path-rate-limit message-type echo-req drop-threshold
    forward 80 reverse 80
set security gprs gtp profile gtp1 path-rate-limit message-type other drop-threshold
    forward 80 reverse 80
```

Step-by-Step Procedure

To configure the GTP message rate and path rate limit:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Set the GTP message rate limit.

   ```
   [edit security gprs gtp profile gtp1]
   user@host# set rate-limit 300
   ```

3. Specify the message type to set the path rate limit for GTP control messages.

   ```
   [edit security gprs gtp profile gtp1]
   user@host# set path-rate-limit message-type
   ```

4. Select GTP control message types.

   ```
   [edit security gprs gtp profile gtp1]
   user@host# set path-rate-limit message-type create-req
   user@host# set path-rate-limit message-type delete-req
   user@host# set path-rate-limit message-type echo-req
   user@host# set path-rate-limit message-type other
   ```

5. Set the alarm threshold for the GTP control message types.

   ```
   [edit security gprs gtp profile gtp1 path-rate-limit]
   user@host# set message-type create-req alarm threshold
   user@host# set message-type delete-req alarm threshold
   user@host# set message-type echo-req alarm threshold
   user@host# set message-type other alarm threshold
   ```

Copyright © 2016, Juniper Networks, Inc.

6. Limit the control messages in the forward direction.

```
[edit security gprs gtp profile gtp1 path-rate-limit message-type]
user@host# set create-req alarm threshold forward 50
user@host# set delete-req alarm threshold forward 50
user@host# set echo-req alarm threshold forward 50
user@host# set other alarm threshold forward 50
```

7. Limit the control messages in the reverse direction.

```
[edit security gprs gtp profile gtp1 path-rate-limit message-type]
user@host# set create-req alarm threshold reverse 50
user@host# set delete-req alarm threshold reverse 50
user@host# set echo-req alarm threshold reverse 50
user@host# set other alarm threshold reverse 50
```

8. Set the drop threshold for the GTP control message types.

```
[edit security gprs gtp profile gtp1 path-rate-limit]
user@host# set message-type create-req drop threshold
user@host# set message-type delete-req drop threshold
user@host# set message-type echo-req drop threshold
user@host# set message-type other drop threshold
```

9. Limit the control messages in the forward direction.

```
[edit security gprs gtp profile gtp1 path-rate-limit message-type]
user@host# set create-req drop threshold forward 80
user@host# set delete-req drop threshold forward 80
user@host# set echo-req drop threshold forward 80
user@host# set other drop threshold forward 80
```

10. Limit the control messages in the reverse direction.

```
[edit security gprs gtp profile gtp1 path-rate-limit message-type]
user@host# set create-req drop threshold reverse 80
user@host# set delete-req drop threshold reverse 80
user@host# set echo-req drop threshold reverse 80
user@host# set other drop threshold reverse 80
```

Results

From configuration mode, confirm your configuration by entering the **show security gprs gtp profile *profile-name*** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security gprs gtp profile p1
    rate-limit 300;
    path-rate-limit {
      message-type create-req {
        drop-threshold {
          forward 80;
          reverse 80;
        }
        alarm-threshold {
          forward 50;
          reverse 50;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      message-type delete-req {
        drop-threshold {
          forward 80;
          reverse 80;
        }
        alarm-threshold {
          forward 50;
          reverse 50;
        }
      }
      message-type echo-req {
        drop-threshold {
          forward 80;
          reverse 80;
        }
        alarm-threshold {
          forward 50;
          reverse 50;
        }
      }
      message-type other {
        drop-threshold {
          forward 80;
          reverse 80;
        }
        alarm-threshold {
          forward 50;
          reverse 50;
        }
      }
    }
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose Verify that the GTP message rate and path rate limit configuration is correct.

Action From operational mode, enter the **show security gprs gtp counters path-rate-limit** command.

```
Path-rate-limit counters:
                                    Drop                    Alarm
Create Request                       20                      50
Delete Request                       20                      50
Echo Request                         20                      50
Others                               20                      50
```

Meaning The **show security gprs gtp counters path-rate-limit** command displays the number of packets received since the alarm threshold or the drop threshold value was reached. If you configure the **alarm-threshold** value as 50 and the **drop-threshold** value as 80 for the Create Request message, and if the device receives 100 packets in a second or minute, then the Drop number will be 20 and the Alarm number will be 50.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message-Rate Limiting on page 2239
- Supported GTP Message Types on page 2235

Example: Enabling GTP Sequence Number Validation

This example shows how to enable GTP sequence number validation feature.

- Requirements on page 2244
- Overview on page 2244
- Configuration on page 2244
- Verification on page 2245

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set the gtp profile as gtp1 and you also enable the sequence number validation feature.

Configuration

Step-by-Step Procedure To enable GTP sequence number validation feature:

1. Set the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Enable the sequence number validation.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 seq-number-validated
   ```

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

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


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Understanding GTP IP Fragmentation

A GPRS tunneling protocol (GTP) packet consists of the message body and three headers: GTP, UDP, and IP. If the resulting IP packet is larger than the maximum transmission unit (MTU) on the transferring link, the sending Serving GPRS Support Node (SGSN) or gateway GPRS support node (GGSN) performs an IP fragmentation.

By default, the device buffers IP fragments until it receives a complete GTP message, and then inspects the GTP message.

Related Documentation

- Understanding GTP Information Elements on page 2247
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 99

Configuring GTP Information Elements

- Understanding GTP Information Elements on page 2247
- Understanding GTP APN Filtering on page 2248
- Example: Setting a GTP APN and a Selection Mode on page 2249
- Understanding IMSI Prefix Filtering of GTP Packets on page 2250
- Example: Setting a Combined IMSI Prefix and APN Filter on page 2251
- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Supported R6 Information Elements on page 2252
- Example: Removing R6 Information Elements from GTP Messages on page 2256
- Supported R7 Information Elements on page 2257
- Example: Removing R7 Information Elements from GTP Messages on page 2258
- Supported R8 Information Elements on page 2259
- Example: Removing R8 Information Elements from GTP Messages on page 2259
- Supported R9 Information Elements on page 2260
- Example: Removing R9 Information Elements from GTP Messages on page 2261
- Understanding GTPv1 Information Element Removal on page 2262
- Example: Removing GTPv1 Information Elements Using IE Number on page 2262

Understanding GTP Information Elements

Information elements (IEs) are included in all GPRS tunneling protocol (GTP) control message packets. IEs provide information about GTP tunnels, such as creation, modification, deletion, and status. Junos OS supports IEs consistent with Third–Generation Partnership Project (3GPP) Release 6, Release 7, Release 8, and Release 9. If you have contractual agreements with operators running earlier releases of 3GPP, you can reduce network overhead by restricting control messages containing unsupported IEs.

NOTE: If a new information element (IE) is introduced, there will be no drop in GTP messages because GTP passes the messages even if it encounters unknown new IEs.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding GTP IP Fragmentation on page 2245
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Understanding GTP APN Filtering

An access point name (APN) is an information element (IE) included in the header of a GPRS tunneling protocol (GTP) packet that provides information about how to reach a network. An APN comprises two elements:

- Network ID—Identifies the name of an external network such as example.com.
- Operator ID—Uniquely identifies the operators' public land mobile network (PLMN) such as mnc123.mcc456.

By default, the device permits all APNs. However, you can configure the device to perform APN filtering to restrict access to roaming subscribers to external networks.

To enable APN filtering, you must specify one or more APNs. To specify an APN, you need to know the domain name of the network (for example, example.com) and, optionally, the operator ID. Because the domain name (network ID) portion of an APN can potentially be very long and contain many characters, you can use the wildcard (*) as the first character of the APN. The wildcard indicates that the APN is not limited only to example.com but also includes all the characters that might precede it.

You may also set a *selection mode* for the APN. The selection mode indicates the origin of the APN and whether or not the Home Location Register (HLR) has verified the user subscription. You set the selection mode according to the security needs of your network. Possible selection modes include the following:

- Mobile Station—Mobile station-provided APN, subscription not verified.

  This selection mode indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.

- Network—Network-provided APN, subscription not verified.

  This selection mode indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.

- Verified—MS or network-provided APN, subscription verified.

  This selection mode indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.

APN filtering applies only to create-pdp-request messages. When performing APN filtering, the device inspects GTP packets to—look for APNs that match APNs that you set. If the APN of a GTP packet matches an APN that you specified, the device then

Copyright © 2016, Juniper Networks, Inc.

verifies the selection mode and only forwards the GTP packet if both the APN and the selection mode match the APN and the selection mode that you specified. Because APN filtering is based on perfect matches, using the wildcard (*) when setting an APN suffix can prevent the inadvertent exclusion of APNs that you would otherwise authorize.

Additionally, the device can filter GTP packets based on the combination of an International Mobile Subscriber Identity (IMSI) prefix and an APN. When you filter GTP packets based on an IMSI prefix, you must also specify an APN.

An APN string is case-insensitive. For instance, in the following example you set two APN strings, WWW.EXAMPLE.COM.CN and www.example.com.cn, with the same IMSI prefix value. In this configuration, the lowercase string will display after the uppercase string, and the packet will be dropped.

user@host# show configuration security gprs gtp | display set

set security gprs gtp profile test apn WWW.EXAMPLE.COM.CN imsi-prefix * action pass

set security gprs gtp profile test apn www.example.com.cn imsi-prefix * action drop

If an APN is configured with two IMSI prefix entries, then the IMSI prefix with the longest match takes priority. For example, see the following configuration:

user@host# show configuration security gprs gtp | display set

set security gprs gtp profile test apn WWW.EXAMPLE.COM.CN imsi-prefix 12345678 action pass

set security gprs gtp profile test apn www.example.com.cn imsi-prefix 12345 action drop

If an incoming packet value matches the IMSI prefix value 12345678, then the packet will pass. The IMSI prefix value 12345678 takes precedence over the IMSI prefix value 12345, as the longest matched IMSI prefix takes priority.

Related Documentation

- Example: Setting a GTP APN and a Selection Mode on page 2249
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Example: Setting a GTP APN and a Selection Mode

This example shows how to set a GTP APN and a selection mode.

- Requirements on page 2250
- Overview on page 2250

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 2250
- Verification on page 2250

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set a GTP APN as example.com.mnc123.mcc456.gprs and use the wildcard (*) character. You also set the IMSI prefix and set the selection mode as network.

Configuration

Step-by-Step Procedure

To configure a GTP APN and a selection mode:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Set a selection mode for the APN.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 apn
     *example.com.mnc123.mcc456.gprs imsi-prefix * action selection net
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP APN Filtering on page 2248
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Understanding IMSI Prefix Filtering of GTP Packets

A GPRS support node (GSN) identifies a mobile station (MS) by its International Mobile Station Identity (IMSI). An IMSI consists of three elements: the mobile country code (MCC), the mobile network code (MNC), and the Mobile Subscriber Identification Number (MSIN). The MCC and MNC combined constitute the IMSI prefix and identify the mobile subscriber's home network, or public land mobile network (PLMN).

Copyright © 2016, Juniper Networks, Inc.

By setting IMSI prefixes, you can configure the device to deny GPRS tunneling protocol (GTP) traffic coming from nonroaming partners. By default, a device does not perform IMSI prefix filtering on GTP packets. By setting IMSI prefixes, you configure the device to filter create-pdp-request messages and permit only GTP packets with IMSI prefixes that match the ones you set. The device allows GTP packets with IMSI prefixes that do not match any of the IMSI prefixes that you set. To block GTP packets with IMSI prefixes that do not match any of the IMSI prefixes set, use an explicit wildcard for the IMSI filter, and the drop action should be the last IMSI prefix filtering policy.

When you filter GTP packets based on an IMSI prefix, you must also specify an APN.

Related Documentation

- Example: Setting a Combined IMSI Prefix and APN Filter on page 2251
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Example: Setting a Combined IMSI Prefix and APN Filter

This example shows how to set and combine IMSI prefix and APN filter.

- Requirements on page 2251
- Overview on page 2251
- Configuration on page 2251
- Verification on page 2252

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set example.com.mnc123.mcc456.gprs as an APN and use the wildcard(*). You permit all selection modes for this APN. You also set the IMSI prefix for a known PLMN, which is 246565. The MCC-MNC pair can be five or six digits.

Configuration

Step-by-Step Procedure

To set and combine IMSI prefix and APN filter:

1. Set the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Set the selection mode for APN.

   ```
   [edit]
   ```

```
user@host# set security gprs gtp profile gtp1 apn
  *example.com.mnc123.mcc456.gprs imsi-prefix 246565* action pass
```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding IMSI Prefix Filtering of GTP Packets on page 2250
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Understanding R6, R7, R8, and R9 Information Elements Removal

The Third-Generation Partnership Project (3GPP) R6, R7, R8, and R9 information elements (IEs) removal feature allows you to retain interoperability in roaming between Second-Generation Partnership Project (2GPP) and 3GPP networks. You can configure the GPRS tunneling protocol (GTP)-aware Juniper Networks device, residing on the border of a public land mobile network (PLMN) and a GPRS Roaming Exchange (GRX) and acting as a Gp firewall, to remove 3GPP-specific attributes from the GTP packet header when the packet passes into a 2GPP network. You can configure the device to remove the RAT, RAI, Common Flags, ULI, MS Time Zone, IMEI-SV, and access point name (APN) restriction IEs from GTP messages prior to forwarding these messages to the gateway GPRS support node (GGSN).

Related Documentation

- Example: Removing R6 Information Elements from GTP Messages on page 2256
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Supported R6 Information Elements

Junos OS supports all 3GPP R6 IEs for GTP), as listed in Table 247.

Table 247: Supported Information Elements

IE Type Value	Information Element
1	Cause

Copyright © 2016, Juniper Networks, Inc.

Table 247: Supported Information Elements *(continued)*

IE Type Value	Information Element
2	International Mobile Subscriber Identity (IMSI)
3	Routing Area Identity (REI)
4	Temporary Logical Link Identity (TLLI)
5	Packet TMSI (P-TMSI)
8	Reordering Required
9	Authentication Triplet
11	MAP Cause
12	P-TMSI Signature
13	MS Validated
14	Recovery
15	Selection Mode
16	Tunnel Endpoint Identifier Data I
17	Tunnel Endpoint Identifier Control Plane
18	Tunnel Endpoint Identifier Data II
19	Teardown ID
20	NSAPI
21	RANAP Cause
22	RAB Context
23	Radio Priority SMS
24	Radio Priority
25	Packet Flow ID
26	Charging Characteristics
27	Trace Reference
28	Trace Type

Copyright © 2016, Juniper Networks, Inc.

Table 247: Supported Information Elements *(continued)*

IE Type Value	Information Element
29	MS Not Reachable Reason
127	Charging ID
128	End User Address
129	MM Context
130	PDP Context
131	Access Point Name
132	Protocol Configuration Options
133	GSN Address
134	MS International PSTN/ISDN Number (MSISDN)
135	Quality of Service Profile
136	Authentication Quintuplet
137	Traffic Flow Template
138	Target Identification
139	UTRAN Transparent Container
140	RAB Setup Information
141	Extension Header Type List
142	Trigger Id
143	OMC Identity
144	RAN Transparent Container
145	PDP Context Prioritization
146	Additional RAB Setup Information
147	SGSN Number
148	Common Flags
149	APN Restriction

Copyright © 2016, Juniper Networks, Inc.

Table 247: Supported Information Elements *(continued)*

IE Type Value	Information Element
150	Radio Priority LCS
151	RAT Type
152	User Location Information
153	MS Time Zone
154	IMEI-SV
155	CAMEL Charging Information Container
156	MBMS UE Context
157	Temporary Mobile Group Identity (TMGI)
158	RIM Routing Address
159	MBMS Protocol Configuration Options
160	MBMS Service Area
161	Source TNC PDCP context Information
162	Additional Trace Information
163	Hop Counter
164	Selected PLMN ID
165	MBMS Session Identifier
166	MBMS2G/3G Indicator
167	Enhanced NSAPI
168	MBMS Session Duration
169	Additional MBMS Trace Information
173	BSS Container
174	Cell Identification
175	PDU Numbers
176	BSSGP Cause

Copyright © 2016, Juniper Networks, Inc.

Table 247: Supported Information Elements *(continued)*

IE Type Value	Information Element
178	RIM Routing Address Discriminator
179	List of setup PFCS
180	PS Hand-over XID Parameters
188	Reliable INTER RAT HANDOVER INFO
251	Charging Gateway Address
255	Private Extension

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Example: Removing R6 Information Elements from GTP Messages on page 2256
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Example: Removing R6 Information Elements from GTP Messages

This example shows how to remove R6 information elements from GTP messages.

- Requirements on page 2256
- Overview on page 2256
- Configuration on page 2256
- Verification on page 2257

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the Gp interface of the security device to remove newly added R6 IEs (RAT, Common Flags, ULI, IMEI-SV, MS Time Zone, and APN restrictions) from the GTP message.

Configuration

Step-by-Step Procedure

To remove R6 information elements from GTP messages:

1. Specify the GTP profile.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security gprs gtp profile gtp1
```

2. Specify the information element.

```
[edit]
user@host# set security gprs gtp profile gtp1 remove-ie version v1 release R6
```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Supported R7 Information Elements

Junos OS supports all 3GPP R7 IEs for GTP, as listed in Table 248.

Table 248: Supported Information Elements

IE Type Value	Information Element
172	PS Handover Request Context
181	MS Info Change Reporting Action
182	Direct Tunnel Flags
183	Correlation-ID
184	Bearer Control Mode

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Example: Removing R7 Information Elements from GTP Messages on page 2258
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Example: Removing R7 Information Elements from GTP Messages

This example shows how to remove R7 information elements from GTP messages.

- Requirements on page 2258
- Overview on page 2258
- Configuration on page 2258
- Verification on page 2258

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the Gp interface of the security device to remove newly added R7 IEs from the GTP message.

Configuration

Step-by-Step Procedure

To remove R7 information elements from GTP messages:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Specify the information element.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 remove-ie version v1 release R7
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Supported R8 Information Elements

Junos OS supports all 3GPP R8 IEs for GTP, as listed in Table 249.

Table 249: Supported Information Elements

IE Type Value	Information Element
189	RFSP Index

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Example: Removing R8 Information Elements from GTP Messages on page 2259
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Example: Removing R8 Information Elements from GTP Messages

This example shows how to remove R8 information elements from GTP messages.

- Requirements on page 2259
- Overview on page 2259
- Configuration on page 2259
- Verification on page 2260

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the Gp interface of the security device to remove newly added R8 IEs from the GTP message.

Configuration

Step-by-Step Procedure

To remove R8 information elements from GTP messages:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Specify the information element.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 remove-ie version v1 release R8
   ```

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

Copyright © 2016, Juniper Networks, Inc.

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Supported R9 Information Elements

Junos OS supports all 3GPP R9 IEs for GTP, as listed in Table 250.

Table 250: Supported Information Elements

IE Type Value	Information Element
190	Fully Qualified Domain Name (FQDN)
191	Evolved Allocation/Retention Priority 1
192	Evolved Allocation/Retention Priority 2
193	Extended Common Flags
194	User CSG Information (UCI)
195	CSG Information Reporting Action
196	CSG ID
197	CSG Membership Indication (CMI)
198	Aggregate Maximum Bit Rate (AMBR)

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Example: Removing R9 Information Elements from GTP Messages on page 2261
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Example: Removing R9 Information Elements from GTP Messages

This example shows how to remove R9 information elements from GTP messages.

- Requirements on page 2261
- Overview on page 2261
- Configuration on page 2261
- Verification on page 2261

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure the Gp interface of the security device to remove newly added R9 IEs from the GTP message.

Configuration

Step-by-Step Procedure

To remove R9 information elements from GTP messages:

1. Specify the GTP profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1
   ```

2. Specify the information element.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp1 remove-ie version v1 release R9
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding R6, R7, R8, and R9 Information Elements Removal on page 2252
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Understanding GTPv1 Information Element Removal

The number of network elements in a mobile network is expanding with the introduction of multiple releases of 3GPP specifications. Every release introduces newer information elements (IEs) that are not defined in the prior releases. Therefore mobile networks have diverse set of network elements creating inter operability problems between different releases of the devices. You can configure the GPRS tunneling protocol (GTP) firewall to remove information elements (IE) by release with the following command.

set security gprs gtp profile *gtp1* remove-ie.

However newer IEs that will be introduced in the future releases might also cause inter operability problems. Each information element has a unique ID, the IE number. IE numbers range from 1 to 255. You can configure the GTP firewall to remove specific IEs using the user-configured IE number.

When you configure the IE removal, the GTP firewall deletes the corresponding IEs of the GTPv1 messages; updates the length of the GTP, the UDP, and the IP; and then passes the GTPv1 message. The GTP firewall also updates the cyclic redundancy check (CRC) code. IE removal by IE number supports all IEs, ranging from 1 to 255.

You can remove the IE removal configuration with the following commands:

delete security gprs gtp profile *gtp1* remove-ie—Deletes the IE removal configuration for the GTP profile GTP1.

delete security gprs gtp profile *gtp1* remove-ie version *v1* number *4*—Deletes the IE removal configuration for GTP profile with version v1 and IE number 4.

NOTE: **The IE removal feature supports GTPv1 only.**

Related Documentation

- Understanding GTP Information Elements on page 2247
- Example: Removing GTPv1 Information Elements Using IE Number on page 2262

Example: Removing GTPv1 Information Elements Using IE Number

This example shows how to configure the GPRS tunelling protocol (GTP) interface of the security device to remove user-configured IEs from GTP messages.

- Requirements on page 2262
- Overview on page 2263
- Configuration on page 2263

Requirements

Enable GTP on the device.

Copyright © 2016, Juniper Networks, Inc.

set security gprs gtp enable

Overview

In this example, you configure IE removal for the GTP profile called gtp1. The IEs are removed using the user-configured IE number 4.

Configuration

CLI Quick Configuration

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

```
set security gprs gtp profile gtp1
set security gprs gtp profile gtp1 remove-ie version v1 number 4
```

Step-by-Step Procedure

To configure the GTP interface of the security device to remove user-configured IEs from the GTP message:

1. Specify the GTP profile.

   [edit]

   user@host# **set security gprs gtp profile gtp1**

2. Specify the IE number.

   [edit security gprs gtp profile gtp1]

   user@host# **set remove-ie version v1 number 4**

Results

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

```
gtp {
    profile gtp1 {
        remove-ie {
            version v1 {
                number 4;
            }
        }
    }
}
```

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

Related Documentation

- Understanding GTP Information Elements on page 2247
- Understanding GTPv1 Information Element Removal on page 2262

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 100

Configuring NAT for GTP

- Understanding NAT for GTP on page 2265
- Example: Configuring GTP Inspection in NAT on page 2266
- Understanding Network Address Translation-Protocol Translation on page 2270
- Example: Enhancing Traffic Engineering by Configuring NAT-PT Between an IPv4 and an IPv6 Endpoint with SCTP Multihoming on page 2271

Understanding NAT for GTP

A General Packet Radio Service (GPRS) interface supports both GPRS tunneling protocol (GTP) inspection and Network Address Translation (NAT) simultaneously in the same routing instance. When GTP packets configured with static NAT are inspected in a network, only addresses within IP headers are translated. The addresses within their payloads are not translated. For each endpoint, the related GTP session must belong to the same zone and virtual router. This means the header source IP, C-tunnel IP, and U-tunnel IP in the payload are defined in the same scope for a packet.

NOTE: **When you enable NAT, only the outer IP packet has to be translated. The embedded IP addresses are not translated.**

During a GTP packet flow, the source IP address and destination IP address cannot be translated to NAT simultaneously. When you delete or deactivate NAT rule configuration on a device, the NAT rule related GSN and GTP tunnels are deleted. If the NAT rule related GSN number and tunnel number are huge, this deleting process will take several minutes.

Related Documentation

- Understanding GTP IP Fragmentation on page 2245
- Understanding GTP Inspection Objects on page 2229
- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring GTP Inspection in NAT

This example shows how to configure a NAT rule to map a private IP (one that is inside the network and not routable) to a public IP (one that is outside of the network and is routable). It also shows how to inspect GTP traffic between an internal and external network.

- Requirements on page 2266
- Overview on page 2266
- Configuration on page 2266
- Verification on page 2270

Requirements

Before you begin, the device must be restarted after GTP is enabled. By default, GTP is disabled on the device.

Overview

In this example, you configure interfaces as ge-0/0/0 and ge-0/0/1, with addresses 10.0.0.254/8 and 123.0.0.254/8. You then configure the security zone and static NAT. You enable the GTP service in the security policies to allow bidirectional traffic between two networks, and you check the traffic between the internal and external network.

Configuration

CLI Quick Configuration

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

```
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.254/8
set interfaces ge-0/0/1 unit 0 family inet address 123.0.0.254/8
set security zones security-zone zone1 interfaces ge-0/0/0.0 host-inbound-traffic
  system-services all
set security zones security-zone zone1 host-inbound-traffic protocols all
set security zones security-zone other-zone interfaces ge-0/0/1.0 host-inbound-traffic
  system-services all
set security zones security-zone other-zone host-inbound-traffic protocols all
set security address-book global address gsn1 10.0.0.1/8
set security address-book global address other-gsn 20.0.0.1/8
set security nat static rule-set rs1 from zone other-zone
set security nat static rule-set rs1 rule r1 match destination-address 123.0.0.1/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 10.0.0.1/32
set security nat proxy-arp interface ge-0/0/0.0 address 123.0.0.1/32
set security gprs gtp enable
set security gprs gtp profile gtp1
set security gprs gtp profile gtp1 timeout 1
set security gprs gtp profile gtp1 seq-number-validated
set security policies from-zone zone1 to-zone other-zone policy out-gtp match
  source-address gsn1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone zone1 to-zone other-zone policy out-gtp match
  destination-address other-gsn
set security policies from-zone zone1 to-zone other-zone policy out-gtp match application
  junos-gprs-gtp
set security policies from-zone zone1 to-zone other-zone policy out-gtp then permit
  application-services gprs-gtp-profile gtp1
set security policies from-zone other-zone to-zone zone1 policy in-gtp match
  source-address other-gsn
set security policies from-zone other-zone to-zone zone1 policy in-gtp match
  destination-address gsn1
set security policies from-zone other-zone to-zone zone1 policy in-gtp match application
  junos-gprs-gtp
set security policies from-zone other-zone to-zone zone1 policy in-gtp then permit
  application-services gprs-gtp-profile gtp1
```

Step-by-Step Procedure

To configure GTP inspection in NAT:

1. Configure interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.254/8
user@host# set interfaces ge-0/0/1 unit 0 family inet address 123.0.0.254/8
```

2. Configure and security zones

```
[edit security]
user@host# set zones security-zone zone1 interfaces ge-0/0/0.0
  host-inbound-traffic system-services all
user@host# set zones security-zone zone1 host-inbound-traffic protocols all
user@host# set zones security-zone other-zone interfaces ge-0/0/1.0
  host-inbound-traffic system-services all
user@host# set zones security-zone other-zone host-inbound-traffic protocols all
```

3. Define the address book.

```
[edit security]
user@host# set address-book global address gsn1 10.0.0.1/8
user@host# set address-book global address other-gsn 20.0.0.1/8
```

4. Define NAT rule.

```
[edit security nat]
user@host# set static rule-set rs1 from zone other-zone
user@host# set static rule-set rs1 rule r1 match destination-address 123.0.0.1/32
user@host# set static rule-set rs1 rule r1 then static-nat prefix 10.0.0.1/32
user@host# set proxy-arp interface ge-0/0/0.0 address 123.0.0.1/32
```

5. Enable GTP profile.

```
[edit security gprs gtp]
user@host# set enable
user@host# set profile gtp1
user@host# set profile gtp1 timeout 1
user@host# set profile gtp1 seq-number-validated
```

6. Check GTP traffic.

```
[edit security policies]
user@host# set from-zone zone1 to-zone other-zone policy out-gtp match
  source-address gsn1
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from-zone zone1 to-zone other-zone policy out-gtp match
  destination-address other-gsn
user@host# set from-zone zone1 to-zone other-zone policy out-gtp match application
  junos-gprs-gtp
user@host# set from-zone zone1 to-zone other-zone policy out-gtp then permit
  application-services gprs-gtp-profile gtp1
user@host# set from-zone other-zone to-zone zone1 policy in-gtp match
  source-address other-gsn
user@host# set from-zone other-zone to-zone zone1 policy in-gtp match
  destination-address gsn1
user@host# set from-zone other-zone to-zone zone1 policy in-gtp match application
  junos-gprs-gtp
user@host# set from-zone other-zone to-zone zone1 policy in-gtp then permit
  application-services gprs-gtp-profile gtp1
```

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

```
[edit]
user@host# show security
gprs {
    gtp {
        enable;
        profile gtp1 {
            timeout 1;
            seq-number-validated;
        }
    }
}
address-book {
    global {
        address gsn1 10.0.0.1/8;
        address other-gsn 20.0.0.1/8;
    }
}
nat {
    static {
        rule-set rs1 {
            from zone other-zone;
            rule r1 {
                match {
                    destination-address 123.0.0.1/32;
                }
                then {
                    static-nat {
                        prefix {
                            10.0.0.1/32;
                        }
                    }
                }
            }
        }
    }
    proxy-arp {
        interface ge-0/0/0.0 {
            address {
                123.0.0.1/32;
```

 Copyright © 2016, Juniper Networks, Inc.

```
            }
        }
    }
}
policies {
    from-zone zone1 to-zone other-zone {
        policy out-gtp {
            match {
                source-address gsn1;
                destination-address other-gsn;
                application junos-gprs-gtp;
            }
            then {
                permit {
                    application-services {
                        gprs-gtp-profile gtp1;
                    }
                }
            }
        }
    }
    from-zone other-zone to-zone zone1 {
        policy in-gtp {
            match {
                source-address other-gsn;
                destination-address gsn1;
                application junos-gprs-gtp;
            }
            then {
                permit {
                    application-services {
                        gprs-gtp-profile gtp1;
                    }
                }
            }
        }
    }
}
zones {
    security-zone trust {
        host-inbound-traffic {
            system-services {
                all;
            }
            protocols {
                all;
            }
        }
        interfaces {
            ge-0/0/0.0;
        }
    }
    security-zone zone1 {
        host-inbound-traffic {
            protocols {
                all;
            }
        }
        interfaces {
            ge-0/0/0.0;
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    security-zone other-zone {
        host-inbound-traffic {
            protocols {
                all;
            }
        }
        interfaces {
            ge-0/0/1.0 {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                }
            }
        }
    }
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying GTP Inspection on NAT

Purpose Verify the GTP traffic between the internal network and the external network.

Action From operational mode, enter the **show security** command.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Understanding GTP Message Filtering on page 2233
- Supported GTP Message Types on page 2235

Understanding Network Address Translation-Protocol Translation

Network Address Translation-Protocol Translation (NAT-PT) is a protocol translation mechanism that can be done in two directions, from IPv4 address format to IPv6 address format and vice versa. NAT-PT binds the addresses in the IPv6 network with addresses in the IPv4 network and vice versa to provide transparent routing for the datagrams traversing between address realms.

In each direction, the static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes a destination IP address translation in one direction and a source IP address translation in the opposite direction.

The main advantage of NAT-PT is that the end devices and networks can run either IPv4 addresses or IPv6 addresses and traffic can be started from any side.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- GPRS Overview on page 2215
- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Features Overview on page 2323
- Example: Enhancing Traffic Engineering by Configuring NAT-PT Between an IPv4 and an IPv6 Endpoint with SCTP Multihoming on page 2271

Example: Enhancing Traffic Engineering by Configuring NAT-PT Between an IPv4 and an IPv6 Endpoint with SCTP Multihoming

This example shows how to enhance traffic engineering by configuring NAT-PT between an IPv4 endpoint and an IPv6 endpoint. NAT-PT is a protocol translation mechanism that allows communication between IPv6-only and IPv4-only nodes through protocol-independent translation of IPv4 and IPv6 datagrams, requiring no state information for the session. NAT-PT binds the addresses in the IPv6 network with addresses in the IPv4 network and vice versa to provide transparent routing for the datagrams traversing between address realms. The main advantage of NAT-PT is that the end devices and networks can run either IPv4 addresses or IPv6 addresses and traffic can be started from any side.

- Requirements on page 2271
- Overview on page 2271
- Configuration on page 2272
- Verification on page 2276

Requirements

This example uses the following hardware and software components:

- A high-end SRX Series device
- Endpoint A connected to an SRX Series device using two IPv6 addresses
- Endpoint B connected to an SRX Series device using two IPv4 addresses

Overview

In this example, you configure NAT-PT between an IPv4 endpoint and an IPv6 endpoint. Endpoint A is connected to the SRX Series device using two IPv6 addresses and endpoint B is connected to the SRX Series device using two IPv4 addresses.

You can configure the SRX Series device to translate the IP header and IP address list (located in the INIT/INT-ACK message) between an IPv4 address format and an IPv6 address format. In each direction, static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the opposite direction.

Figure 101 illustrates the network topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 101: NAT-PT Between an IPv4 Endpoint and an IPv6 Endpoint

For configuring NAT-PT details between IPv4 and IPv6 endpoints, see Table 251.

Table 251: Configuring NAT-PT Details Between IPv4 and IPv6 Endpoints

Endpoints	Address One	Address Two
A (IPv6)	2a.1.1.1/96	2c.3.3.3/96
B (IPv4)	2.2.2.2/24	4.4.4.4/34

Configuration

CLI Quick Configuration

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

```
set interfaces ge-4/0/0 unit 0 family inet address 1.1.1.100/24
set interfaces ge-4/0/0 unit 0 family inet6 address 2a::1:1:100/96
set interfaces ge-4/0/1 unit 0 family inet address 2.2.2.100/24
set interfaces ge-4/0/1 unit 0 family inet6 address 2b::2:2:100/96
set interfaces ge-4/0/2 unit 0 family inet address 3.3.3.100/24
set interfaces ge-4/0/2 unit 0 family inet6 address 2c::3:3:100/96
set interfaces ge-4/0/3 unit 0 family inet address 4.4.4.100/24
set interfaces ge-4/0/3 unit 0 family inet6 address 2d::4:4:100/96
set security zones security-zone sctp_zone1 host-inbound-traffic system-services all
set security zones security-zone sctp_zone1 host-inbound-traffic protocols all
set security zones security-zone sctp_zone1 interfaces ge-4/0/0.0
set security zones security-zone sctp_zone1 interfaces ge-4/0/2.0
set security zones security-zone sctp_zone2 host-inbound-traffic system-services all
set security zones security-zone sctp_zone2 host-inbound-traffic protocols all
set security zones security-zone sctp_zone2 interfaces ge-4/0/1.0
set security zones security-zone sctp_zone2 interfaces ge-4/0/3.0
set security nat static rule-set sctp-natpt-from-zone1 from zone sctp_zone1
set security nat static rule-set sctp-natpt-from-zone1 rule r1-dst match destination-address
  2b::2:2:2/128
set security nat static rule-set sctp-natpt-from-zone1 rule r1-dst then static-nat prefix
  2.2.2.2/32
set security nat static rule-set sctp-natpt-from-zone1 rule r3-dst match
  destination-address 2d::4:4:4/128
set security nat static rule-set sctp-natpt-from-zone1 rule r3-dst then static-nat prefix
  4.4.4.4/32
set security nat static rule-set sctp-natpt-from-zone2 from zone sctp_zone2
set security nat static rule-set sctp-natpt-from-zone2 rule r2-dst match
  destination-address 1.1.1.1/32
```


Copyright © 2016, Juniper Networks, Inc.

```
set security nat static rule-set sctp-natpt-from-zone2 rule r2-dst then static-nat prefix
  2a::1:1:1/128
set security nat static rule-set sctp-natpt-from-zone2 rule r4-dst match
  destination-address 3.3.3.3/32
set security nat static rule-set sctp-natpt-from-zone2 rule r4-dst then static-nat prefix
  2c::3:3:3/128
```

Step-by-Step Procedure

To configure NAT-PT between an IPv4 endpoint and an IPv6 endpoint:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-4/0/0 unit 0 family inet address 1.1.1.100/24
user@host# set ge-4/0/0 unit 0 family inet6 address 2a::1:1:100/96
user@host# set ge-4/0/1 unit 0 family inet address 2.2.2.100/24
user@host# set ge-4/0/1 unit 0 family inet6 address 2b::2:2:100/96
user@host# set ge-4/0/2 unit 0 family inet address 3.3.3.100/24
user@host# set ge-4/0/2 unit 0 family inet6 address 2c::3:3:100/96
user@host# set ge-4/0/3 unit 0 family inet address 4.4.4.100/24
user@host# set ge-4/0/3 unit 0 family inet6 address 2d::4:4:100/96
```

2. Configure zones.

```
[edit security zones]
user@host# set security-zone sctp_zone1 host-inbound-traffic system-services all
user@host# set security-zone sctp_zone1 host-inbound-traffic protocols all
user@host# set security-zone sctp_zone1 interfaces ge-4/0/0.0
user@host# set security-zone sctp_zone1 interfaces ge-4/0/2.0
user@host# set security-zone sctp_zone2 host-inbound-traffic system-services all
user@host# set security-zone sctp_zone2 host-inbound-traffic protocols all
user@host# set security-zone sctp_zone2 interfaces ge-4/0/1.0
user@host# set security-zone sctp_zone2 interfaces ge-4/0/3.0
```

3. Configure rules for the first static NAT zone.

```
[edit security nat]
user@host# set static rule-set sctp-natpt-from-zone1 from zone sctp_zone1
```

4. Specify the static NAT rule match criteria for the traffic coming from zone 1.

```
[edit security nat]
user@host# set static rule-set sctp-natpt-from-zone1 rule r1-dst match
  destination-address 2b::2:2:2/128
user@host# set static rule-set sctp-natpt-from-zone1 rule r1-dst then static-nat
  prefix 2.2.2.2/32
user@host# set static rule-set sctp-natpt-from-zone1 rule r3-dst match
  destination-address 2d::4:4:4/128
user@host# set static rule-set sctp-natpt-from-zone1 rule r3-dst then static-nat
  prefix 4.4.4.4/32
```

5. Configure rules for the second static NAT zone.

```
[edit security nat]
user@host# set static rule-set sctp-natpt-from-zone2 from zone sctp_zone2
```

6. Specify the static NAT rule match criteria for the traffic coming from zone 2.

```
[edit security nat]
user@host# set static rule-set sctp-natpt-from-zone2 rule r2-dst match
  destination-address 1.1.1.1/32
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set static rule-set sctp-natpt-from-zone2 rule r2-dst then static-nat
  prefix 2a::1:1:1/128
user@host# set static rule-set sctp-natpt-from-zone2 rule r4-dst match
  destination-address 3.3.3.3/32
user@host# set static rule-set sctp-natpt-from-zone2 rule r4-dst then static-nat
  prefix 2c::3:3:3/128
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show security nat static** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces
  ge-4/0/0 {
  unit 0 {
    family inet {
      address 1.1.1.100/24;
    }
    family inet6 {
      address 2a::1:1:100/96;
    }
  }
}
ge-4/0/1 {
  unit 0 {
    family inet {
      address 2.2.2.100/24;
    }
    family inet6 {
      address 2b::2:2:100/96;
    }
  }
}
ge-4/0/2 {
  unit 0 {
    family inet {
      address 3.3.3.100/24;
    }
    family inet6 {
      address 2c::3:3:100/96;
    }
  }
}
ge-4/0/3 {
  unit 0 {
    family inet {
      address 4.4.4.100/24;
    }
    family inet6 {
      address 2d::4:4:100/96;
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security zones
security-zone sctp_zone1 {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
      ge-4/0/0.0;
      ge-4/0/2.0;
   }
}
security-zone sctp_zone2 {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
      ge-4/0/1.0;
      ge-4/0/3.0;
   }
}
```

```
[edit]
user@host# show security nat static
rule-set sctp-natpt-from-zone1 {
   from zone sctp_zone1;
   rule r1-dst {
      match {
         destination-address 2b::2:2:2/128;
      }
      then {
         static-nat {
            prefix {
               2.2.2.2/32;
            }
         }
      }
   }
   rule r3-dst {
      match {
         destination-address 2d::4:4:4/128;
      }
      then {
         static-nat {
            prefix {
               4.4.4.4/32;
            }
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
  rule-set sctp-natpt-from-zone2 {
    from zone sctp_zone2;
    rule r2-dst {
      match {
        destination-address 1.1.1.1/32;
      }
      then {
        static-nat {
          prefix {
            2a::1:1:1/128;
          }
        }
      }
    }
    rule r4-dst {
      match {
        destination-address 3.3.3.3/32;
      }
      then {
        static-nat {
          prefix {
            2c::3:3:3/128;
          }
        }
      }
    }
  }
```

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

Verification

Verifying the Configuration

Purpose Verify that the NAT-PT configuration between an IPv4 endpoint and an IPv6 endpoint is correct.

Action From operational mode, enter the **show security zones** and **show security nat static rule all** commands.

```
user@host> show security zones

Security zone: sctp_zone1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 2
  Interfaces:
    ge-4/0/0.0
    ge-4/0/2.0

Security zone: sctp_zone2
  Send reset for non-SYN session TCP packets: Off
```


Copyright © 2016, Juniper Networks, Inc.

```
  Policy configurable: Yes
  Interfaces bound: 2
  Interfaces:
    ge-4/0/1.0
    ge-4/0/3.0

user@host> show security nat static rule all
Total static-nat rules: 4
Total referenced IPv4/IPv6 ip-prefixes: 4/4

Static NAT rule: r1-dst               Rule-set: sctp-natpt-from-zone1
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : sctp_zone1
  Destination addresses     : 2b::2:2:2
  Host addresses            : 2.2.2.2
  Netmask                   : 128
  Host routing-instance     : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0

Static NAT rule: r3-dst               Rule-set: sctp-natpt-from-zone1
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : sctp_zone1
  Destination addresses     : 2d::4:4:4
  Host addresses            : 4.4.4.4
  Netmask                   : 128
  Host routing-instance     : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0

Static NAT rule: r2-dst               Rule-set: sctp-natpt-from-zone2
  Rule-Id                   : 3
  Rule position             : 3
  From zone                 : sctp_zone2
  Destination addresses     : 1.1.1.1
  Host addresses            : 2a::1:1:1
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0

Static NAT rule: r4-dst               Rule-set: sctp-natpt-from-zone2
  Rule-Id                   : 4
  Rule position             : 4
  From zone                 : sctp_zone2
  Destination addresses     : 3.3.3.3
  Host addresses            : 2c::3:3:3
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The **show security zones** command displays all the zones configured and the interfaces associated with the zone. The **show security nat static rule all** command displays all the static NAT rules configured.

Related Documentation

- Understanding Network Address Translation-Protocol Translation on page 2270

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 101

Configuring GGSN

- Understanding GGSN Redirection on page 2279
- GGSN Pooling Scenarios Overview on page 2279
- Example: Configuring a GGSN Custom Policy on page 2283
- Example: Configuring Custom Applications on page 2286

Understanding GGSN Redirection

Junos OS supports GPRS tunneling protocol (GTP) traffic and gateway GPRS support node (GGSN) redirection. A GGSN (X) can send create-pdp-context responses in which it can specify different GGSN IP addresses (GGSN Y and GGSN Z) for subsequent GTP-C and GTP-U messages. Consequently, the SGSN sends the subsequent GGSN tunneling protocol, control (GTP-C) and GGSN tunneling protocol, user plane (GTP-U) messages to GGSNs Y and Z, instead of X.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

GGSN Pooling Scenarios Overview

The General Packet Radio Service (GPRS) tunneling protocol (GTP) supports different Gateway GPRS Support Node (GGSN) IP addresses during a tunnel creation procedure. There are two GGSN pooling scenarios that support Serving GPRS Support Node (SGSN) roaming.

- Understanding GGSN Pooling for Scenario 1 on page 2279
- Understanding GGSN Pooling for Scenario 2 on page 2281

Understanding GGSN Pooling for Scenario 1

In Figure 102, a packet data protocol (PDP) context request is sent from SGSN A to GGSN B during a GTP tunnel creation procedure. After sending the PDP context request message, GGSN D records the request information and it uses a different destination IP address from the request packet's destination IP address to send the response message to SGSN A.

Copyright © 2016, Juniper Networks, Inc.

In this scenario, two GTP packet messages are sent to Services Processing Unit 1 (SPU1) and SPU2 by the central point, and the messages are processed by SPU1 and SPU2 individually. The session is created on SPU1 and SPU 2 for each GTP packet. SPU1 records the request packet information and SPU2 records the response packet information.

The PDP response message sent from GGSN D to SGSN A is dropped because of a lack of request information. Thus the GTP tunnel is not established.

Figure 102: GGSN Pooling Scenario 1

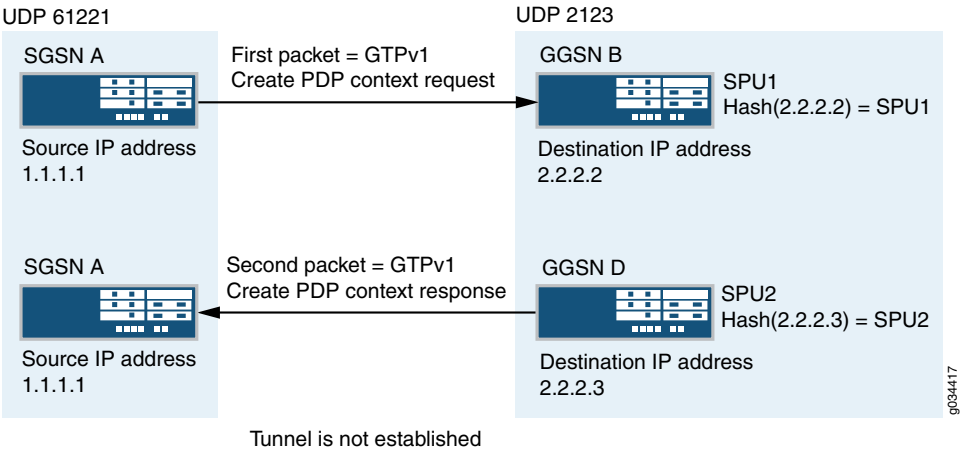

NOTE: SPU2 cannot retrieve request information from SPU1.

Install Request Information to Remote SPU

In this scenario, the PDP response packet was dropped because of a lack of request information, and the GTP tunnel was not established. This can be resolved by installing the request information on the correct SPU.

In Figure 103, when creating a tunnel, the response packet's GGSN IP address changes, triggering a new session, and the central point distributes the message to another SPU.

The response packet always sends to the request packet's source address to the SPU. This helps to install the request information to the remote SPU for the next response packet.

Install the request information into the predictable SPU, HASH (req-src-ip) function while processing the request packet. After the expected SPU number (Hash (1.1.1.1) = SPU2) is calculated by the source IP address of the PDP request message, the request information is installed in the remote SPU2 through the Juniper Message Passing Interface (JMPI).

Copyright © 2016, Juniper Networks, Inc.

Figure 103: Functionality : GGSN Pooling Scenario 1

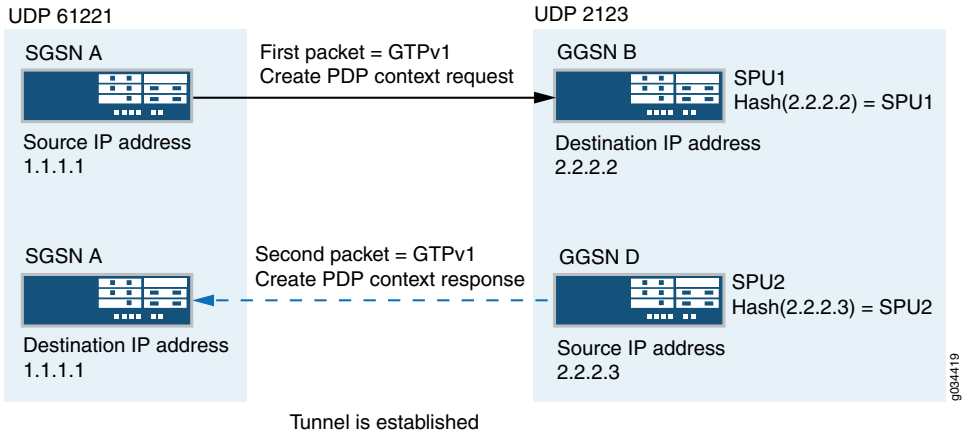

Now the request information is installed on local SPU1 and remote SPU2, so the PDP response message is allowed.

Workarounds for Scenario 1

In scenario 1, the PDP context request message sent from SGSN A reached the Junos OS default application **junos-gprs-gtp** and the GTP inspection was enabled for PDP context request message. However, the PDP context response message sent from GGSN D cannot reach the Junos OS default application **junos-gprs-gtp**. Thus the packet will not be inspected by the GTP module.

As a workaround, you need to enable GTP inspection for the PDP context response message by configuring the custom policy and applications. See the following examples:

- Example: Configuring a GGSN Custom Policy on page 2283
- Example: Configuring Custom Applications on page 2286

Understanding GGSN Pooling for Scenario 2

In Figure 104, a packet data protocol (PDP) context request is sent from SGSN A to GGSN B during a GTP tunnel creation procedure. After receiving the PDP context request message, GGSN B sends the PDP context response message to SGSN A. After receiving the PDP context response message, the GTP-C tunnel is created between SGSN C and GGSN D. Then SGSN C sends an update PDP context request message to GGSN D using different source and destination IP addresses from the request packet's IP header.

In scenario 2, the SGSN A creates the first GTP context request and sends it to the SPU by the central point. The session is created for the request packet on SPU1. The response packet sent from GGSN B to SGSN A reaches the session correctly.

The request packet sent from SGSN A indicates that GTP-C is established on control IP 1.1.1.2 and the GTP-U is established on data IP 1.1.1.3. Likewise, the response message sent from GGSN indicates that GTP-C is established on control IP 2.2.2.3 and GTP-U is established on data IP 2.2.2.4.

Copyright © 2016, Juniper Networks, Inc.

The GTP-C and GTP-U tunnels are created on local SPU1 after all the endpoints are established. However, the tunnel is not established on SPU 2, so the PDP update request message received from SPU2 is dropped.

Figure 104: GGSN Pooling Scenario 2

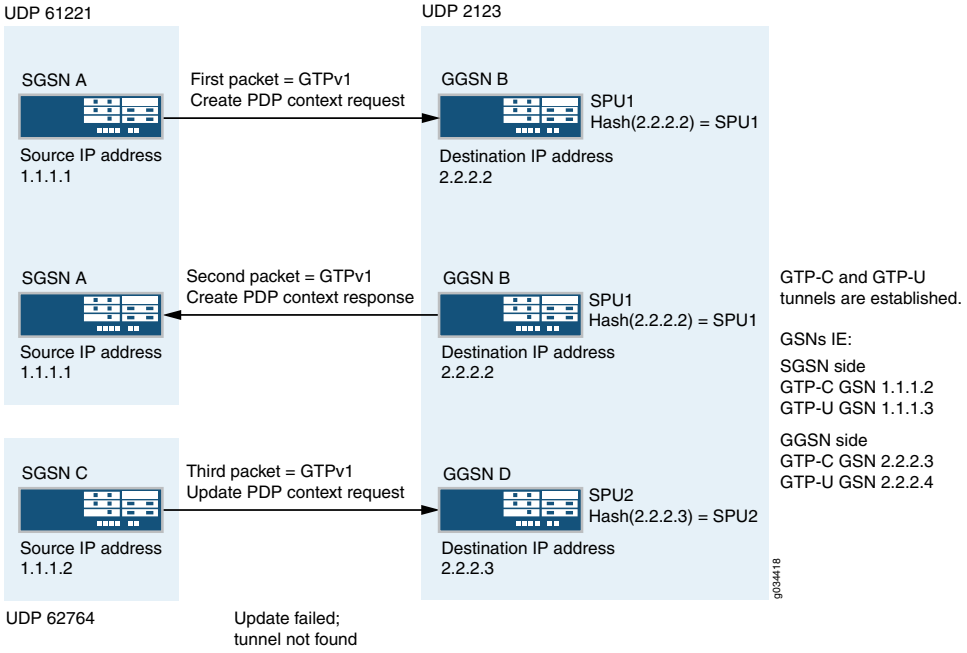

Install Tunnel Information to Remote SPU

In scenario 2, the update request packet is dropped because of a lack of tunnel information. This can be resolved by installing the tunnel information to the correct SPU after the request and response packets are processed. The SPU that receives the response packet installs the tunnel information on the local or remote SPU.

In Figure 105, after the tunnel is established, the control messages are sent to the control tunnel endpoint. The destination IP address of all the control messages should be the control tunnel's GGSN endpoint IP address. This helps to calculate the remote SPU number in advance for the subsequent control message.

Install the tunnel information into the predictable SPU. After the SPU number is calculated by the control tunnel GGSN endpoint IP, the tunnel information is installed in the remote SPU through JMPI.

Copyright © 2016, Juniper Networks, Inc.

Figure 105: Functionality : GGSN Pooling Scenario 2

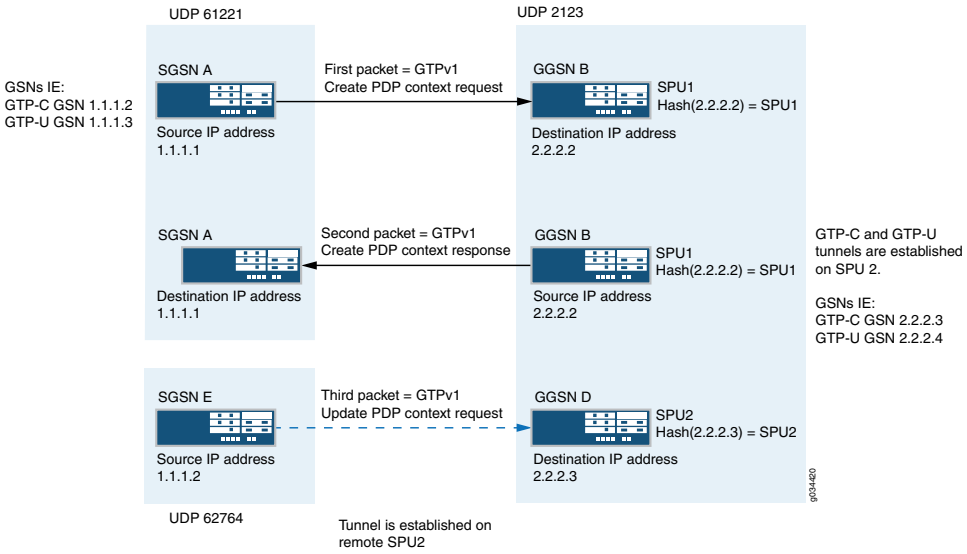

Now the tunnel information is installed on remote SPU2, so the PDP update response message is allowed.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTP on page 2223
- Supported GTP Message Types on page 2235

Example: Configuring a GGSN Custom Policy

This example shows how to configure a Gateway GPRS Support Node (GGSN) custom policy to support GGSN pooling scenario 1.

- Requirements on page 2283
- Overview on page 2284
- Configuration on page 2284
- Verification on page 2285

Requirements

This example uses the following hardware and software components:

- A high-end SRX Series device
- A PC
- Junos OS Release 12.1X44-D10

Before you begin, you should be familiar with GGSN pooling scenarios. See "GGSN Pooling Scenarios Overview" on page 2279.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you set security zones from zone ggsn and to zone sgsn. Next you set the GGSN policy name to ggsn-pool-g2s. You set the name of the match source address to ggsn-1 and the match destination address to sgsn-1.

Then you set the port based application to src_2123 and src_3386. You set the action type to permit. Then you set the application services name to gprs-gtp-profile and the GTP profile name to test. Finally, you set the default policy name to deny-all.

Configuration

Configuring a GGSN Custom Policy

CLI Quick Configuration

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

```
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
  source-address ggsn-1
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
  destination-address sgsn-1
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match application
  src_2123
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match application
  src_3386
set security policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s then permit
  application-services gprs-gtp-profile test
set security policies default-policy deny-all
```

Step-by-Step Procedure

To configure a GGSN custom policy:

1. Configure the GGSN custom policy.

   ```
   [edit security ]
   user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s
   ```

2. Configure the source address.

   ```
   [edit security]
   user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
     source-address ggsn-1
   ```

3. Configure the destination address.

   ```
   [edit security]
   user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
     destination-address sgsn-1
   ```

4. Configure the policy applications.

   ```
   [edit security]
   user@host#set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
     application src_2123
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s match
  application src_3386
```

5. Configure the activity type and specify the GTP profile name.

```
[edit security]
user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s then
  permit
user@host# set policies from-zone ggsn to-zone sgsn policy ggsn-pool-g2s then
  permit application-services gprs-gtp-profile test
```

6. Configure the default policy.

```
[edit security]
user@host#  set policies default-policy deny-all
```

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

```
[edit]
user@host# show security policies
  from-zone zone-name to-zone zone-name {
    from-zone ggsn to-zone sgsn {
      policy ggsn-pool-g2s {
        match {
          source-address ggsn-1;
          destination-address sgsn-1;
          application [ src_2123 src_3386 ];
        }
        then {
          permit {
            application-services {
              gprs-gtp-profile test;
            }
          }
        }
      }
      }
    }
    default-policy {
      deny-all;
    }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying the Configuration on page 2285

Verifying the Configuration

Purpose Verify that the GGSN custom policy configuration is correct.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security** command.

Sample Output

```
user@host>show security policies
From zone: sgsn, To zone: ggsn
  Policy: ggsn-pool-g2s, State: enabled, Index: 5, Scope Policy: 0, Sequence
number: 1
    Source addresses: ggsn1
    Destination addresses: sgsn1
    Applications: src_2123 src_3386
    Action: permit, application services: gprs-gtp-profile test
    Default policy: Deny-all
```

This output shows a summary of policy configuration.

Related Documentation

- Example: Configuring Custom Applications on page 2286
- GGSN Pooling Scenarios Overview on page 2279

Example: Configuring Custom Applications

This example shows how to configure custom applications to support GGSN pooling scenario 1.

- Requirements on page 2286
- Overview on page 2286
- Configuration on page 2287

Requirements

This example uses the following hardware and software components:

- A high-end SRX Series device
- A PC
- Junos OS Release 12.1X44-D10

Before you begin, configure the required GGSN policy. See "Example: Configuring a GGSN Custom Policy" on page 2283.

Overview

In this example, you create applications src_2123 and src_3386 to identify source ports 2123 and 3386 for both TCP and UDP.

First you configure a custom application called src_2123.You set the application protocol to gprs-gtp-c. Then you set the networking protocol type to UDP. You set the source port to 2123 and the destination port to 0-0.

Then you configure another custom application called src_3386. You set the application protocol to gprs-gtp-v0. Then you set the networking protocol type to UDP. Finally, you set the source port to 3386 and the destination port to 0-0.

 Copyright © 2016, Juniper Networks, Inc.

Configuration

Configuring Custom Applications

CLI Quick Configuration

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

```
set applications application src_2123 application-protocol gprs-gtp-c
set applications application src_2123 protocol udp
set applications application src_2123 source-port 2123
set applications application src_2123 destination-port 0-0
set applications application src_3386 application-protocol gprs-gtp-v0
set applications application src_3386 protocol udp
set applications application src_3386 source-port 3386
set applications application src_3386 destination-port 0-0
```

Step-by-Step Procedure

To configure custom policy applications:

1. Configure the first custom application and application protocol name.

   ```
   [edit applications]
   user@host# set application src_2123 application-protocol gprs-gtp-c
   ```

2. Configure the networking protocol type.

   ```
   [edit applications]
   user@host# set application src_2123 protocol udp
   ```

3. Configure the source port number.

   ```
   [edit applications]
   user@host#set application src_2123 source-port 2123
   ```

4. Configure the TCP or UDP destination port number.

   ```
   [edit applications]
   user@host# set application src_2123 destination-port 0-0
   ```

5. Configure the second custom application and application protocol name.

   ```
   [edit applications]
   user@host# set application src_3386 application-protocol gprs-gtp-v0
   ```

6. Configure the networking protocol type.

   ```
   [edit applications]
   user@host# set application src_3386 protocol udp
   ```

7. Configure the source port number.

   ```
   [edit applications]
   user@host#  set application src_3386 source-port 3386
   ```

8. Configure the destination port number.

   ```
   [edit applications]
   user@host# set application src_3386 destination-port 0-0
   ```

Copyright © 2016, Juniper Networks, Inc.

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

```
[edit]
user@host# show applications
  application src_2123 {
    application-protocol gprs-gtp-c;
    protocol udp;
    source-port 2123;
    destination-port 0-0;
  }
  application src_3386 {
    application-protocol gprs-gtp-v0;
    protocol udp;
    source-port 3386;
    destination-port 0-0;
  }
```

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

Related Documentation

- Example: Configuring a GGSN Custom Policy on page 2283
- GGSN Pooling Scenarios Overview on page 2279

Copyright © 2016, Juniper Networks, Inc.

PART 31

Configuring GPRS Tunnel Protocol v2

- Configuring GTPv2 on page 2291
- Configuring GTPv2 Message Filtering on page 2303
- GTPv2 Information Elements Overview on page 2313

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 102

Configuring GTPv2

- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTP Path Restart on page 2296
- Example: Restarting a GTPv2 Path on page 2297
- Understanding GTPv2 Tunnel Cleanup on page 2298
- Example: Setting the Timeout Value for GTPv2 Tunnels on page 2299
- Understanding GTPv2 Traffic Logging on page 2300
- Example: Enabling GTPv2 Traffic Logging on page 2301

Understanding GTPv2

The GPRS tunneling protocol (GTP) establishes a GTP tunnel between a Serving GPRS Support Node (SGSN) and a Gateway GPRS Support Node (GGSN) for individual Mobile Stations (MS). In previous releases, only GTP version 0 (GTPv0) and GTP version 1 (GTPv1) were deployed. GTP version 2 (GTPv2) is implemented in the Junos operating system (Junos OS) Release 11.4.

GTPv2 is part of Long Term Evolution (LTE), a fourth generation (4G) wireless broadband technology developed by Third-Generation Partnership Project (3GPP). 3GPP is the standard body for developing GPRS standards. LTE is designed to increase the capacity and speed of mobile telephone networks. GTPv2 is a protocol designed for LTE networks. An LTE network comprises network elements, LTE interfaces, and protocols.

GTPv0 and GTPv1 are implemented using SGSNs and GGSNs. However, in GTPv2, the traditional SGSNs and GGSNs are replaced by three logical nodes—a serving gateway (SGW), a packet data network gateway (PGW), and a mobility management entity (MME).

Figure 106 shows the following LTE interfaces where SRX Series devices are deployed in the public land mobile network (PLMN).

Copyright © 2016, Juniper Networks, Inc.

Figure 106: LTE Interfaces

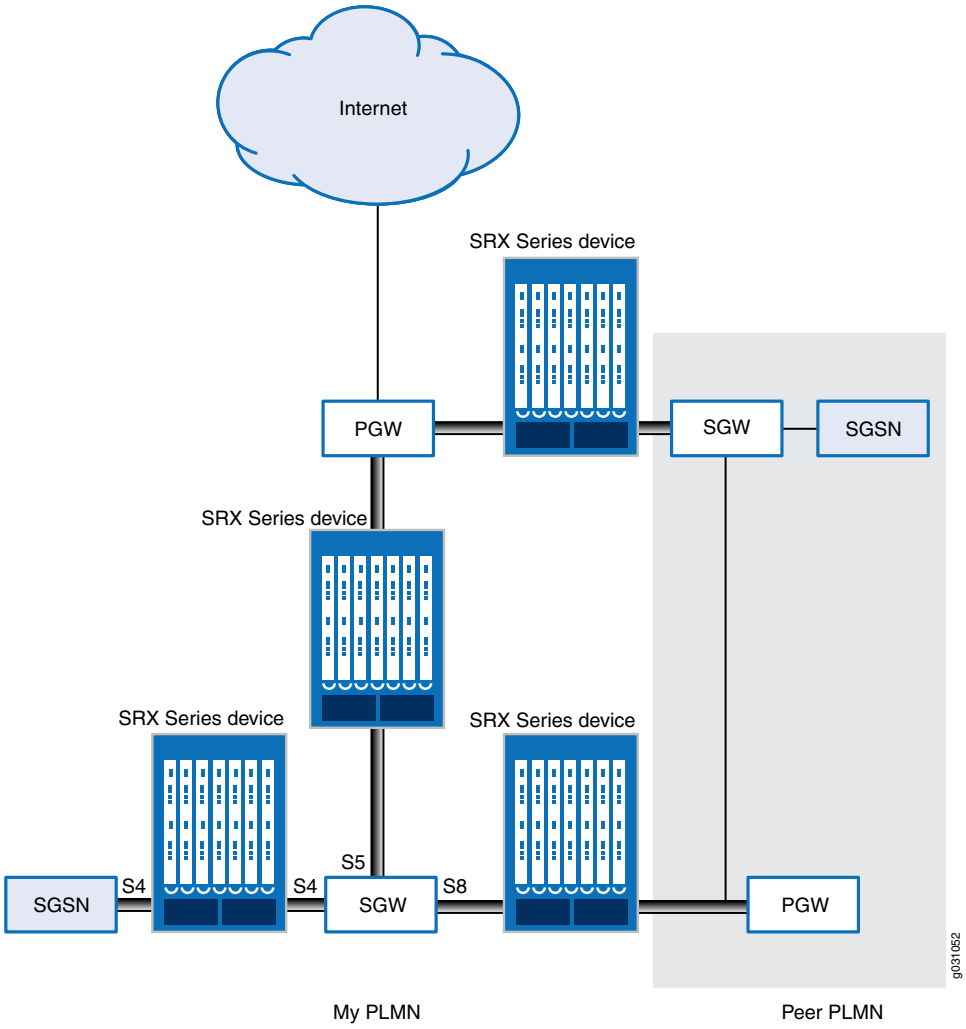

- S5—This interface connects an SGW and a PGW. It provides user plane tunneling and tunnel management capability between the SGW and the PGW. It is also used for SGW relocation that happens because of user equipment mobility or SGW connection to a non-collocated PGW. The S5 interface is equivalent to the Gn interface in a Third Generation (3G) mobile network.
- S8—This interface connects an SGW in a visited PLMN (VPLM) and a PGW in a home PLMN (HPLMN). S8 is the inter-PLMN variant of S5. The S8 interface is equivalent to the Gp interface in a 3G mobile network.
- S4—This interface connects an S4 SGSN and an SGW. It provides related control and mobility support between GPRS core network and 3GPP Anchor function. It also provides user plane tunneling if direct tunneling is not established. The S4 interface does not have any equivalent interface in the 3G mobile network, because it provides interoperability between 3G and 4G networks.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- GPRS Overview on page 2215
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Understanding Policy-Based GTPv2

GPRS tunneling protocol version 2 (GTPv2) implements a policy mechanism that checks every GTPv2 packet against security policies that regulate GTPv2 traffic. Based on the security policy, the packet is then forwarded, dropped, or tunneled.

A GTPv2 security policy allows you to forward, deny, or tunnel GTPv2 traffic. However, the security policy does not enable GTPv2 traffic inspection on the device. To enable traffic inspection, you must apply a GTPv2 inspection object to a security policy. A GTPv2 inspection object is a set of configuration parameters for processing GTPv2 traffic.

You can apply only one GTPv2 inspection object per security policy. However, you can apply an inspection object to multiple security policies.

NOTE: **By default, a GTPv2 inspection object is not applied to a security policy. You need to explicitly apply an inspection object to a security policy.**

Using GTPv2 security policies, you can permit or deny GTPv2 tunnel establishment from certain peers, such as a serving gateway (SGW). You can configure GTPv2 security policies that specify multiple source and destination addresses, address groups, or an entire zone.

Related Documentation

- Example: Enabling GTPv2 Inspection in Policies on page 2293
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Example: Enabling GTPv2 Inspection in Policies

This example shows how to enable GTPv2 inspection in policies.

- Requirements on page 2294
- Overview on page 2294
- Configuration on page 2294
- Verification on page 2296

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, the device must be restarted after GTPv2 is enabled. By default, GTPv2 is disabled on the device.

Overview

In this example, you configure interfaces as ge-0/0/1 and ge-0/0/2, and assign them the interface addresses 4.0.0.254/8 and 5.0.0.254/8, respectively. You then configure the security zones and specify the global addresses as 4.0.0.5/32 and 5.0.0.6/32, respectively. You enable GTPv2 inspection in security policies to allow bidirectional traffic between two networks within the same public land mobile network (PLMN).

Configuration

CLI Quick Configuration

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

```
set security gprs gtp profile gtp2
set interfaces ge-0/0/1 unit 0 family inet address 4.0.0.254/8
set interfaces ge-0/0/2 unit 0 family inet address 5.0.0.254/8
set security zones security-zone sgw1 interfaces ge-0/0/1.0 host-inbound-traffic
  system-services all
set security zones security-zone sgw1 host-inbound-traffic protocols all
set security zones security-zone pgw1 interfaces ge-0/0/2.0 host-inbound-traffic
  system-services all
set security zones security-zone pgw1 host-inbound-traffic protocols all
set security address-book global address local-sgw1 4.0.0.5/32
set security address-book global address remote-pgw1 5.0.0.6/32
set security policies from-zone sgw1 to-zone pgw1 policy sgw1_to_pgw1 match
  source-address local-sgw1 destination-address remote-pgw1 application junos-gprs-gtp
set security policies from-zone sgw1 to-zone pgw1 policy sgw1_to_pgw1 then permit
  application-services gprs-gtp-profile gtp2
set security policies from-zone pgw1 to-zone sgw1 policy pgw1_to_sgw1 match
  source-address remote-pgw1 destination-address local-sgw1 application junos-gprs-gtp
set security policies from-zone pgw1 to-zone sgw1 policy pgw1_to_sgw1 then permit
  application-services gprs-gtp-profile gtp2
```

Step-by-Step Procedure

To configure GTPv2 inspection in policies:

1. Enable GTPv2.

```
[edit]
user@host# set security gprs gtp enable
user@host# commit
user@host# exit
user@host# request system reboot
```

NOTE: After enabling GTPv2, you must reboot the device for GTPv2 inspection to take effect.

Copyright © 2016, Juniper Networks, Inc.

2. Create the GTPv2 inspection object.

```
[edit]
user@host# set security gprs gtp profile gtp2
```

3. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 4.0.0.254/8
user@host# set ge-0/0/2 unit 0 family inet address 5.0.0.254/8
```

4. Configure the security zones.

```
[edit security zones]
user@host# set security-zone sgw1 interfaces ge-0/0/1.0
user@host# set security-zone sgw1 host-inbound-traffic system-services all
user@host# set security-zone sgw1 host-inbound-traffic protocols all
user@host# set security-zone pgw1 interfaces ge-0/0/2.0
user@host# set security-zone pgw1 host-inbound-traffic system-services all
user@host# set security-zone pgw1 host-inbound-traffic protocols all
```

5. Specify the addresses.

```
[edit security address-book global]
user@host# set address local-sgw1 4.0.0.5/32
user@host# set address remote-pgw1 5.0.0.6/32
```

6. Enable GTPv2 inspection in the security policies.

```
[edit security policies]
user@host# set from-zone sgw1 to-zone pgw1 policy sgw1_to_pgw1 match
  source-address local-sgw1 destination-address remote-pgw1 application
  junos-gprs-gtp
user@host# set from-zone sgw1 to-zone pgw1 policy sgw1_to_pgw1 then permit
  application-services gprs-gtp-profile gtp2
user@host# set from-zone pgw1 to-zone sgw1 policy pgw1_to_sgw1 match
  source-address remote-pgw1 destination-address local-sgw1 application
  junos-gprs-gtp
user@host# set from-zone pgw1 to-zone sgw1 policy pgw1_to_sgw1 then permit
  application-services gprs-gtp-profile gtp2
```

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

```
[edit]
user@host# show security policies
from-zone sgw1 to-zone pgw1 {
  policy sgw1_to_pgw1 {
  match {
  source-address local-sgw1;
    destination-address remote-pgw1;
    application junos-gprs-gtp;
    }
    then {
      permit {
      application-services {
      gprs-gtp-profile gtp2;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        }
      }
      }
    }
    from-zone pgw1 to-zone sgw1 {
      policy pgw1_to_sgw1 {
      match {
      source-address remote-pgw1;
        destination-address local-sgw1;
        application junos-gprs-gtp;
      }
        then {
          permit {
          application-services {
          gprs-gtp-profile gtp2;
        }
      }
      }
      }
    }
    default-policy {
      permit-all;
    }
```

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

Verification

Confirm that the configuration is working properly.

Verifying GTPv2 Inspection in Policies

Purpose Verify that GTPv2 inspection is enabled.

Action From operational mode, enter the **show security policies** command.

Related Documentation

- Understanding Policy-Based GTPv2 on page 2293
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Understanding GTP Path Restart

Restarting a GPRS tunneling protocol (GTP) path terminates all GTP tunnels between two devices. Each GTP gateway is associated with a restart number. You can obtain a restart number from the Recovery information element (IE) of a GTP message.

Copyright © 2016, Juniper Networks, Inc.

You can detect a restart by comparing the locally stored restart number with the newly obtained one. The locally stored restart number is a nonzero value and does not match with the new restart number.

You can use the **set security gprs gtp profile name restart-path (echo | create | all)** configuration statement to restart a GTP path.

After you configure this command, the device detects the changed restart number obtained from the Recovery IE in the messages. You can use the **echo** option to obtain a new restart number from echo messages, the **create** option to obtain a restart number from create-session messages, or the **all** option to obtain a new restart number from all types of GTP messages.

Related Documentation

- Example: Restarting a GTPv2 Path on page 2297
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303
- Understanding GTPv2 Information Elements on page 2313

Example: Restarting a GTPv2 Path

This example shows how to restart a GTPv2 path.

- Requirements on page 2297
- Overview on page 2297
- Configuration on page 2298
- Verification on page 2298

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

NOTE: For brevity, this example uses GTPv2.

In this example, you restart the GTPv2 path for the GTPv2 inspection object named gtp2. You obtain a new restart number from the Recovery information element (IE) in an echo message.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To restart the GTPv2 path:

1. Specify the GTPv2 profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2
   ```

2. Restart the path.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 restart-path echo
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTP Path Restart on page 2296
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303
- Understanding GTPv2 Information Elements on page 2313

Understanding GTPv2 Tunnel Cleanup

A GPRS tunneling protocol version 2 (GTPv2) tunnel enables transmission of GTPv2 traffic between GPRS support nodes (GSNs).

While transmitting traffic, GTPv2 tunnels might hang for a number of reasons. For example, delete-pdp-request messages might get lost in the network, or a GSN might not shut down properly. In such a case, you can remove hanging GTPv2 tunnels either automatically or manually.

To remove a hanging GTPv2 tunnel automatically, you need to set a GTPv2 tunnel timeout value on the device. The device automatically identifies and removes a tunnel that is idle for the period specified by the timeout value. The default GTPv2 tunnel timeout value is 36 hours.

You can use the **set security gprs gtp profile name timeout** configuration statement to configure this value on the device. The timeout range is 1 through 1000 hours.

Copyright © 2016, Juniper Networks, Inc.

To remove a hanging GTPv2 tunnel manually, you need to use the **clear security gprs gtp tunnel** operational mode command.

NOTE: The default GTP tunnel timeout value is changed from 24 hours to 36 hours in Junos operating system (Junos OS) Release 11.4. The reason being that GTP does not inspect GTP-U (user plane) traffic in Junos OS Release 11.4.

Related Documentation

- Example: Setting the Timeout Value for GTPv2 Tunnels on page 2299
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Example: Setting the Timeout Value for GTPv2 Tunnels

This example shows how to set the timeout value for GTPv2 tunnels.

- Requirements on page 2299
- Overview on page 2299
- Configuration on page 2299
- Verification on page 2300

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set the tunnel timeout value to 40 hours for the GTPv2 inspection object named gtp2.

Configuration

Step-by-Step Procedure

To configure the GTPv2 tunnel timeout value:

1. Specify the GTPv2 profile.

```
[edit]
user@host# set security gprs gtp profile gtp2
```

2. Specify the timeout value.

```
[edit]
user@host# set security gprs gtp profile gtp2 timeout 40
```

Copyright © 2016, Juniper Networks, Inc.

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTPv2 Tunnel Cleanup on page 2298
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Understanding GTPv2 Traffic Logging

You can use the console or syslog to view GPRS tunneling protocol version 2 (GTPv2) traffic logs. You can configure the device to log GTPv2 packets based on their status. GTPv2 packet status can be any of the following:

- Forwarded—GTPv2 packet was forwarded because it was valid.
- State-invalid—GTPv2 packet was dropped because it failed stateful inspection or a sanity check. In case of a sanity check failure, the packet is marked as sanity.
- Prohibited—GTPv2 packet was dropped because it failed message length, message type, or International Mobile Subscriber Identity (IMSI) prefix checks.
- Rate-limited—GTPv2 packet was dropped because it exceeded the maximum rate limit of the destination GPRS support node (GSN).

By default, GTPv2 logging is disabled on the device. You can use the **set security gprs gtp profile name log** configuration statement to enable GTPv2 logging on the device.

Related Documentation

- Example: Enabling GTPv2 Traffic Logging on page 2301
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303
- Understanding GTPv2 Information Elements on page 2313

Copyright © 2016, Juniper Networks, Inc.

Example: Enabling GTPv2 Traffic Logging

This example shows how to enable GTPv2 traffic logging on a device.

- Requirements on page 2301
- Overview on page 2301
- Configuration on page 2301
- Verification on page 2301

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable GTPv2 traffic logging for forwarded GTPv2 packets.

Configuration

Step-by-Step Procedure

To enable GTPv2 traffic logging for forwarded GTPv2 packets:

1. Specify the GTPv2 profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2
   ```

2. Enable logging for GTPv2 forwarded packets.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 log forwarded basic
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTPv2 Traffic Logging on page 2300
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303
- Understanding GTPv2 Information Elements on page 2313

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 103

Configuring GTPv2 Message Filtering

- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303
- Example: Permitting and Denying GTPv2 Message Types on page 2307
- Understanding GTPv2 Message-Length Filtering on page 2308
- Example: Setting GTPv2 Message Lengths on page 2308
- Understanding GTPv2 Message-Type Filtering on page 2309
- Understanding GTPv2 Message-Rate Limiting on page 2310
- Example: Limiting the GTPv2 Message Rate on page 2311

Understanding GTPv2 Message Filtering

When a device receives a GPRS tunneling protocol version 2 (GTPv2) packet, it checks the packet against GTPv2 policies configured on the device. If the packet matches a policy, then the device inspects the packet based on the GTPv2 inspection object applied to the policy. If the packet fails to match any of the configuration parameters, it is dropped. If the packet matches all the configuration parameters, it is forwarded.

Related Documentation

- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message-Length Filtering on page 2308
- Understanding GTPv2 Message-Type Filtering on page 2309
- Supported GTPv2 Message Types on page 2303

Supported GTPv2 Message Types

Table 252 lists the message types supported in GTPv2. You can use these message types to configure GTPv2 message-type filtering.

Copyright © 2016, Juniper Networks, Inc.

Table 252: GTPv2 Messages

Message	Message Type
bearer resource command	bearer-resource
bearer resource failure	bearer-resource
change notification request	change-notification
change notification response	change-notification
context request	context
context response	context
context acknowledgement	context
configuration transfer tunnel	config-transfer
create bearer request	create-bearer
create bearer response	create-bearer
create indirect data forwarding tunnel request	create-data-forwarding
create indirect data forwarding tunnel response	create-data-forwarding
create forwarding tunnel request	create-tnl-forwarding
create forwarding tunnel response	create-tnl-forwarding
create session request	create-session
create session response	create-session
CS paging indication	cs-paging
delete bearer request	delete-bearer
delete bearer response	delete-bearer
delete bearer command	delete-command
delete bearer failure	delete-command
delete indirect data forwarding tunnel request	delete-data-forwarding
delete indirect data forwarding tunnel response	delete-data-forwarding
delete PDN connection set request	delete-pdn

Copyright © 2016, Juniper Networks, Inc.

Table 252: GTPv2 Messages *(continued)*

Message	Message Type
delete PDN connection set response	delete-pdn
delete session request	delete-session
delete session response	delete-session
detach notification	detach
detach acknowledgement	detach
downlink data notification	downlink-notification
downlink data acknowledgement	downlink-notification
downlink data notification failure	downlink-notification
echo request	echo
echo response	echo
forward access context notification	fwd-access
forward access context acknowledgement	fwd-access
forward relocation request	fwd-relocation
forward relocation response	fwd-relocation
forward relocation complete	fwd-relocation
forward relocation acknowledgement	fwd-relocation
identification request	identification
identification response	identification
MBMS session start request	mbms-sess-start
MBMS session start response	mbms-sess-start
MBMS session stop request	mbms-sess-stop
MBMS session stop response	mbms-sess-stop
MBMS session update request	mbms-sess-update
MBMS session update response	mbms-sess-update

Copyright © 2016, Juniper Networks, Inc.

Table 252: GTPv2 Messages *(continued)*

Message	Message Type
modify bearer request	modify-bearer
modify bearer response	modify-bearer
modify bearer command	modify-command
modify bearer failure	modify-command
release access-bearer request	release-access
release access-bearer response	release-access
relocation cancel request	relocation-cancel
relocation cancel response	relocation-cancel
resume notification	resume
resume acknowledgement	resume
stop paging indication	stop-paging
suspend notification	suspend
suspend acknowledgement	suspend
trace session activation	trace-session
trace session deactivation	trace-session
update bearer request	update-bearer
update bearer response	update-bearer
update PDN connection set request	update-pdn
update PDN connection set response	update-pdn
version not supported	ver-not-supported

Related Documentation

- Understanding GTPv2 Message Filtering on page 2303
- Understanding GTPv2 Message-Type Filtering on page 2309
- Example: Permitting and Denying GTPv2 Message Types on page 2307s
- GPRS Overview on page 2215

Copyright © 2016, Juniper Networks, Inc.

- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTP Message Types on page 2235

Example: Permitting and Denying GTPv2 Message Types

This example shows how to permit and deny GTPv2 message types.

- Requirements on page 2307
- Overview on page 2307
- Configuration on page 2307
- Verification on page 2307

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, for the gtp2 profile, you configure the device to drop the echo and create-session message types for version 2.

Configuration

Step-by-Step Procedure

To permit and deny GTPv2 message types:

1. Specify the GTPv2 profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2
   ```

2. Drop the echo messages for version 2.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 drop echo 2
   ```

3. Drop the create-session messages for version 2.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 drop create-session 2
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding GTPv2 Message-Type Filtering on page 2309
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Understanding GTPv2 Message-Length Filtering

You can configure a device to drop GPRS tunneling protocol version 2 (GTPv2) packets that do not meet the specified minimum or maximum message lengths. In the GTPv2 header, the message length field indicates the length, in octets, of the GTPv2 payload. The message length does not include the length of the GTPv2 header itself, the UDP header, the TCP header, or the IP header.

The default minimum and maximum GTPv2 message lengths are 0 and 65,535 bytes, respectively. Therefore, you can specify the GTPv2 message length range from 0 to 65,535 bytes.

Related Documentation

- Example: Setting GTPv2 Message Lengths on page 2308
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Example: Setting GTPv2 Message Lengths

This example shows how to set GTPv2 message lengths.

- Requirements on page 2308
- Overview on page 2309
- Configuration on page 2309
- Verification on page 2309

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you configure the minimum GTPv2 message length to 10 octets and the maximum GTPv2 message length to 1500 octets for the GTPv2 inspection object named gtp2.

Configuration

Step-by-Step Procedure

To configure GTPv2 message lengths:

1. Specify the GTPv2 profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2
   ```

2. Specify the minimum message length.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 min-message-length 10
   ```

3. Specify the maximum message length.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 max-message-length 1500
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTPv2 Message-Length Filtering on page 2308
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Understanding GTPv2 Message-Type Filtering

You can configure a device to filter GPRS tunneling protocol version 2 (GTPv2) packets based on their message types. By default, the device permits all GTPv2 message types.

You permit and deny message types based on the GTP version number. For example, you can deny message types for one version while you permit them for the other version.

You can use the **set security gprs gtp profile profile name drop message-type number** configuration statement to discard GTPv2 message types. If the version number is not

Copyright © 2016, Juniper Networks, Inc.

mentioned, message types for all versions are discarded. If a configured message type is not valid for the particular GTP version, the specific configuration does not take effect.

NOTE: **Message types valid for GTP version 1 (GTPv1) might not be valid for GTPv2, and vice versa.**

A GTPv2 message type includes one or many messages. When you permit or deny a message type, you automatically permit or deny all messages of the specified message type. For example, if you drop the identification message type, then you automatically drop the identification-request and identification-response messages. Also, if you drop the create-pdp message type for version 2, then only the create-pdp-request and create-pdp-response messages for version 2 are dropped.

Related Documentation

- Example: Permitting and Denying GTPv2 Message Types on page 2307
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Understanding GTPv2 Message-Rate Limiting

You can configure a device to limit the rate of control traffic going to GPRS tunneling protocol version 2 (GTPv2) gateways. GTPv2 gateways are important resources in a public land mobile network (PLMN) and require protection. You can use the **set security gprs gtp profile name rate-limit messages per second** configuration statement to limit the rate of control traffic to the GTPv2 gateways.

Related Documentation

- Example: Limiting the GTPv2 Message Rate on page 2311
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Copyright © 2016, Juniper Networks, Inc.

Example: Limiting the GTPv2 Message Rate

This example shows how to limit the GTPv2 message rate.

- Requirements on page 2311
- Overview on page 2311
- Configuration on page 2311
- Verification on page 2311

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you limit the rate of incoming GTPv2 messages to 200 packets per second to any GTPv2 gateway.

Configuration

Step-by-Step Procedure

To configure the GTPv2 message rate limit:

1. Specify the GTPv2 profile.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2
   ```

2. Set the rate limit.

   ```
   [edit]
   user@host# set security gprs gtp profile gtp2 rate-limit 200
   ```

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

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTPv2 Message-Rate Limiting on page 2310
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Understanding GTPv2 Message Filtering on page 2303
- Supported GTPv2 Message Types on page 2303

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 104

GTPv2 Information Elements Overview

- Understanding GTPv2 Information Elements on page 2313
- Understanding GTPv2 IMSI Prefix and APN Filtering on page 2313
- Example: Setting a Combined GTPv2 IMSI Prefix and APN Filter on page 2315

Understanding GTPv2 Information Elements

Information elements (IEs) are included in all GPRS tunneling protocol version 2 (GTPv2) control message packets. IEs provide information about GTPv2 tunnels, such as creation, modification, deletion, and status. The Junos operating system (Junos OS) supports IEs consistent with the Third-Generation Partnership Project (3GPP) Release 8.

Related Documentation

- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Understanding GTPv2 IMSI Prefix and APN Filtering

A GPRS support node (GSN) identifies a Mobile Station (MS) by its International Mobile Subscriber Identity (IMSI). An IMSI comprises three elements: the mobile country code (MCC), the mobile network code (MNC), and the Mobile Subscriber Identification Number (MSIN). The MCC is a three-digit number, and the MNC is a two-digit or three-digit number. The MCC and MNC combined constitute the IMSI prefix and identify the mobile subscriber's home network or public land mobile network (PLMN). Therefore, the IMSI prefix acts as the PLMN identifier and is used to identify valid roaming partners.

By default, a device does not perform IMSI prefix filtering on GPRS tunneling protocol version 2 (GTPv2) packets. By setting IMSI prefixes, you configure the device to filter create-session-request messages and permit only GTPv2 packets with IMSI prefixes that match the ones you set.

When you filter GTPv2 packets based on an IMSI prefix, you must also specify an access point name (APN).

Copyright © 2016, Juniper Networks, Inc.

An APN is an information element (IE) included in the header of a GTPv2 packet that provides information about how to reach a network. An APN comprises two elements:

- Network ID—Identifies the name of an external network, such as example.com.
- Operator ID—Uniquely identifies the operators' PLMN, such as mnc123.mcc789.gprs.

For example, example.com.mnc123.mcc789.gprs is an APN for reaching the example.com network through the mnc123.mcc789.gprs operator.

By default, a device does not perform APN filtering on GTPv2 packets. However, you can configure the device to perform APN filtering to restrict access to roaming subscribers to external networks.

You can use the **set security gprs gtp profile profile name apn pattern-string imsi-prefix imsi-prefix-digits action (pass |drop |selection)** configuration statement to filter packets based on the combination of an IMSI prefix and an APN.

To specify an APN, you need to know the network ID or the domain name of the network (for example, example.com) and, optionally, the operator ID. Because the network ID portion of an APN can be very long, you can use the wildcard (*) as the first character of the APN string. For example, if you use *.example.com as the network ID, the wildcard indicates that the APN is not limited only to example.com but also includes all the characters that might precede it.

You can use the **selection** option to set a *selection mode* for the APN. The selection mode indicates the origin of the APN and whether or not the Home Location Register (HLR) has verified the user subscription. You set the selection mode according to the security needs of your network. Possible selection modes include the following:

- ms—MS-provided APN, subscription is not verified.
- net—Network-provided APN, subscription is not verified.
- vrf—MS-provided or network-provided APN, subscription is verified.

You can use the **drop** option to drop all APNs and the **pass** option to pass all APNs for any selection mode.

When performing APN filtering, the device inspects packets to look for APNs that match APNs that you set. If the APN of a packet matches an APN that you specified, then the device verifies the selection mode and forwards the GTPv2 packet.

NOTE: **The device only forwards the GTPv2 packet if both the APN and the selection mode match the APN and the selection mode that you specified.**

Because APN filtering is based on perfect matches, using the wildcard (*) when setting an APN suffix can prevent the inadvertent exclusion of APNs that you would otherwise authorize.

Copyright © 2016, Juniper Networks, Inc.

NOTE: IMSI prefix and APN filtering apply to create-session-request messages only.

Related Documentation

- Example: Setting a Combined IMSI Prefix and APN Filter on page 2251
- Understanding GTPv2 Information Elements on page 2313
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Example: Setting a Combined GTPv2 IMSI Prefix and APN Filter

This example shows how to set a combined GTPv2 IMSI prefix and APN filter.

- Requirements on page 2315
- Overview on page 2315
- Configuration on page 2315
- Verification on page 2316

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set example.com.mnc123.mcc789.gprs as an APN and use the wildcard (*). You set the selection mode as network for this APN. You also set the IMSI prefix as 123678.

Configuration

Step-by-Step Procedure

To set the combined IMSI prefix and APN filter:

1. Specify the GTPv2 profile.

```
[edit]
user@host# set security gprs gtp profile gtp2
```

2. Set the selection mode for the APN.

```
[edit]
user@host# set security gprs gtp profile gtp2 apn
  *example.com.mnc123.mcc789.gprs imsi-prefix 123678 action selection net
```

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

Copyright © 2016, Juniper Networks, Inc.

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

Verification

To verify the configuration is working properly, enter the **show security gprs** command.

Related Documentation

- Understanding GTPv2 IMSI Prefix and APN Filtering on page 2313
- Understanding GTPv2 Information Elements on page 2313
- GPRS Overview on page 2215
- Understanding GTPv2 on page 2291
- Understanding Policy-Based GTPv2 on page 2293
- Example: Enabling GTPv2 Inspection in Policies on page 2293
- Supported GTPv2 Message Types on page 2303

Copyright © 2016, Juniper Networks, Inc.

PART 32

Configuring Stream Control Transmission Protocol

- Configuring SCTP on page 2319

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 105

Configuring SCTP

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Features Overview on page 2323
- Understanding Central Point Architecture Support for SCTP on page 2324
- Understanding SCTP Behavior in Chassis Cluster on page 2324
- Understanding SCTP Multihoming on page 2326
- Understanding SCTP Multichunk Inspection on page 2327
- SCTP Packet Structure Overview on page 2327
- SCTP Configuration Overview on page 2329
- Example: Configuring a GPRS SCTP Profile for Policy-Based Inspection to Reduce Security Risks on page 2329
- Example: Configuring a Security Policy to Permit or Deny SCTP Traffic on page 2332

Understanding Stream Control Transmission Protocol

Stream Control Transmission Protocol (SCTP) is an IP Transport Layer protocol. SCTP exists at an equivalent level with User Datagram Protocol (UDP) and Transmission Control Protocol (TCP), which provides transport layer functions to many Internet applications. SCTP is a reliable transport protocol operating on top of a connectionless packet network such as IP and supports data transfer across the network in single IP or multi-IP cases.

The SCTP module inspects IPv4 and IPv6 traffic and checks all segments of the SCTP packet. The packet is then permitted or dropped based on the policy. For IPv6 traffic, the SCTP module inspects every extension header until it finds the SCTP header, and then only the SCTP header is processed and all the other headers are ignored.

SCTP is used for applications where monitoring and detection of loss of session is required. The SCTP path or session failure detection mechanism, for example, the heartbeat, monitors the connectivity of the session.

Figure 107 illustrates the SCTP 4-way handshake and TCP 3-way handshake.

Copyright © 2016, Juniper Networks, Inc.

Figure 107: SCTP 4-way Handshake and TCP 3-way Handshake

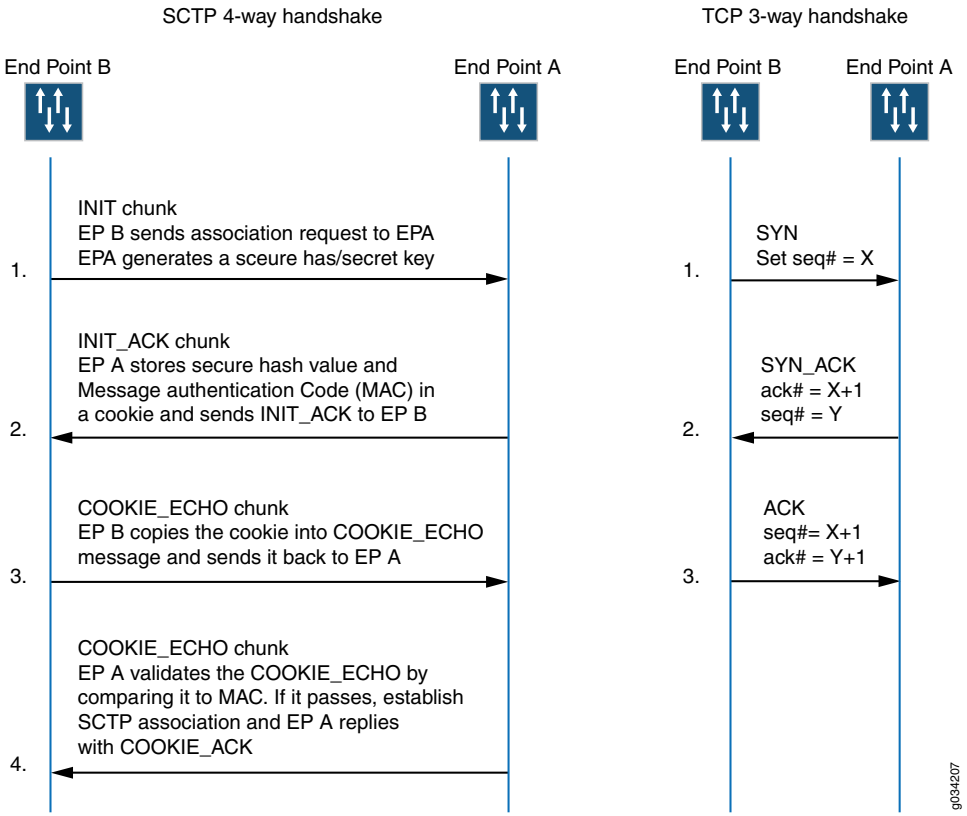

SCTP provides the following services:

- Aggregate Server Access Protocol (ASAP)
- Bearer-independent Call Control (BICC)
- Direct Data Placement Segment chunk (DDP-segment)
- Direct Data Placement Stream session control (DDP-stream)
- Diameter in a DTLS/SCTP DATA chunk (Diameter-DTLS)
- Diameter in a SCTP DATA chunk (Diameter-SCTP)
- DPNSS/DASS 2 extensions to IUA Protocol (DUA)
- Endpoint Handlescape Redundancy Protocol (ENRP)
- H.248 Protocol (H248)
- H.323 Protocol (H323)
- ISDN User Adaptation Layer (IUA)
- MTP2 User Peer-to-Peer Adaptation Layer (M2PA)
- MTP2 User Adaptation Layer (M2UA)
- MTP3 User Adaptation Layer (M3UA)

Copyright © 2016, Juniper Networks, Inc.

- Other unspecified-configured SCTP payload protocols (Others)
- Q.IPC
- Reserved
- S1 Application Protocol (S1AP)
- Simple Middlebox Configuration (SIMCO)
- SCCP User Adaptation Layer (SUA)
- Transport Adapter Layer Interface (TALI)
- V5.2 User Adaptation Layer (V5UA)
- X2 Application Protocol (X2AP)

SCTP can transport signaling messages to and from Signaling System 7 (SS7) for 3G mobile networks through M3UA, M2UA, or SUA. SCTP is a packet-based transport protocol. SCTP provide reliable and secure transport, minimized end-to-end delay, short failover time in case of network failures and both sequence and no-sequence transport.

SCTP is optimized to:

- Avoid the multithread infrastructure problems, when the traffic is high
- Improve the SCTP association searching rate (association lookup process speed is increased) by SCTP hash table optimization on the SPU
- Improve FSM for retransmission cases

SCTP has the following limitations and constraints:

- IP Addresses
  - A maximum of eight source IP addresses and eight destination IP addresses are allowed in an SCTP communication.
  - Only static IP NAT is supported; the interface packets (from one side: client or server) coming in must belong to the same zone.
- Policies
  - Dynamic policy is not supported. You must configure all policies for SCTP sessions.
  - When policies are deleted, the related sessions and associations are cleared.
  - You configure one policy to permit SCTP traffic from all client IPs to all server IPs, and another policy to permit SCTP traffic from server IPs to client IPs. If one policy has an SCTP profile, then the same SCTP profile is needed for the reverse policy.
  - If you configure different policies for each session belonging to one association, there will be multiple policies related to one association, and the SCTP packet management (drop, rate limit, and so on) uses the profile attached to the handling SCTP session's policy.
- SCTP enable/disable is controlled by whether there is a SCTP profile configured.

Copyright © 2016, Juniper Networks, Inc.

- If no profile is attached to a policy, SCTP packets are forwarded without inspection.
- If a profile with the **nat-only** option is attached to a policy,then only NAT translation is done on the SCTP packets matching the policy. If a profile does not have the **nat-only** option set, then both NAT translation and SCTP inspection are done on each SCTP packet matching the policy.
- If you disable SCTP, all associations are deleted, and subsequent SCTP packets are passed or dropped according to the policy.
- If you enable SCTP, all existing SCTP sessions must be cleared or the traffic matching old sessions will be forwarded without any inspection from the SCTP module.

  If you want to enable SCTP again, all the running SCTP communications will be dropped, because no associations exist. New SCTP communications can establish an association and perform the inspections.

  NOTE: **Clear old SCTP sessions when SCTP is reenabled; doing this will avoid any impact caused by the old SCTP sessions on the new SCTP communications.**

- If you add an SCTP profile to an existing policy, you must do one of the following: clear related sessions or remove the old policy and create a new policy.
- If you change the timeout value in the SCTP profile, the configured handshake and the timeout value in existing associations will not change.

- SCTP Rate Limiting
  - Any change in the rate-limiting configuration will not affect the subsequent traffic of existing associations. It will apply to the newly established associations.
  - The supported protocol decimal value is from 0 to 63. This value includes 48 IANA assigned protocols and 16 unassigned protocols.
  - A maximum of 80 addresses are rate limited in one profile.
  - A maximum of 10 protocols are rate limited for one address in one profile.
  - The supported rate limit value is from 1 to 12000.
- SCTP Payload Protocol Blocking
  - Any change in the protocol-blocking configuration immediately impacts the subsequent traffic of existing associations.
  - The supported protocol decimal value is from 0 to 63. This value includes 48 IANA assigned protocols and 16 unassigned protocols.
- An SCTP endpoint can be a multihomed host with either all IPv4 addresses or all IPv6 addresses. An SCTP endpoint also supports NAT-PT in two directions, from an IPv4 address format to an IPv6 address format, and vice versa. SCTP module does not support IPv4 or IPv6 mixed-up multihoming and IPv4 or IPv6 mixed-up NAT-PT.

Copyright © 2016, Juniper Networks, Inc.

- For static NAT to work, the interfaces packets (from one side: client or server side) coming in must belong to the same zone.
- For multihome cases, only IPv4 address parameter or IPv6 address parameter in INIT or INI-ACK is supported.
- Only static NAT is supported for SCTP.
- Only established SCTP associations are synchronized to peer sessions.
- SCTP sessions are not deleted with associations; they time out in 30 minutes, which is the default value. The timeout value is configurable and can be changed.
- If the 4-way handshake process is not handled on one node, and is handled instead on two nodes (for example, two sessions on two nodes in active/active mode) or if the cluster is in failover before the 4-way handshake is completed, the association will not be established successfully.
- One SPU supports a maximum of 20,000 associations and a maximum of 1,280,000 SCTP sessions.

  In some cases, the associations might not be distributed to SPUs very evenly because the ports' hash result on the central point is uneven. For example, this event can occur when only two peers of ports are used, and one peer has 100 associations, but another peer has only one association. In this case, the associations cannot be distributed evenly on the firewall with more than one SPU.
- Unified in-service software upgrade (ISSU) to earlier Junos OS releases is not supported.
- The M3UA/SCCP message parsing is checked , but the M3UA/SCCP stateful inspection is not checked.
- Only ITU-T Rec. Q.711-Q.714 (07/96) standard is supported. ANSI, ETSI, China, and other standards are not supported.
- Only RFC 4960 is supported.
- VPN session affinity does not support GPRS tunneling protocol (GTP) and Stream Control Transmission Protocol (SCTP).

Related Documentation

- SCTP Configuration Overview on page 2329
- SCTP Packet Structure Overview on page 2327
- SCTP Features Overview on page 2323
- Understanding SCTP Behavior in Chassis Cluster on page 2324

SCTP Features Overview

The following are the important features of SCTP:

Copyright © 2016, Juniper Networks, Inc.

- Multihoming support where one or both endpoints of a connection can consist of more than one IP address. This enables transparent failover between redundant network paths.
- Delivery of data in chunks within an independent stream eliminates unnecessary head-of-line blocking.
- Path selection and monitoring functionality to select a primary data transmission path and test the connectivity of the transmission path.
- Validation and acknowledgment mechanisms protect against flooding attacks and provide notification of duplicated or missing data chunks.
- Improved error detection suitable for jumbo Ethernet frames.

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Packet Structure Overview on page 2327
- Understanding SCTP Behavior in Chassis Cluster on page 2324
- Example: Configuring a Security Policy to Permit or Deny SCTP Traffic on page 2332

Understanding Central Point Architecture Support for SCTP

A Stream Control Transmission Protocol (SCTP) association is a connection between two SCTP endpoints. Each SCTP endpoint identifies the association with a tag. During an SCTP association setup, two SCTP endpoints exchange their own tags for receiving packets. During the exchange of packets between two SCTP endpoints, both the source address and the destination address can change in the association life cycle.

Prior to Junos OS Release 15.1X49-D40, all sessions of a given SCTP association are hashed to the same Services Processing Unit (SPU) by the fixed per-association SCTP port pair. However, in some cases, multiple SCTP associations share the same port pair, resulting in a bad load-balancing situation with all traffic being handled by a single SPU. Starting from Junos OS Release 15.1X49-D40, to handle the load-balancing issue, tag-based hash distribution is used to ensure even distribution of SCTP traffic from different associations among all SPUs. A 32-bit connection tag is introduced that uniquely identifies the SCTP sessions. The connection tag for SCTP is the vTag and the connection ID remains 0 if the connection tag is not used by the sessions.

Related Documentation

- Understanding Enhancements to Central Point Architecture for the SRX5000 Line on page 1670

Understanding SCTP Behavior in Chassis Cluster

In a chassis cluster configuration mode, the SCTP configuration and the established SCTP association is synced with the peer device. The SCTP module supports both active-active and active-passive modes.

Copyright © 2016, Juniper Networks, Inc.

The established SCTP association sends a creation or deletion message to the peer whenever an association is created or deleted on the active device. The secondary device adds or deletes an association respectively upon receiving the message from the established SCTP association. SCTP module then registers the corresponding callback function to receive and handle this message. There is no continuous timer sync between the two associations.

SCTP module will register a cold start sync function when a secondary device joins the cluster or reboots. The SCTP cold start function is called to sync all SCTP associations with the peer devices at the same time.

After the switchover, the established SCTP associations will remain functioning, but the associations in the progress of establishment will be lost and the establishment procedure needs to be re-initiated. It is also possible that the associations in the progress of teardown miss the ack message and leaves unestablished SCTP associations in the firewall. These associations will be cleaned up when the timer expires (5 hours by default) due to no activity in the association.

NOTE:

- You should configure all policies for your required SCTP sessions. For example, suppose you have endpoints A and B. Endpoint A has one SCTP association with x number of IPs (IP_a1, IP_a2, IP_a3...IP_ax). Endpoint B has one SCTP association with y number of IPs (IP_b1, IP_b2, IP_b3...IP_by.) The policy on the security device should permit all possible x*y paths in both directions.
- When an SCTP association is removed, the related SCTP sessions still exist and time out by themselves.

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Packet Structure Overview on page 2327
- SCTP Features Overview on page 2323

Copyright © 2016, Juniper Networks, Inc.

Understanding SCTP Multihoming

A Stream Control Transmission Protocol (SCTP) endpoint can be a multihomed host with either all IPv4 addresses or all IPv6 addresses. In Figure 108, endpoint A is connected to an SRX Series device with two IPv4 addresses, and endpoint B is connected to an SRX Series device with two IPv4 addresses. Therefore, endpoint A and endpoint B can set up an association using four different pairs of IP addresses, resulting in four valid paths for communication.

Figure 108: SCTP Multihoming with Two IPv4 Endpoints

In Figure 109, endpoint A is connected to an SRX Series device with two IPv6 addresses, and endpoint B is connected to an SRX Series device with two IPv6 addresses. Therefore, endpoint A and endpoint B can set up an association using four different pairs of IP addresses, resulting in four valid paths for communication.

Figure 109: SCTP Multihoming with Two IPv6 Endpoints

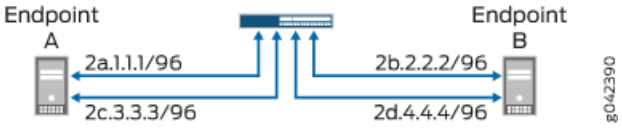

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Features Overview on page 2323
- Understanding Network Address Translation-Protocol Translation on page 2270

Copyright © 2016, Juniper Networks, Inc.

Understanding SCTP Multichunk Inspection

The Stream Control Transmission Protocol (SCTP) firewall checks all chunks in a message and then permits or drops the packet based on the policy. Use the **set security gprs sctp multichunk-inspection enable** command to enable SCTP multichunk inspection to check all chunks in a message. Use the **delete security gprs sctp multichunk-inspection enable** or **set security gprs sctp multichunk-inspection disable** command to disable SCTP multichunk inspection to check only the first chunk.

After enabling SCTP multichunk inspection, the SCTP firewall checks all chunks in a message and permits or drops the packet. The SCTP firewall drops the packet in the following cases:

- The layout of the SCTP chunks do not follow RFC 4960.
- A control chunk cannot pass the inspection of the SCTP finite state machine (FSM) or sanity checks.
- A data chunk is not allowed to pass the SCTP profile because of the SCTP FSM or sanity checks.
- A data chunk is not allowed to pass through the SCTP profile because of protocol blocking or rate limiting. The SCTP firewall resets this chunk to a null protocol data unit (PDU) and continues to check the next chunk. A data chunk is set to a null PDU based on the following rules:
  - When you set the null PDU value to **0xFFFF** using the **set security gprs sctp nullpdu protocol ID-0xFFFF** command, then the payload protocol identifier value is replaced with **0xFFFF** and the user data field is not modified.
  - When you set the null PDU value to **0x0000** using the **set security gprs sctp nullpdu protocol ID-0x0000** command, then the payload protocol identifier value is replaced with **0x0000** and the first four bytes of the user data field is replaced with zeroes.

  If all chunks in a packet are null PDUs, the SCTP firewall drops the packet.

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Packet Structure Overview on page 2327

SCTP Packet Structure Overview

An SCTP packet consists of the following sections:

- Common header section
- Data chunk section

Figure 110 illustrates the structure of the SCTP packet.

Copyright © 2016, Juniper Networks, Inc.

Figure 110: SCTP Packet Structure

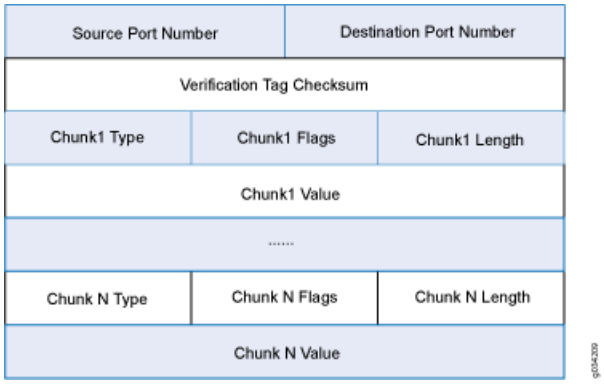

Common header section—All SCTP packets require a common header section. This section occupies the first 12 bytes of the packet. Table 253 describes the fields in the common header section:

Table 253: Common Header Fields

Field	Description
Source port number	Identifies the sending port.
Destination port number	Identifies the receiving port. The hosts use the destination port number to route the packet to the appropriate destination or an application.
Verification tag	Distinguishes stale packets from a previous connection. This is a 32-bit random value created during initialization.
Checksum	Uses the cyclic redundancy check (CRC32) algorithm to detect errors that might have been introduced during data transmission.

Data chunk section—This section occupies the remaining portion of the packet. Table 254 describes the fields in the data chunk section:

Table 254: Data Chunk Fields

Field	Description
Chunk Type	Identifies the contents of the chunk value field. This is 1- byte long.
Chunk Flags	Consists of 8 flag-bits whose definition varies with the chunk type. The default value is zero. This indicates that no application identifier is specified by the upper layer for the data.
Chunk Length	Specifies the total length of the chunk in bytes. This field is 2 - bytes long. If the chunk does not form a multiple of 4 bytes (that is, the length is not a multiple of 4) it is implicitly padded with zeros which are not included in the chunk length.
Chunk Value	A general purpose data field.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The resource manager (RM) allows 8 source IP addresses and 8 destination IP addressses during an SCTP communication.

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Features Overview on page 2323
- Understanding SCTP Behavior in Chassis Cluster on page 2324

SCTP Configuration Overview

You must configure at least one SCTP profile to enable the security device to perform stateful inspection on all SCTP traffic. The stateful inspection of SCTP traffic will drop some anomalous SCTP packets.

The SCTP firewall supports deeper inspection of the profiles:

- Packet filtering—The profile configuration of drop packets for special SCTP payload protocol and M3UA service enables packet filtering.
- Limit-rate—Controls the M3UA and SCCP packets rate per association.

The SCTP deeper inspection requires the following settings:

- Creating a SCTP profile
- Configuring the filtering and limit parameters
- Binding the SCTP profile to a policy

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Packet Structure Overview on page 2327
- SCTP Features Overview on page 2323
- Understanding SCTP Behavior in Chassis Cluster on page 2324
- Example: Configuring a Security Policy to Permit or Deny SCTP Traffic on page 2332

Example: Configuring a GPRS SCTP Profile for Policy-Based Inspection to Reduce Security Risks

In the GPRS architecture, the fundamental cause of security threats to an operator's network is the inherent lack of security in the GPRS tunneling protocol (GTP). This example

Copyright © 2016, Juniper Networks, Inc.

shows how to configure a GPRS SCTP profile for policy-based inspection to reduce the GTP's security risks.

- Requirements on page 2330
- Overview on page 2330
- Configuration on page 2330
- Verification on page 2331

Requirements

Before you begin, understand the GPRS SCTP hierarchy and its options.

Overview

In this example, you configure a GPRS SCTP profile by setting the limit rate parameter and the payload protocol parameter for SCTP inspection. If your policy includes the **nat-only** option, the payload IP addresses are translated, but they are not inspected.

NOTE: The SCTP commands can be applied only to the policy configured with an SCTP profile.

If you remove the SCTP profile from the policy, the packets are forwarded without any inspection, and the IP address list in the packet payload will not be translated, even if the related static NAT is configured.

Configuration

CLI Quick Configuration

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

```
set security gprs sctp profile roam2att limit rate address 10.1.1.0 sccp 100
set security gprs sctp profile roam2att limit rate address 10.1.1.0 ssp 10
set security gprs sctp profile roam2att limit rate address 10.1.1.0 sst 50
set security gprs sctp profile roam2att drop payload-protocol all
set security gprs sctp profile roam2att permit payload-protocol dua
```

Step-by-Step Procedure

To configure a GPRS SCTP profile:

1. Configure the limit rate parameter.

NOTE: The limit rate is per association.

```
[edit security gprs sctp profile roam2att]
user@host# set limit rate address 10.1.1.0 sccp 100
user@host# set limit rate address 10.1.1.0 ssp 10
user@host# set limit rate address 10.1.1.0 sst 50
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure the payload protocol to drop all SCTP payload messages.

   ```
   [edit security gprs sctp profile roam2att]
   user@host# set drop payload-protocol all
   ```

3. Configure the payload protocol to allow certain SCTP payload messages.

   ```
   [edit security gprs sctp profile roam2att]
   user@host# set permit payload-protocol dua
   ```

Results From configuration mode, confirm your configuration by entering the **show security gprs** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security gprs
sctp {
  profile roam2att {
    drop {
      payload-protocol all;
    }
    permit {
      payload-protocol dua;
    }
    limit {
      rate {
        address 10.1.1.0 {
          sccp 100;
          ssp 10;
          sst 50;
        }
      }
    }
  }
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying SCTP Profile Configuration

Purpose Verify the SCTP profile configuration.

Action From configuration mode, enter the **show configuration security gprs sctp profile roam2att** command.

```
user@host> show configuration security gprs sctp profile roam2att
drop {
    payload-protocol all;
}
permit {
    payload-protocol dua;
```

Copyright © 2016, Juniper Networks, Inc.

```
}
limit {
    rate {
        address 10.1.1.0 {
            sccp 100;
            ssp 10;
            sst 50;
        }
    }
}
```

Meaning The output displays information about the SCTP payload messages allowed and SCTP payload messages that are dropped. Verify the following information:

- Dropped SCTP payload messages
- Allowed SCTP payload messages

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Packet Structure Overview on page 2327
- SCTP Features Overview on page 2323
- Understanding SCTP Behavior in Chassis Cluster on page 2324
- Example: Configuring a Security Policy to Permit or Deny SCTP Traffic on page 2332

Example: Configuring a Security Policy to Permit or Deny SCTP Traffic

This example shows how to configure a security policy to permit or deny SCTP traffic.

- Requirements on page 2332
- Overview on page 2333
- Configuration on page 2334
- Verification on page 2336

Requirements

Before you begin:

- Create zones. See "Example: Creating Security Zones" on page 1051.
- Configure an address book and create addresses for use in the policy. See "Example: Configuring Address Books and Address Sets" on page 1076.
- Create an application (or application set) that indicates that the policy applies to traffic of that type. See "Example: Configuring Applications and Application Sets" on page 1172.
- Configure a GPRS SCTP profile. See "Example: Configuring a GPRS SCTP Profile for Policy-Based Inspection to Reduce Security Risks" on page 2329.

Copyright © 2016, Juniper Networks, Inc.

Overview

The SCTP firewall implements a policy mechanism that is administratively used to determine the packets that can be passed or dropped. Policies can be configured for multiple addresses, address groups, or the entire zone.

NOTE: In situations where only a few ports are used for SCTP traffic, the SCTP associations are not evenly distributed to Services Processing Units (SPUs). This occurs in the following cases:

- Uneven hash results on the association ports pairs.
- The number of port pairs is less than, or not much greater than, the number of SPUs.

This configuration example shows how to:

- Deny SCTP traffic from the trust zone to the IP address 10.1.1.0/24 in the untrust zone.
- Permit SCTP traffic from an IP address 10.1.1.1/32 in the trust zone to the untrust zone with the SCTP configuration specified in the roam2att profile.

Figure 111 shows the SCTP firewall implementation.

Copyright © 2016, Juniper Networks, Inc.

Figure 111: SCTP Firewall Implementation

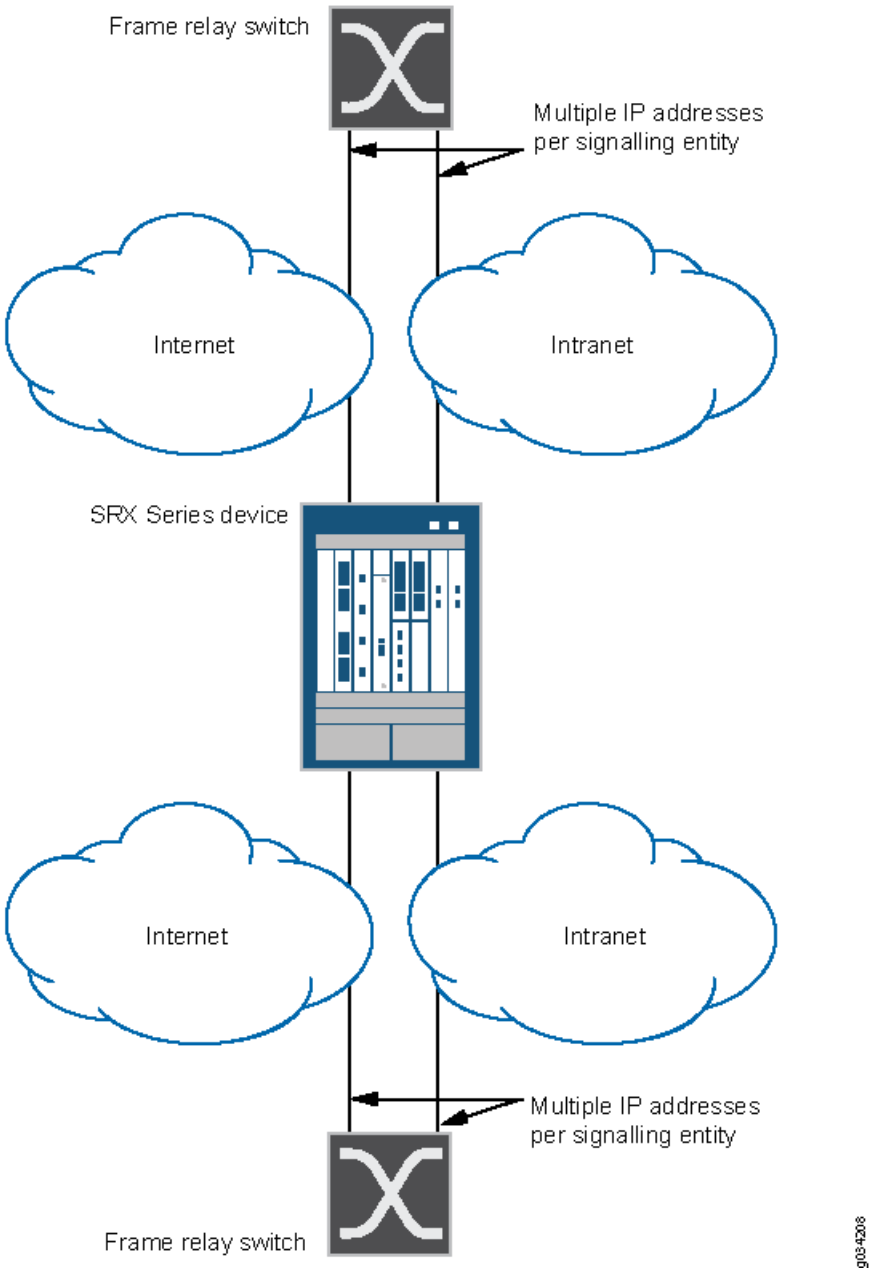

Configuration

CLI Quick Configuration

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

```
set security zones security-zone trust interfaces ge-0/0/2
set security zones security-zone untrust interfaces ge-0/0/1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone untrust policy deny-all match source-address
  any
set security policies policy from-zone trust to-zone untrust policy deny-all match
  destination-address 10.1.1.0/24
set security policies policy from-zone trust to-zone untrust policy deny-all match
  application junos-gprs-sctp
set security policies from-zone trust to-zone untrust policy deny-all then deny
set security policies from-zone trust to-zone untrust policy allow-att-roaming match
  source-address 10.1.2.0/24
set security policies from-zone trust to-zone untrust policy allow-att-roaming match
  destination-address any
set security policies policy from-zone trust to-zone untrust policy allow-att-roaming
  match application junos-gprs-sctp
set security policies from-zone trust to-zone untrust policy allow-att-roaming then permit
  application-services gprs-sctp-profile roam2att
```

Step-by-Step Procedure

To configure a security policy to permit or deny SCTP traffic:

1. Configure the interfaces and security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/0/2
user@host# set security-zone untrust interfaces ge-0/0/1
```

2. Create the security policy to permit traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy allow-att-roaming match source-address 10.1.2.0/24
user@host# set policy allow-att-roaming match destination-address any
user@host# set policy allow-att-roaming match application junos-gprs-sctp
user@host# set policy allow-att-roaming then permit application-services
  gprs-sctp-profile roam2att
```

3. Create the security policy to deny traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy deny-all match source-address any
user@host# set policy deny-all match destination-address 10.1.1.0/24
user@host# set policy deny-all match application junos-gprs-sctp
user@host# set policy deny-all then deny
```

Results

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

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy deny-all {
    match {
      source-address any;
      destination-address 10.1.1.0/24;
      application junos-gprs-sctp;
    }
    then {
      deny;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    policy allow-att-roaming {
      match {
        source-address 10.1.2.0/24;
        destination-address any;
        application junos-gprs-sctp;
      }
      then {
        permit {
          application-services {
            gprs-sctp-profile roam2att;
          }
        }
      }
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

Verifying SCTP Configuration

Purpose Verify the policy inspection configuration.

Action From operational mode, enter **show configuration |display set |match profile**

Related Documentation

- Understanding Stream Control Transmission Protocol on page 2319
- SCTP Configuration Overview on page 2329
- SCTP Packet Structure Overview on page 2327
- SCTP Features Overview on page 2323
- Understanding SCTP Behavior in Chassis Cluster on page 2324

Copyright © 2016, Juniper Networks, Inc.

PART 33

Configuration Statements and Operational Commands

- Configuration Statements on page 2339
- Operational Commands on page 2409

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 106

Configuration Statements

- Security Configuration Statement Hierarchy on page 2340
- [edit security gprs] Hierarchy Level on page 2341
- [edit security forwarding-process] Hierarchy Level on page 2345
- action (APN GTP) on page 2347
- alarm-threshold (Security GPRS) on page 2348
- apn on page 2349
- application-services (Security Forwarding Process) on page 2350
- association-timeout on page 2351
- create-req on page 2351
- delete-req on page 2352
- drop (Security GTP) on page 2353
- drop (Security SCTP) on page 2358
- drop-threshold (Security GPRS) on page 2360
- echo-req on page 2361
- enable (GPRS GTP) on page 2361
- enable-gtpu-distribution on page 2362
- end-user-address-validated (GTP) on page 2362
- forward on page 2363
- gprs on page 2364
- gprs-gtp-profile on page 2367
- gprs-sctp-profile on page 2368
- gtp on page 2369
- gtp-in-gtp-denied on page 2371
- handover-on-roaming-intf on page 2372
- handshake-timeout on page 2372
- imsi-prefix on page 2373
- limit (Security SCTP) on page 2374
- log (Security GTP) on page 2378

Copyright © 2016, Juniper Networks, Inc.

- log (Security SCTP) on page 2379
- max-message-length on page 2380
- message-type on page 2381
- min-message-length on page 2382
- multichunk-inspection on page 2382
- nullpdu on page 2383
- number on page 2383
- other on page 2384
- path-rate-limit on page 2385
- permit (Security SCTP) on page 2386
- profile (Security GTP) on page 2387
- profile (Security SCTP) on page 2390
- rate-limit (Security GTP) on page 2394
- remove-ie on page 2395
- req-timeout on page 2395
- restart-path on page 2396
- reverse on page 2397
- sctp on page 2398
- seq-number-validated (GTP) on page 2403
- timeout (Security GTP) on page 2403
- traceoptions (Security GTP) on page 2404
- traceoptions (Security SCTP) on page 2406
- u-tunnel-validated (GTP) on page 2407
- version (Security GTP) on page 2408

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823

 Copyright © 2016, Juniper Networks, Inc.

- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

[edit security gprs] Hierarchy Level

```
security {
```

```
gprs {
  gtp {
    enable;
    profile profile-name {
      apn pattern-string {
        mcc-mnc mcc-mnc-number {
          action {
            drop;
            pass;
            selection (ms|net|vrf);
          }
        }
      }
      drop {
        aa-create-pdp (0 | 1 | 2 | all);
        aa-delete-pdp (0 | 1 | 2 | all);
        bearer-resource (0 | 1 | 2 | all);
        change-notification (0 | 1 | 2 | all);
        config-transfer (0 | 1 | 2 | all);
        context (0 | 1 | 2 | all);
        create-bearer (0 | 1 | 2 | all);
        create-data-forwarding (0 | 1 | 2 | all);
        create pdp (0 | 1 | 2 | all);
        create-session (0 | 1 | 2 | all);
        create-tnl-forwarding (0 | 1 | 2 | all);
        cs-paging (0 | 1 | 2 | all);
        data-record (0 | 1 | 2 | all);
        delete-bearer (0 | 1 | 2 | all);
        delete-command (0 | 1 | 2 | all);
        delete-data-forwarding (0 | 1 | 2 | all);
        delete-pdn (0 | 1 | 2 | all);
        delete-pdp (0 | 1 | 2 | all);
        delete-session (0 | 1 | 2 | all);
        detach (0 | 1 | 2 | all);
        downlink-notification (0 | 1 | 2 | all);
        echo (0 | 1 | 2 | all);
        error-indication (0 | 1 | 2 | all);
        failure-report (0 | 1 | 2 | all);
        fwd-access (0 | 1 | 2 | all);
        fwd-relocation (0 | 1 | 2 | all);
        fwd-srns-context (0 | 1 | 2 | all);
        g-pdu (0 | 1 | 2 | all);
        identification (0 | 1 | 2 | all);
        mbms-sess-start (0 | 1 | 2 | all);
        mbms-sess-stop (0 | 1 | 2 | all);
        mbms-sess-update (0 | 1 | 2 | all);
        modify-bearer (0 | 1 | 2 | all);
        modify-command (0 | 1 | 2 | all);
        node-alive (0 | 1 | 2 | all);
        note-ms-present (0 | 1 | 2 | all);
        pdu-notification (0 | 1 | 2 | all);
        ran-info (0 | 1 | 2 | all);
        redirection (0 | 1 | 2 | all);
        release-access (0 | 1 | 2 | all);
        relocation-cancel (0 | 1 | 2 | all);
        resume (0 | 1 | 2 | all);
```


Copyright © 2016, Juniper Networks, Inc.

```
    send-route (0 | 1 | 2 | all);
    sgsn-context (0 | 1 | 2 | all);
    stop-paging (0 | 1 | 2 | all);
    supported-extension (0 | 1 | 2 | all);
    suspend (0 | 1 | 2 | all);
    trace-session (0 | 1 | 2 | all);
    update-bearer (0 | 1 | 2 | all);
    update-pdn (0 | 1 | 2 | all);
    update-pdp (0 | 1 | 2 | all);
    ver-not-supported (0 | 1 | 2 | all);
}
gtp-in-gtp-denied;
handover-on-roaming-intf;
log {
    forwarded (basic | detail);
    prohibited (basic | detail);
    rate-limited {
        (basic | detail);
        frequency-number number;
    }
    state-invalid (basic | detail);
}
max-message-length number;
min-message-length number;
path-rate-limit {
    message-type {
        create-req {
            alarm-threshold {
                forward number;
                reverse number;
            }
            drop-threshold {
                forward number;
                reverse number;
            }
        }
        delete-req {
            alarm-threshold {
                forward number;
                reverse number;
            }
            drop-threshold {
                forward number;
                reverse number;
            }
        }
        echo-req {
            alarm-threshold {
                forward number;
                reverse number;
            }
            drop-threshold {
                forward number;
                reverse number;
            }
        }
```

```
            other {
              alarm-threshold {
                forward number;
                reverse number;
              }
              drop-threshold {
                forward number;
                reverse number;
              }
            }
          }
        }
        rate-limit limit;
        remove-ie {
          version v1 {
          number ie-number;
          release (R6 | R7 | R8 | R9);
          }
        }
        req-timeout;
        restart-path (all | create | echo);
        timeout (value);
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
      }
    }
    sctp {
      log {
        association;
        configuration;
        control-message-all;
        control-message-drop;
      }
      multichunk-inspection (enable | disable);
      nullpdu {
        protocol (ID-0x0000 | ID-0xFFFF);
      }
      profile profile-name {
        association-timeout time-in-minutes;
        drop {
          m3ua-service {
            isup;
            sccp;
            tup;
          }
          payload-protocol {
            all;
```


Copyright © 2016, Juniper Networks, Inc.

```
                asap;
                bicc;
                ddp-segment;
                ddp-stream;
                dua;
                enrp;
                h248;
                h323;
                iua;
                m2pa;
                m2ua;
                m3ua;
                qipc;
                reserved;
                simco;
                sua;
                tali;
                v5ua;
              }
            }
            handshake-timeout time-in-seconds;
            limit {
              rate {
                address ip-address {
                  sccp rate-limit;
                  ssp rate-limit;
                  sst rate-limit;
                }
                sccp rate-limit;
                ssp rate-limit;
                sst rate-limit;
              }
            }
          }
          traceoptions {
            file {
              filename;
              files number;
              match regular-expression;
              (no-world-readable | world-readable);
              size maximum-file-size;
            }
            flag flag;
            no-remote-trace;
          }
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security forwarding-process] Hierarchy Level

```
security {
```

```
forwarding-process {
  application-services {
    enable-gtpu-distribution;
    maximize-alg-sessions;
    maximize-idp-sessions {
      inline-tap;
      weight (equal | firewall | idp);
    }
    packet-ordering-mode {
      (hardware | software);
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

action (APN GTP)

Syntax

```
action {
  drop;
  pass;
  selection (ms|net|vrf);
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* apn *pattern-string* mcc-mnc *mcc-mnc-number*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure GTP profile access point (AP) name filtering action to allow or deny access to specific access points.

Options

- **drop**—Specify to deny GTP packets from all selection modes for the specified access points.
- **pass**—Specify to permit GTP packets from all selection modes for the access points.
- **selection**—Specify one of the following selection modes for the access points:
  - **ms**—The access point name is provided by a mobile station, and the user-subscription is not verified.
  - **net**—The access point name is provided by a network, and the user subscription is not verified.
  - **vrf**— The access point name is provided by a network or an MS, and the user-subscription is verified.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

alarm-threshold (Security GPRS)

Syntax

```
alarm-threshold {
    forward number;
    reverse number;
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type (create-req | delete-req | echo-req | other)]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify an alarm threshold for path rate limit.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

apn

Syntax

```
apn pattern-string {
  imsi-prefix imsi-prefix-digits {
    action {
      drop;
      pass;
      selection (ms|net|vrf);
    }
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4. Option mcc-mnc replaced with imsi-prefix in Junos OS Release 12.1X44-D10.

Description Allow or deny access to specific access point names (APNs).

Options

- *pattern-string*—Set an APN suffix, such as “example.net.mcc123.mnc456.gprs”.
- *imsi-prefix-digits*—Specify an IMSI prefix.
- **drop**—Specify to deny GTP packets from all selection modes for the APN.
- **pass**—Specify to permit GTP packets from all selection modes for the APN.
- **selection**—Specify one of the following selection modes for the APN:
  - **ms**—The APN is provided by a Mobile Station (MS), and the user-subscription is not verified.
  - **net**—The APN is provided by a network, and the user-subscription is not verified.
  - **vrf**—The APN is provided by a network or an MS, and the user-subscription is verified.

NOTE: Because APN filtering is based on a perfect match, using the wildcard (*) when setting an APN suffix can prevent the inadvertent exclusion of APNs you would otherwise authorize. The security device automatically denies all other APNs that do not match, if the action is pass. You can only use the wildcard as the first character in the APN string.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

application-services (Security Forwarding Process)

Syntax

```
application-services {
  enable-gtpu-distribution;
  maximize-alg-sessions;
  maximize-idp-sessions {
    inline-tap;
    weight (equal | firewall | idp);
  }
  packet-ordering-mode {
    (hardware | software);
  }
}
```

Hierarchy Level [edit security forwarding-process]

Release Information Statement introduced in Junos OS Release 9.6. Statement updated in Junos OS Release 10.4. Statement updated in Junos OS Release 15.1X49-D40 with the **enable-gtpu-distribution** option.

Description You can configure the device to switch from an integrated firewall mode to maximize intrusion detection and prevention (IDP) mode to increase the capacity of IDP processing with the **maximize-idp-sessions** option. When you maximize IDP, you are decoupling IDP processes from firewall processes, allowing the device to support the same number of firewall and IDP sessions.

You can configure maximum Application Layer Gateway (ALG) sessions by using the **maximize-alg-sessions** option. By default, the session capacity number for Real-Time Streaming Protocol (RTSP), FTP, and Trivial File Transfer Protocol (TFTP) ALG sessions is 10,000 per flow Services Processing Unit (SPU). You must reboot the device (and its peer in chassis cluster mode) for the configuration to take effect. The **maximize-alg-sessions** option now enables you to increase defaults as follows:

- RTSP, FTP, and TFTP ALG session capacity: 25,000 per flow SPU
- TCP proxy connection capacity: 40,000 per flow SPU

NOTE: **Flow session capacity is reduced to half per flow SPU; therefore the aforementioned capacity numbers will not change on central point flow.**

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

association-timeout

Syntax association-timeout *time-in-minutes*;

Hierarchy Level [edit security gprs sctp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2. The association timeout range increased in Junos OS Release 12.1X45-D10.

Description Set the association timeout for Stream Control Transmission Protocol (SCTP).

Options ***time-in-minutes***—Number of minutes of association time that elapse before the session is terminated.

Range: 10 through 6000 (100 hours).

Default: 300 minutes (5 hours).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

create-req

Syntax

```
create-req {
  alarm-threshold {
    forward number;
    reverse number;
  }
  drop-threshold {
    forward number;
    reverse number;
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit the number of packets per second for GTP create request.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

delete-req

Syntax

```
delete-req {
  alarm-threshold {
    forward number;
    reverse number;
  }
  drop-threshold {
    forward number;
    reverse number;
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit the number of packets per second for GTP delete request.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

drop (Security GTP)

Syntax

```
drop {
    aa-create-pdp (0 | 1 | 2 | all);
    aa-delete-pdp (0 | 1 | 2 | all);
    bearer-resource (0 | 1 | 2 | all);
    change-notification (0 | 1 | 2 | all);
    config-transfer (0 | 1 | 2 | all);
    context (0 | 1 | 2 | all);
    create-bearer (0 | 1 | 2 | all);
    create-data-forwarding (0 | 1 | 2 | all);
    create pdp (0 | 1 | 2 | all);
    create-session (0 | 1 | 2 | all);
    create-tnl-forwarding (0 | 1 | 2 | all);
    cs-paging (0 | 1 | 2 | all);
    data-record (0 | 1 | 2 | all);
    delete-bearer (0 | 1 | 2 | all);
    delete-command (0 | 1 | 2 | all);
    delete-data-forwarding (0 | 1 | 2 | all);
    delete-pdn (0 | 1 | 2 | all);
    delete-pdp (0 | 1 | 2 | all);
    delete-session (0 | 1 | 2 | all);
    detach (0 | 1 | 2 | all);
    downlink-notification (0 | 1 | 2 | all);
    echo (0 | 1 | 2 | all);
    error-indication (0 | 1 | 2 | all);
    failure-report (0 | 1 | 2 | all);
    fwd-access (0 | 1 | 2 | all);
    fwd-relocation (0 | 1 | 2 | all);
    fwd-srns-context (0 | 1 | 2 | all);
    g-pdu (0 | 1 | 2 | all);
    identification (0 | 1 | 2 | all);
    mbms-sess-start (0 | 1 | 2 | all);
    mbms-sess-stop (0 | 1 | 2 | all);
    mbms-sess-update (0 | 1 | 2 | all);
    modify-bearer (0 | 1 | 2 | all);
    modify-command (0 | 1 | 2 | all);
    node-alive (0 | 1 | 2 | all);
    note-ms-present (0 | 1 | 2 | all);
    pdu-notification (0 | 1 | 2 | all);
    ran-info (0 | 1 | 2 | all);
    redirection (0 | 1 | 2 | all);
    release-access (0 | 1 | 2 | all);
    relocation-cancel (0 | 1 | 2 | all);
    resume (0 | 1 | 2 | all);
    send-route (0 | 1 | 2 | all);
    sgsn-context (0 | 1 | 2 | all);
    stop-paging (0 | 1 | 2 | all);
    supported-extension (0 | 1 | 2 | all);
    suspend (0 | 1 | 2 | all);
    trace-session (0 | 1 | 2 | all);
    update-bearer (0 | 1 | 2 | all);
    update-pdn (0 | 1 | 2 | all);
    update-pdp (0 | 1 | 2 | all);
```

Copyright © 2016, Juniper Networks, Inc.

```
  ver-not-supported (0 | 1 | 2 | all);
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. New GTP message types added in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4

Description Drop GTP message types. Specify **All** to drop messages for all GTP versions. Specify **0**, **1**, or **2** to drop messages for GTP versions 0, 1, or 2, respectively.

Options The following lists CLI keywords that each represent a GTP message type.

You must specify *(0 | 1 | 2 | all)* to specify the GTP release version number for the specified message type. The possible versions are 0, 1, 2, or all.

NOTE: A GTP message type includes one or many messages. When you drop a message type, you automatically drop all messages of the specified type.

- **aa-create-pdp** —Represents Create AA PDP Context Request and Create AA PDP Context Response messages.
- **aa-delete-pdp** —Represents Delete AA PDP Context Request and Delete AA PDP Context Response messages.
- **bearer-resource**—Represents Bearer Resource Command and Bearer Resource Failure messages.

Copyright © 2016, Juniper Networks, Inc.

- **change-notification**—Represents Change Notification Request and Change Notification Response messages.
- **context**—Represents Context Request and Context Response messages.
- **config-transfer**—Represents Configuration Transfer Tunnel messages.
- **create-bearer**—Represents Create Bearer Request and Create Bearer Response messages.
- **create-data-forwarding**—Represents Create Indirect Data Forwarding Request and Create Indirect Data Forwarding Response messages.
- **create-tnl-forwarding**—Represents Create Forwarding Tunnel Request and Create Forwarding Tunnel Response messages.

Copyright © 2016, Juniper Networks, Inc.

- **create-pdp**—Represents Create PDP Context Request and Create PDP Context Response messages.
- **create-session**—Represents Create Session Request and Create Session Response messages.
- **cs-paging**—Represents CS Paging Indication messages.
- **data-record**—Represents Data Record Request and Data Record Response messages.
- **delete-bearer**—Represents Delete Bearer Request and Delete Bearer Response messages.
- **delete-command**—Represents Delete Bearer Command and Delete Bearer Failure messages.
- **delete-data-forwarding**—Represents Delete Indirect Data Forwarding Request and Delete Indirect Data Forwarding Response messages.
- **delete-pdn**—Represents Delete PDN Connection Set Request and Delete PDN Connection Set Response messages.
- **delete-pdp**—Represents Delete PDP Context Request and Delete PDP Context Response messages.
- **delete-session**—Represents Delete Session Request and Delete Session Response messages.
- **detach**—Represents Detach Notification and Detach Acknowledgement messages.
- **downlink-notification**—Represents Downlink Data Notification, Downlink Data Acknowledgement, and Downlink Data Notification Failure Indication messages.
- **echo**—Represents Echo Request and Echo Response messages.
- **error-indication**—Represents Error Indication messages.
- **failure-report**—Represents Failure Report Request and Failure Report Response messages.
- **fwd-access**—Represents Forward Access Context Notification and Forward Access Context Acknowledgment messages.
- **fwd-relocation**—Represents Forward Relocation Request, Forward Relocation Response, Forward Relocation Complete, and Forward Relocation Complete Acknowledge messages.
- **fwd-srns-context**—Represents Forward SRNS Context Request and Forward SRNS Context Response messages.
- **g-pdu**—Represents G-PDU and T-PDU messages.
- **identification**—Represents Identification Request and Identification Response messages.
- **mbms-sess-start**—Represents MBMS Session Start Request and MBMS Session Start Response messages.
- **mbms-sess-stop**—Represents MBMS Session Stop Request and MBMS Session Stop Response messages.

Copyright © 2016, Juniper Networks, Inc.

- **mbms-sess-update**—Represents MBMS Session Update Request and MBMS Session Update Response messages.
- **modify-bearer**—Represents Modify Bearer Request and Modify Bearer Response messages.
- **modify-command**—Represents Modify Bearer Command and Modify Bearer Failure messages.
- **node-alive**—Represents Node Alive Request and Node Alive Response messages.
- **note-ms-present**—Represents Note MS GPRS Present Request and Note MS GPRS Present Response messages.
- **pdu-notification**—Represents PDU Notification request and PDU Notification response messages.
- **ran-info**—Represents Ran Info Relay messages.
- **redirection**—Represents Redirection Request and Redirection Response messages.
- **relocation-cancel**—Represents Relocations Cancel Request and Relocation Cancel Response messages.
- **resume**—Represents Resume Notification and Resume Acknowledgement messages.
- **send-route**—Represents Send Route Info Request and Send Route Info Response messages.
- **sgsn-context**—Represents SGSN Context Request and SGSN Context Response messages.
- **stop-paging**—Represents Stop Paging Indication messages.
- **supported-extension**—Represents Supported Extension Headers Notification messages.
- **suspend**—Represents Suspend Notification and Suspend Acknowledgement messages.
- **trace-session**—Represents Trace Session Activation and Trace Session Deactivation messages.
- **update-bearer**—Represents Update Bearer Request and Update Bearer Response messages.
- **update-pdn**—Represents Update PDN Set Connection Request and PDN Set Connection Response messages.
- **update-pdp**—Represents Update PDP Request and Update PDP Response messages.
- **ver-not-supported**—Represents Version Not Supported messages.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

drop (Security SCTP)

Syntax

```
drop {
  m3ua-service {
    isup;
    sccp;
    tup;
  }
  payload-protocol {
    id;
    all;
    asap;
    bicc;
    ddp-segment;
    ddp-stream;
    diameter-dtls;
    diameter-sctp;
    dua;
    enrp;
    h248;
    h323;
    iua;
    m2pa;
    m2ua;
    m3ua;
    qipc;
    reserved;
    s1ap;
    simco;
    sua;
    tali;
    v5ua;
    x2ap;
  }
}
```

Hierarchy Level [edit security gprs sctp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2. Support for the **payload-protocol** statement was modified n Junos OS Release 12.1X46-D10.

Description Display information about the configuration of the current Stream Control Transmission Protocol (SCTP) inspection.

Options

- **m3ua-services**—M3UA data service indicator. The following values are supported:
  - **isup**—ISDN Upper Part.
  - **sccp**—Signaling Connection Control Part.
  - **tup**—Telephone User Part.
- **payload-protocol**—SCTP payload protocol identifier. The following values are supported.

Copyright © 2016, Juniper Networks, Inc.

- **id**—Specify payload protocol ID.
- **all**—All SCTP payload protocol identifiers (id:0~63).
- **asap**—Aggregate Server Access Protocol.
- **bicc**—Bearer Independent Call Control.
- **ddp-segement**—Direct Data Placement Segment chunk.
- **ddp-stream**—Direct Data Placement Stream session control.
- **diameter-dtls**—Diameter in a DTLS/SCTP DATA chunk.
- **diameter-sctp**—Diameter in a SCTP DATA chunk.
- **dua**—DPNSS/DASS 2 extensions to IUA Protocol.
- **enrp**—Endpoint Handlespace Redundancy Protocol.
- **h248**—H.248 Protocol.
- **h323**—H.323 Protocol.
- **iua**—ISDN User Adaptation Layer.
- **m2pa**—MTP2 User Peer-to-Peer Adaption Layer.
- **m2ua**—MTP2 User Adaption Layer.
- **m3ua**—MTP3 User Adaption Layer.
- **qipc**—Q.IPC.
- **reserved**—Reserved by SCTP.
- **s1ap**—S1 Application Protocol (S1AP).
- **simco**—Simple Middlebox Configuration.
- **sua**—SCCP User Adaption Layer.
- **tali**—Transport Adapter Layer Interface.
- **v5ua**—v5.2 User Adaption Layer.
- **x2ap**—X2 Application Protocol (X2AP).

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

drop-threshold (Security GPRS)

Syntax

```
drop-threshold {
    forward number;
    reverse number;
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type (create-req | delete-req | echo-req | other)]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify drop threshold for path rate limit.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

echo-req

Syntax

```
echo-req {
  alarm-threshold {
    forward number;
    reverse number;
  }
  drop-threshold {
    forward number;
    reverse number;
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit the number of packets per minute for GTP echo request.

Options *number*—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

enable (GPRS GTP)

Syntax enable;

Hierarchy Level [edit security gprs gtp]

Release Information Statement introduced in Junos OS Release 10.0.

Description Enable GPRS tunneling protocol functionality.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

enable-gtpu-distribution

Syntax enable-gtpu-distribution;

Hierarchy Level [edit security forwarding-process application-services]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Enable GPRS tunneling protocol, user plane(GTP-U) session distribution to distribute GTP-U traffic handled by a gateway GPRS support node (GGSN) and a Serving GPRS Support Node (SGSN) pair on all Services Processing Units (SPUs) by switching using tunnel-based session distribution where the GTP-U traffic of different tunnels is spread across different SPUs.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Central Point Architecture Support for GTP on page 2218
- Security Configuration Statement Hierarchy on page 603

end-user-address-validated (GTP)

Syntax end-user-address-validated;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the validated address of the end user.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

forward

Syntax forward *number*;

Hierarchy Level
[edit security gprs gtp profile *profile-name* path-rate-limit message-type create-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type delete-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type echo-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type other (alarm-threshold | drop-threshold)]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit messages in the forward direction.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

gprs

Syntax

```
gprs {
    gtp {
        enable;
        profile profile-name {
            apn pattern-string {
                imsi-prefix imsi-prefix-digits {
                    action {
                        drop;
                        pass;
                        selection (ms|net|vrf);
                    }
                }
            }
            drop {
                aa-create-pdp (0 | 1 | 2 | all);
                aa-delete-pdp (0 | 1 | 2 | all);
                bearer-resource (0 | 1 | 2 | all);
                change-notification (0 | 1 | 2 | all);
                config-transfer (0 | 1 | 2 | all);
                context (0 | 1 | 2 | all);
                create-bearer (0 | 1 | 2 | all);
                create-data-forwarding (0 | 1 | 2 | all);
                create pdp (0 | 1 | 2 | all);
                create-session (0 | 1 | 2 | all);
                create-tnl-forwarding (0 | 1 | 2 | all);
                cs-paging (0 | 1 | 2 | all);
                data-record (0 | 1 | 2 | all);
                delete-bearer (0 | 1 | 2 | all);
                delete-command (0 | 1 | 2 | all);
                delete-data-forwarding (0 | 1 | 2 | all);
                delete-pdn (0 | 1 | 2 | all);
                delete-pdp (0 | 1 | 2 | all);
                delete-session (0 | 1 | 2 | all);
                detach (0 | 1 | 2 | all);
                downlink-notification (0 | 1 | 2 | all);
                echo (0 | 1 | 2 | all);
                error-indication (0 | 1 | 2 | all);
                failure-report (0 | 1 | 2 | all);
                fwd-access (0 | 1 | 2 | all);
                fwd-relocation (0 | 1 | 2 | all);
                fwd-srns-context (0 | 1 | 2 | all);
                g-pdu (0 | 1 | 2 | all);
                identification (0 | 1 | 2 | all);
                mbms-sess-start (0 | 1 | 2 | all);
                mbms-sess-stop (0 | 1 | 2 | all);
                mbms-sess-update (0 | 1 | 2 | all);
                modify-bearer (0 | 1 | 2 | all);
                modify-command (0 | 1 | 2 | all);
                node-alive (0 | 1 | 2 | all);
                note-ms-present (0 | 1 | 2 | all);
                pdu-notification (0 | 1 | 2 | all);
                ran-info (0 | 1 | 2 | all);
```


Copyright © 2016, Juniper Networks, Inc.

```
            redirection (0 | 1 | 2 | all);
            release-access (0 | 1 | 2 | all);
            relocation-cancel (0 | 1 | 2 | all);
            resume (0 | 1 | 2 | all);
            send-route (0 | 1 | 2 | all);
            sgsn-context (0 | 1 | 2 | all);
            stop-paging (0 | 1 | 2 | all);
            supported-extension (0 | 1 | 2 | all);
            suspend (0 | 1 | 2 | all);
            trace-session (0 | 1 | 2 | all);
            update-bearer (0 | 1 | 2 | all);
            update-pdn (0 | 1 | 2 | all);
            update-pdp (0 | 1 | 2 | all);
            ver-not-supported (0 | 1 | 2 | all);
          }
          gtp-in-gtp-denied;
          handover-on-roaming-intf;
          log {
            forwarded (basic | detail);
            prohibited (basic | detail);
            rate-limited {
              (basic | detail);
              frequency-number number;
            }
            state-invalid (basic | detail);
          }
          max-message-length number;
          min-message-length number;
          rate-limit limit;
          remove-ie {
            version v1 {
            number ie-number;
            release (R6 | R7 | R8 | R9);
            }
          }
          req-timeout;
          restart-path (all | create | echo);
          timeout (value);
        }
        traceoptions {
          file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
        }
      }
      sctp {
        log {
          association;
          configuration;
          control-message-all;
```

```
    control-message-drop;
    data-message-drop;
    rate-limit;
  }
  profile profile-name {
    association-timeout time-in-minutes;
    drop {
      m3ua-service {
        isup;
        sccp;
        tup;
      }
      payload-protocol {
        all;
        asap;
        bicc;
        ddp-segment;
        ddp-stream;
        dua;
        enrp;
        h248;
        h323;
        iua;
        m2pa;
        m2ua;
        m3ua;
        qipc;
        reserved;
        simco;
        sua;
        tali;
        v5ua;
      }
    }
    handshake-timeout time-in-seconds;
    limit {
      rate {
        address ip-address {
          sccp rate-limit;
          ssp rate-limit;
          sst rate-limit;
        }
        sccp rate-limit;
        ssp rate-limit;
        sst rate-limit;
      }
    }
    nat-only;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        flag flag;
        no-remote-trace;
      }
    }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 10.0. Statement modified in Junos OS Release 15.1X49-D40.

Description Configure General Packet Radio Service (GPRS) features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

gprs-gtp-profile

Syntax gprs-gtp-profile *profile-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Specify the name of the GPRS tunneling protocol profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

gprs-sctp-profile

Syntax gprs-sctp-profile *profile-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Specify the name of the GPRS stream control protocol profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

gtp

Syntax

```
gtp {
    enable;
    profile profile-name {
        apn pattern-string {
            imsi-prefix imsi-prefix-digits {
                action {
                    drop;
                    pass;
                    selection (ms|net|vrf);
                }
            }
        }
        drop {
            aa-create-pdp (0 | 1 | 2 | all);
            aa-delete-pdp (0 | 1 | 2 | all);
            bearer-resource (0 | 1 | 2 | all);
            change-notification (0 | 1 | 2 | all);
            config-transfer (0 | 1 | 2 | all);
            context (0 | 1 | 2 | all);
            create-bearer (0 | 1 | 2 | all);
            create-data-forwarding (0 | 1 | 2 | all);
            create pdp (0 | 1 | 2 | all);
            create-session (0 | 1 | 2 | all);
            create-tnl-forwarding (0 | 1 | 2 | all);
            cs-paging (0 | 1 | 2 | all);
            data-record (0 | 1 | 2 | all);
            delete-bearer (0 | 1 | 2 | all);
            delete-command (0 | 1 | 2 | all);
            delete-data-forwarding (0 | 1 | 2 | all);
            delete-pdn (0 | 1 | 2 | all);
            delete-pdp (0 | 1 | 2 | all);
            delete-session (0 | 1 | 2 | all);
            detach (0 | 1 | 2 | all);
            downlink-notification (0 | 1 | 2 | all);
            echo (0 | 1 | 2 | all);
            error-indication (0 | 1 | 2 | all);
            failure-report (0 | 1 | 2 | all);
            fwd-access (0 | 1 | 2 | all);
            fwd-relocation (0 | 1 | 2 | all);
            fwd-srns-context (0 | 1 | 2 | all);
            g-pdu (0 | 1 | 2 | all);
            identification (0 | 1 | 2 | all);
            mbms-sess-start (0 | 1 | 2 | all);
            mbms-sess-stop (0 | 1 | 2 | all);
            mbms-sess-update (0 | 1 | 2 | all);
            modify-bearer (0 | 1 | 2 | all);
            modify-command (0 | 1 | 2 | all);
            node-alive (0 | 1 | 2 | all);
            note-ms-present (0 | 1 | 2 | all);
            pdu-notification (0 | 1 | 2 | all);
            ran-info (0 | 1 | 2 | all);
            redirection (0 | 1 | 2 | all);
```

Copyright © 2016, Juniper Networks, Inc.

```
            release-access (0 | 1 | 2 | all);
            relocation-cancel (0 | 1 | 2 | all);
            resume (0 | 1 | 2 | all);
            send-route (0 | 1 | 2 | all);
            sgsn-context (0 | 1 | 2 | all);
            stop-paging (0 | 1 | 2 | all);
            supported-extension (0 | 1 | 2 | all);
            suspend (0 | 1 | 2 | all);
            trace-session (0 | 1 | 2 | all);
            update-bearer (0 | 1 | 2 | all);
            update-pdn (0 | 1 | 2 | all);
            update-pdp (0 | 1 | 2 | all);
            ver-not-supported (0 | 1 | 2 | all);
          }
          gtp-in-gtp-denied;
          handover-on-roaming-intf;
          log {
            forwarded (basic | detail);
            prohibited (basic | detail);
            rate-limited {
              (basic | detail);
              frequency-number number;
            }
            state-invalid (basic | detail);
          }
          max-message-length number;
          min-message-length number;
          rate-limit limit;
          remove-ie {
            version v1 {
            number ie-number;
            release (R6 | R7 | R8 | R9);
            }
          }
          req-timeout;
          restart-path (all | create | echo);
          timeout (value);
        }
        traceoptions {
          file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
        }
      }
```

Hierarchy Level [edit security gprs]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 10.0. The **restart-path** option added in Junos OS Release 11.4. New GPRS tunneling protocol (GTP) message types added in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4. Statement modified in Junos OS Release 15.1X49-D40.

Description Use the GTP commands to enable the GTP service, configure GTP objects, set traceoptions, remove GTP inspection object configurations, and obtain configuration information.

A Juniper Networks security device provides firewall protection for the following types of General Packet Radio Service (GPRS) interfaces:

- Gn—The Gn interface is the connection between a Serving GPRS Support Node (SGSN) and a gateway GPRS support node (GGSN) within the same public land mobile network (PLMN).
- Gp—The Gp interface is the connection between two PLMNs.
- Gi—The Gi interface is the connection between a GGSN and the Internet or destination networks connected to a PLMN.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

gtp-in-gtp-denied

Syntax gtp-in-gtp-denied;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Select this option to enable the security device to detect and drop a GTP packet that contains another GTP packet in its message body.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

handover-on-roaming-intf

Syntax handover-on-roaming-intf;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Select this option to enable the security device to receive context and forward relocation messages, inspect the packets, and to set up PDP contexts on the device.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security gprs] Hierarchy Level on page 2341

handshake-timeout

Syntax handshake-timeout *time-in-seconds*;

Hierarchy Level [edit security gprs sctp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Set the handshake time for Stream Control Transmission Protocol (SCTP).

Options ***time-in-seconds***—Number of seconds of handshake time that elapse before the session is terminated. **Range:** 10 to 30 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

imsi-prefix

Syntax

```
imsi-prefix imsi-prefix-digits {
  action {
    drop;
    pass;
    selection (ms|net|vrf);
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* apn *pattern-string*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4. Option mcc-mnc replaced with imsi-prefix in Junos OS Release 12.1X44-D10.

Description Specify an International Mobile Station Identity (IMSI) prefix for filtering GTP packets.

You can also filter GTP packets based on the combination of an IMSI prefix and an access point name (APN).

Options *imsi-prefix-digits*—Specify an IMSI prefix.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

limit (Security SCTP)

Syntax

```
limit {
    address ip-address {
        payload-protocol {
            id {
                rate number;
            }
            asap {
                rate number;
            }
            bicc {
                rate number;
            }
            ddp-segment {
                rate number;
            }
            ddp-stream {
                rate number;
            }
            diameter-dtls {
                rate number;
            }
            diameter-sctp {
                rate number;
            }
            dua {
                rate number;
            }
            enrp {
                rate number;
            }
            h248 {
                rate number;
            }
            h323 {
                rate number;
            }
            iua {
                rate number;
            }
            m2pa {
                rate number;
            }
            m2ua {
                rate number;
            }
            m3ua {
                rate number;
            }
            others {
                rate number;
            }
            qipc {
```


Copyright © 2016, Juniper Networks, Inc.

```
        rate number;
      }
      reserved {
        rate number;
      }
      s1ap {
        rate number;
      }
      simco {
        rate number;
      }
      sua {
        rate number;
      }
      tali {
        rate number;
      }
      v5ua {
        rate number;
      }
      x2ap {
        rate number;
      }
    }
  }
  payload-protocol {
    id {
      rate number;
    }
    asap {
      rate number;
    }
    bicc {
      rate number;
    }
    ddp-segment {
      rate number;
    }
    ddp-stream {
      rate number;
    }
    diameter-dtls {
      rate number;
    }
    diameter-sctp {
      rate number;
    }
    dua {
      rate number;
    }
    enrp {
      rate number;
    }
    h248 {
      rate number;
    }
```

```
        h323 {
            rate number;
        }
        iua {
            rate number;
        }
        m2pa {
            rate number;
        }
        m2ua {
            rate number;
        }
        m3ua {
            rate number;
        }
        others {
            rate number;
        }
        qipc {
            rate number;
        }
        reserved {
            rate number;
        }
        s1ap {
            rate number;
        }
        simco {
            rate number;
        }
        sua {
            rate number;
        }
        tali {
            rate number;
        }
        v5ua {
            rate number;
        }
        x2ap {
            rate number;
        }
    }
    rate {
        address ip-address {
            sccp rate-limit;
            ssp rate-limit;
            sst rate-limit;
        }
        sccp rate-limit;
        ssp rate-limit;
        sst rate-limit;
    }
}
```


Copyright © 2016, Juniper Networks, Inc.

Hierarchy Level [edit security gprs sctp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.2. Statement is modified in Junos OS Release 12.1X46-D10. Support for **address** option accepting both IPv4 and IPv6 formats added in Junos OS Release 12.1X47-D10.

Description Set the rate limit per association for local Services Processing Unit (SPU) packets.

Options

address ***ip-address***—Set Signalling Connection Control Part (SCCP), Subsystem-Prohibited (SSP), and Subsystem Status Test (SST) messages rate limit to an IP address. The IP address can accept either an IPv4 address or an IPv6 address.

sccp ***rate-limit***—Set the SCCP messages rate limit.

ssp ***rate-limit***—Set the SSP messages rate limit.

sst ***rate-limit***—Set the SSP messages rate limit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

log (Security GTP)

Syntax

```
log {
  forwarded (basic | detail);
  prohibited (basic | detail);
  rate-limited {
    (basic | detail);
    frequency-number number;
  }
  state-invalid (basic | detail);
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Configure GTP logs to be viewed from the console or syslog.

NOTE: By default, all logs are disabled on the device.

Options

- **forwarded**—A packet that the security device transmitted because it was valid.
- **prohibited**—A packet that the security device dropped because it was invalid.
- **rate-limited**—A packet that the security device dropped because it exceeded the maximum rate limit of the destination GSN.
  - **frequency-number *number***—Logging frequency over threshold set by rate-limit (2–500).
- **state-invalid**—A packet that the security device dropped because it failed stateful inspection.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

log (Security SCTP)

Syntax

```
log {
    association;
    configuration;
    control-message-all;
    control-message-drop;
    data--message-drop;
    rate-limit;
}
```

Hierarchy Level [edit security gprs sctp]

Release Information Statement introduced in Junos OS Release 10.2. The options **association**, **control-message-all**, **control-message-drop**, and **data-message-drop** added in Junos OS Release 12.1X45-D10.

Description Configure Stream Control Transmission Protocol (SCTP) logs to be viewed from the console or system log.

Options

association—To log association events.

configuration—To log the CLI configuration.

control-message-all—To log both dropped and passed control messages.

control-message-drop—To log the dropped control messages.

data-message-drop—To log the dropped data messages.

rate-limit—To log the rate limit.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

max-message-length

Syntax max-message-length *number*;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Set the maximum message payload length (in bytes) the security device accepts for a GTP message. The default maximum message length is 65,535 bytes. The message length range is from 1 through 65,535 bytes.

Options *number*—Set the maximum message payload length in bytes.

Range: 1 through 65,535 bytes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

message-type

Syntax

```
message-type {
  create-req {
    alarm-threshold {
      forward number;
      reverse number;
    }
    drop-threshold {
      forward number;
      reverse number;
    }
  }
  delete-req {
    alarm-threshold {
      forward number;
      reverse number;
    }
    drop-threshold {
      forward number;
      reverse number;
    }
  }
  echo-req {
    alarm-threshold {
      forward number;
      reverse number;
    }
    drop-threshold {
      forward number;
      reverse number;
    }
  }
  other {
    alarm-threshold {
      forward number;
      reverse number;
    }
    drop-threshold {
      forward number;
      reverse number;
    }
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the group of control messages.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

min-message-length

Syntax min-message-length *number*;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Set the minimum message payload length (in bytes) the security device accepts for a GTP message. The default minimum message length is 0 bytes. The message length range is from 0 through 65,535 bytes.

Options *number*—Set the minimum message payload length in bytes.
Range: 0 through 65,535 bytes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

multichunk-inspection

Syntax multichunk-inspection (enable | disable);

Hierarchy Level [edit security gprs sctp]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Configure the Stream Control Transmission Protocol (SCTP) firewall to enable or disable multichunk inspection.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

nullpdu

Syntax

```
nullpdu {
  protocol (ID-0x0000 | ID-0xFFFF);
}
```

Hierarchy Level [edit security gprs sctp]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Configure the Stream Control Transmission Protocol (SCTP) null protocol data unit (PDU) value.

Options **protocol**—Specify the SCTP null PDU payload protocol identifier.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

number

Syntax number *ie-number*

Hierarchy Level [edit security gprs gtp profile *profile-name* remove-ie version *v1*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the user-configured IE number. IE removal by IE number supports all IEs, ranging from 1 to 255.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

other

Syntax

```
other {
  alarm-threshold {
    forward number;
    reverse number;
  }
  drop-threshold {
    forward number;
    reverse number;
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name* path-rate-limit message-type]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit the number of packets per second for all the other GTPv0/GTPv1-C/GTPv2-C messages.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

path-rate-limit

Syntax

```
path-rate-limit {
  message-type {
    create-req {
      alarm-threshold {
        forward number;
        reverse number;
      }
      drop-threshold {
        forward number;
        reverse number;
      }
    }
    delete-req {
      alarm-threshold {
        forward number;
        reverse number;
      }
      drop-threshold {
        forward number;
        reverse number;
      }
    }
    echo-req {
      alarm-threshold {
        forward number;
        reverse number;
      }
      drop-threshold {
        forward number;
        reverse number;
      }
    }
    other {
      alarm-threshold {
        forward number;
        reverse number;
      }
      drop-threshold {
        forward number;
        reverse number;
      }
    }
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit control messages based on an IP pair.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

permit (Security SCTP)

Syntax

```
permit {
  payload-protocol {
    id;
    all;
    asap;
    bicc;
    ddp-segment;
    ddp-stream;
    diameter-dtls;
    diameter-sctp;
    dua;
    enrp;
    h248;
    h323;
    iua;
    m2pa;
    m2ua;
    m3ua;
    qipc;
    reserved;
    s1ap;
    simco;
    sua;
    tali;
    v5ua;
    x2ap;
  }
}
```

Hierarchy Level
[edit security gprs sctp profile *profile-name*]

Release Information
Statement introduced in Junos OS Release 12.1X46-D10.

Description
Display information about the configuration of the current Stream Control Transmission Protocol (SCTP) inspection.

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

profile (Security GTP)

Syntax

```
profile profile name {
  apn pattern-string {
    imsi-prefix imsi-prefix-digits {
      action {
        drop;
        pass;
        selection (ms|net|vrf);
      }
    }
  }
  drop {
    aa-create-pdp (0 | 1 | 2 | all);
    aa-delete-pdp (0 | 1 | 2 | all);
    bearer-resource (0 | 1 | 2 | all);
    change-notification (0 | 1 | 2 | all);
    config-transfer (0 | 1 | 2 | all);
    context (0 | 1 | 2 | all);
    create-bearer (0 | 1 | 2 | all);
    create-data-forwarding (0 | 1 | 2 | all);
    create pdp (0 | 1 | 2 | all);
    create-session (0 | 1 | 2 | all);
    create-tnl-forwarding (0 | 1 | 2 | all);
    cs-paging (0 | 1 | 2 | all);
    data-record (0 | 1 | 2 | all);
    delete-bearer (0 | 1 | 2 | all);
    delete-command (0 | 1 | 2 | all);
    delete-data-forwarding (0 | 1 | 2 | all);
    delete-pdn (0 | 1 | 2 | all);
    delete-pdp (0 | 1 | 2 | all);
    delete-session (0 | 1 | 2 | all);
    detach (0 | 1 | 2 | all);
    downlink-notification (0 | 1 | 2 | all);
    echo (0 | 1 | 2 | all);
    error-indication (0 | 1 | 2 | all);
    failure-report (0 | 1 | 2 | all);
    fwd-access (0 | 1 | 2 | all);
    fwd-relocation (0 | 1 | 2 | all);
    fwd-srns-context (0 | 1 | 2 | all);
    g-pdu (0 | 1 | 2 | all);
    identification (0 | 1 | 2 | all);
    mbms-sess-start (0 | 1 | 2 | all);
    mbms-sess-stop (0 | 1 | 2 | all);
    mbms-sess-update (0 | 1 | 2 | all);
    modify-bearer (0 | 1 | 2 | all);
    modify-command (0 | 1 | 2 | all);
    node-alive (0 | 1 | 2 | all);
    note-ms-present (0 | 1 | 2 | all);
    pdu-notification (0 | 1 | 2 | all);
    ran-info (0 | 1 | 2 | all);
    redirection (0 | 1 | 2 | all);
    release-access (0 | 1 | 2 | all);
    relocation-cancel (0 | 1 | 2 | all);
```

Copyright © 2016, Juniper Networks, Inc.

```
    resume (0 | 1 | 2 | all);
    send-route (0 | 1 | 2 | all);
    sgsn-context (0 | 1 | 2 | all);
    stop-paging (0 | 1 | 2 | all);
    supported-extension (0 | 1 | 2 | all);
    suspend (0 | 1 | 2 | all);
    trace-session (0 | 1 | 2 | all);
    update-bearer (0 | 1 | 2 | all);
    update-pdn (0 | 1 | 2 | all);
    update-pdp (0 | 1 | 2 | all);
    ver-not-supported (0 | 1 | 2 | all);
  }
  gtp-in-gtp-denied;
  handover-on-roaming-intf;
  log {
    forwarded (basic | detail);
    prohibited (basic | detail);
    rate-limited {
      (basic | detail);
      frequency-number number;
    }
    state-invalid (basic | detail);
  }
  max-message-length number;
  min-message-length number;
  rate-limit limit;
  remove-ie {
    version v1 {
      number ie-number;
      release (R6 | R7 | R8 | R9);
    }
  }
  req-timeout;
  restart-path (all | create | echo);
  timeout (value);
  }
}
```

Hierarchy Level [edit security gprs gtp]

Release Information Statement introduced in Junos OS Release 10.0. The **restart-path** option added in Junos OS Release 11.4. New GPRS tunneling protocol (GTP) message types added in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4. Statement modified in Junos OS Release 15.1X49-D40.

Description Create a profile for the GTP feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

profile (Security SCTP)

Syntax

```
profile profile-name {
  association-timeout time-in-minutes;
  drop {
    m3ua-service {
      isup;
      sccp;
      tup;
    }
    payload-protocol {
      id;
      all;
      asap;
      bicc;
      ddp-segment;
      ddp-stream;
      diameter-dtls;
      diameter-sctp;
      dua;
      enrp;
      h248;
      h323;
      iua;
      m2pa;
      m2ua;
      m3ua;
      qipc;
      reserved;
      s1ap;
      simco;
      sua;
      tali;
      v5ua;
      x2ap;
    }
  }
  handshake-timeout time-in-seconds;
  limit {
    address ip-address {
      payload-protocol {
        id {
          rate number;
        }
        asap {
          rate number;
        }
        bicc {
          rate number;
        }
        ddp-segment {
          rate number;
        }
        ddp-stream {
```


Copyright © 2016, Juniper Networks, Inc.

```
    rate number;
}
diameter-dtls {
    rate number;
}
diameter-sctp {
    rate number;
}
dua {
    rate number;
}
enrp {
    rate number;
}
h248 {
    rate number;
}
h323 {
    rate number;
}
iua {
    rate number;
}
m2pa {
    rate number;
}
m2ua {
    rate number;
}
m3ua {
    rate number;
}
others {
    rate number;
}
qipc {
    rate number;
}
reserved {
    rate number;
}
s1ap {
    rate number;
}
simco {
    rate number;
}
sua {
    rate number;
}
tali {
    rate number;
}
v5ua {
    rate number;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
        x2ap {
            rate number;
        }
    }
}
payload-protocol {
    id {
        rate number;
    }
    asap {
        rate number;
    }
    bicc {
        rate number;
    }
    ddp-segment {
        rate number;
    }
    ddp-stream {
        rate number;
    }
    diameter-dtls {
        rate number;
    }
    diameter-sctp {
        rate number;
    }
    dua {
        rate number;
    }
    enrp {
        rate number;
    }
    h248 {
        rate number;
    }
    h323 {
        rate number;
    }
    iua {
        rate number;
    }
    m2pa {
        rate number;
    }
    m2ua {
        rate number;
    }
    m3ua {
        rate number;
    }
    others {
        rate number;
    }
    qipc {
        rate number;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        reserved {
            rate number;
        }
        s1ap {
            rate number;
        }
        simco {
            rate number;
        }
        sua {
            rate number;
        }
        tali {
            rate number;
        }
        v5ua {
            rate number;
        }
        x2ap {
            rate number;
        }
    }
    rate {
        address ip-address {
            sccp rate-limit;
            ssp rate-limit;
            sst rate-limit;
        }
        sccp rate-limit;
        ssp rate-limit;
        sst rate-limit;
    }
}
nat-only;
permit {
    payload-protocol {
        id;
        all;
        asap;
        bicc;
        ddp-segment;
        ddp-stream;
        diameter-dtls;
        diameter-sctp;
        dua;
        enrp;
        h248;
        h323;
        iua;
        m2pa;
        m2ua;
        m3ua;
        qipc;
        reserved;
        s1ap;
```

```
            simco;
            sua;
            tali;
            v5ua;
            x2ap;
        }
    }
}
```

Hierarchy Level [edit security gprs sctp]

Release Information Statement introduced in Junos OS Release 10.2. Support for the **nat-only** option added in Junos OS Release 12.1X45-D10. Support for the **permit** option is added in Junos OS Release 12.1X46-D10.

Description Create a profile of the Stream Control Transmission Protocol (SCTP) feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

rate-limit (Security GTP)

Syntax rate-limit *value*;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Set the limit rate of control traffic to any GSN defined in a GTP profile.

Options **Range:** 1 through 80,000 messages per second.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

remove-ie

Syntax

```
remove-ie {
  version v1 {
  number ie-number;
  release (R6 | R7 | R8 | R9);
  }
}
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable the security device to detect and remove 3G-specific attributes from the GTP packet header when the packet passes into a 2G network. This allows you to retain interoperability when roaming between 2G and 3G networks.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

req-timeout

Syntax

```
req-timeout;
```

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D35.

Description Specify a GTP request message timeout. The default timeout value is 5 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

restart-path

Syntax restart-path (all | create | echo);

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4. Support for GTPv2 added in Junos OS Release 11.4.

Description Restart a GTP path. Restarting a GTP path deletes all GTP tunnels between two devices.

Options

- **all**—Restart GTP paths by detecting the changed restart number obtained from the Recovery information element (IE) in all GTP messages.
- **create**—Restart GTP paths by detecting the changed restart number obtained from the Recovery IE in create-session messages.
- **echo**—Restart GTP paths by detecting the changed restart number obtained from the Recovery IE in echo messages.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

reverse

Syntax reverse *number*;

Hierarchy Level
[edit security gprs gtp profile *profile-name* path-rate-limit message-type create-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type delete-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type echo-req (alarm-threshold | drop-threshold)]
[edit security gprs gtp profile *profile-name* path-rate-limit message-type other (alarm-threshold | drop-threshold)]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Limit messages in the reverse direction.

Options ***number***—Limit messages in forward or reverse direction.

Range: 1 through 10,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

sctp

Syntax

```
sctp {
  log {
    association;
    configuration;
    control-message-all;
    control-message-drop;
    data-message-drop;
    rate-limit;
  }
  multichunk-inspection (enable | disable);
  nullpdu {
    protocol (ID-0x0000 | ID-0xFFFF);
  }
  profile profile-name {
    association-timeout time-in-minutes;
    drop {
      m3ua-service {
        isup;
        sccp;
        tup;
      }
      payload-protocol {
        id;
        all;
        asap;
        bicc;
        ddp-segment;
        ddp-stream;
        diameter-dtls;
        diameter-sctp;
        dua;
        enrp;
        h248;
        h323;
        iua;
        m2pa;
        m2ua;
        m3ua;
        qipc;
        reserved;
        s1ap;
        simco;
        sua;
        tali;
        v5ua;
        x2ap;
      }
    }
    handshake-timeout time-in-seconds;
    limit {
      address ip-address {
        payload-protocol {
```


Copyright © 2016, Juniper Networks, Inc.

```
id {
    rate number;
}
asap {
    rate number;
}
bicc {
    rate number;
}
ddp-segment {
    rate number;
}
ddp-stream {
    rate number;
}
diameter-dtls {
    rate number;
}
diameter-sctp {
    rate number;
}
dua {
    rate number;
}
enrp {
    rate number;
}
h248 {
    rate number;
}
h323 {
    rate number;
}
iua {
    rate number;
}
m2pa {
    rate number;
}
m2ua {
    rate number;
}
m3ua {
    rate number;
}
others {
    rate number;
}
qipc {
    rate number;
}
reserved {
    rate number;
}
s1ap {
    rate number;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        simco {
            rate number;
        }
        sua {
            rate number;
        }
        tali {
            rate number;
        }
        v5ua {
            rate number;
        }
        x2ap {
            rate number;
        }
    }
}
payload-protocol {
    id {
        rate number;
    }
    asap {
        rate number;
    }
    bicc {
        rate number;
    }
    ddp-segment {
        rate number;
    }
    ddp-stream {
        rate number;
    }
    diameter-dtls {
        rate number;
    }
    diameter-sctp {
        rate number;
    }
    dua {
        rate number;
    }
    enrp {
        rate number;
    }
    h248 {
        rate number;
    }
    h323 {
        rate number;
    }
    iua {
        rate number;
    }
    m2pa {
```


Copyright © 2016, Juniper Networks, Inc.

```
            rate number;
        }
        m2ua {
            rate number;
        }
        m3ua {
            rate number;
        }
        others {
            rate number;
        }
        qipc {
            rate number;
        }
        reserved {
            rate number;
        }
        s1ap {
            rate number;
        }
        simco {
            rate number;
        }
        sua {
            rate number;
        }
        tali {
            rate number;
        }
        v5ua {
            rate number;
        }
        x2ap {
            rate number;
        }
    }
    rate {
        address ip-address {
            sccp rate-limit;
            ssp rate-limit;
            sst rate-limit;
        }
        sccp rate-limit;
        ssp rate-limit;
        sst rate-limit;
    }
}
nat-only;
permit {
    payload-protocol {
        id;
        all;
        asap;
        bicc;
        ddp-segment;
        ddp-stream;
```

```
            diameter-dtls;
            diameter-sctp;
            dua;
            enrp;
            h248;
            h323;
            iua;
            m2pa;
            m2ua;
            m3ua;
            qipc;
            reserved;
            s1ap;
            simco;
            sua;
            tali;
            v5ua;
            x2ap;
          }
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
```

Hierarchy Level [edit security gprs]

Release Information Statement introduced in Junos OS Release 10.2. Support for the **nat-only** option added in Junos OS Release 12.1X45-D10. Support for the **profile** statement added in Junos OS Release 12.1X46-D10.

Description Use the Stream Control Transmission Protocol (SCTP) commands to configure SCTP objects, configure SCTP logs, set trace options, and set address rate limit.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

seq-number-validated (GTP)

Syntax seq-number-validated;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the validated sequence number.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- Example: Enabling GTP Sequence Number Validation on page 2244

timeout (Security GTP)

Syntax timeout *value*;

Hierarchy Level [edit security gprs gtp profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Set the tunnel timeout value in hours. The default is 36 hours.

If a device detects no activity in a tunnel for a specified period, it removes the tunnel from the state table.

Options **Range:** 1 through 1,000 hours.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security GTP)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security gprs gtp]

Release Information Statement introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Enable the device to identify and log the contents of GTP-U or GTP-C messages based on IMSI prefixes or Mobile Station-Integrated Services Data Network (MS-ISDN) identification.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Copyright © 2016, Juniper Networks, Inc.

Syntax: ***x*** **K** to specify KB, ***x*** **m** to specify MB, or ***x*** **g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- **world-readable** | **no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace everything.
  - **chassis-cluster**—Trace chassis cluster events.
  - **configuration**—Trace configuration events.
  - **flow**—Trace flow events.
  - **parser**—Trace parser events.
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security SCTP)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security gprs sctp]

Release Information Statement introduced in Junos OS Release 10.2. The flag statement **detail** introduced in Junos OS Release 12.1X45-D10.

Description Set the trace options for Stream Control Transmission Protocol (SCTP).

Options **file**—Configure the trace file options.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

files ***number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files.

Default: 10 files.

match ***regular-expression***—Refine the output to include lines that contain the regular expression.

size ***maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: ***x*** **K** to specify KB, ***x*** **m** to specify MB, or ***x*** **g** to specify GB

Range: 10 KB through 1 GB.

Default: 128 KB.

Copyright © 2016, Juniper Networks, Inc.

world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- all—Trace everything.
- chassis-cluster—Trace chassis cluster events.
- configuration—Trace configuration events.
- detail—Trace information used for debugging.
- flow—Trace flow events.
- parser—Trace parser events.

no-remote-trace—Set remote tracing as disabled.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

u-tunnel-validated (GTP)

Syntax
```
u-tunnel-validated;
```

Hierarchy Level
[edit security gprs gtp profile *profile-name*]

Release Information
Statement introduced in Junos OS Release 12.1X45-D10.

Description
Specify the validated GTP-U tunnel.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

version (Security GTP)

Syntax

```
version v1 {
  number ie-number;
  release (R6 | R7 | R8 | R9);
}
```

Hierarchy Level [edit security gprs gtp profile profile-name remove-ie]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify GTP version.

Options

- **v1**—GTP version 1.
- **release** —Specify release number. Available options are:
  - Release R6—Specify R6 IE removal.
  - Release R7—Specify R7 IE removal.
  - Release R8—Specify R8 IE removal.
  - Release R9—Specify R9 IE removal.
- **number** —Specify the user-configured IE number. IE removal by IE number supports all IEs, ranging from 1 to 255.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 107

Operational Commands

- clear gtp tunnels
- clear security gprs gtp counters
- clear security gprs sctp association
- clear security gprs sctp counters
- show gtp tunnels
- show security gprs gtp counters
- show security gprs gtp counters path-rate-limit
- show security gprs gtp gsn statistics
- show security gprs sctp association
- show security gprs sctp counters

Copyright © 2016, Juniper Networks, Inc.

clear gtp tunnels

Syntax **clear security gprs gtp tunnel <all | *identifier*>**

Release Information Command introduced in Junos OS Release 10.0. Support for GTPv2 added in Junos OS Release 11.4.

Description Clear all or specified GTP tunnels on the device.

Options

- *identifier*—Clear a single tunnel by entering the tunnel ID. To view current tunnel IDs, type **show security gprs gtp tunnels**.
- all—Clear all existing tunnels.

Required Privilege Level clear

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

clear security gprs gtp counters

Syntax clear security gprs gtp counters *<all | error | ha | message <message-name> | packet | request | tunnel | path-rate-limit>*

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Clear all GTP counters on the device.

Options

- **all**—Clear all GTP counters.
- **error**—Clear GTP error counters.
- **ha**—Clear GTP HA counters.
- **message *message-name***—Clear GTP message counters.
- **packet**—Clear GTP packet counters.
- **request**—Clear GTP request counters.
- **tunnel**—Clear GTP tunnel counters.
- **path-rate-limit**—Clear path-rate-limit counters.

Required Privilege Level clear

Related Documentation

- show security gprs gtp counters on page 2418

List of Sample Output

clear security gprs gtp counters all on page 2411
clear security gprs gtp counters error on page 2411
clear security gprs gtp counters ha on page 2412
clear security gprs gtp counters message v0-create-aa-pdp-req on page 2412
clear security gprs gtp counters packet on page 2412
clear security gprs gtp counters request on page 2412
clear security gprs gtp counters tunnel on page 2412
clear security gprs gtp counters path-rate-limit on page 2412

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security gprs gtp counters all

```
user@host> clear security gprs gtp counters all
All GTP counters have been cleared
```

clear security gprs gtp counters error

```
user@host> clear security gprs gtp counters error
GTP error counter has been cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security gprs gtp counters ha

```
user@host> clear security gprs gtp counters ha
GTP HA counter has been cleared
```

clear security gprs gtp counters message v0-create-aa-pdp-req

```
user@host> clear security gprs gtp counters message v0-create-aa-pdp-req
GTPv0 create AA PDP request message counter has been cleared
```

clear security gprs gtp counters packet

```
user@host> clear security gprs gtp counters packet
GTP packet counter has been cleared
```

clear security gprs gtp counters request

```
user@host> clear security gprs gtp counters request
GTP request counter has been cleared
```

clear security gprs gtp counters tunnel

```
user@host> clear security gprs gtp counters tunnel
GTP tunnel counter has been cleared
```

clear security gprs gtp counters path-rate-limit

```
user@host> clear security gprs gtp counters path-rate-limit
GTP path-rate-limit counter has been cleared
```


Copyright © 2016, Juniper Networks, Inc.

clear security gprs sctp association

Syntax clear security gprs sctp association
<all>
<destination-ip>
<desitnation-port>
<guid>
<init>
<source-ip>
<source-port>

Release Information Command introduced in Junos OS Release 12.1X45-D10.

Description Clear the Stream Control Transmission Protocol (SCTP) association.

Options **none**—Clear the live SCTP associations.

all—Clear all the SCTP associations, both initiated and live. All SCTP traffic is blocked while the associations are being cleared, which can take up to one minute.

destination-ip—Clear the destination IP SCTP association.

destination-port—Clear the destination port SCTP association.

guid—Clear the globally unique identifier SCTP association.

init—Clear the initiated SCTP associations.

source-ip—Clear the source IP address SCTP association.

source-port—Clear the source port SCTP association.

Required Privilege Level clear

Related Documentation
- show security gprs sctp association on page 2427

List of Sample Output clear security gprs sctp association on page 2413

Sample Output

clear security gprs sctp association

```
user@host> clear security gprs sctp association


Clear Association Information for FPC: 2    PIC: 0
Cleared matched SCTP association information:
Has cleared matched association: 0


Clear Association Information for FPC: 2    PIC: 1
Cleared matched SCTP association information:
Has cleared matched association: 9
```

Copyright © 2016, Juniper Networks, Inc.

```
Clear Association Information for FPC: 2    PIC: 2
Cleared matched SCTP association information:
Has cleared matched association: 8


Clear Association Information for FPC: 2    PIC: 3
Cleared matched SCTP association information:
Has cleared matched association: 10


Clear Association Information for FPC: 5    PIC: 0
Cleared matched SCTP association information:
Has cleared matched association: 7


Clear Association Information for FPC: 5    PIC: 1
Cleared matched SCTP association information:
Has cleared matched association: 6
```


Copyright © 2016, Juniper Networks, Inc.

clear security gprs sctp counters

Syntax clear security gprs sctp counters

Release Information Command introduced in Junos OS Release 10.2.

Description Clear the statistics of the dropped Stream Control Transmission Protocol (SCTP) counters.

Options none—Clear all dropped SCTP counters.

Required Privilege Level clear

Related Documentation
- show security gprs sctp counters on page 2429

List of Sample Output clear security gprs sctp counters on page 2415

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security gprs sctp counters

```
user@host> clear security gprs sctp counters
```

show gtp tunnels

Syntax **show security gprs gtp tunnels (brief | summary | detail)**

Release Information Command introduced in Junos OS Release 10.0. Support for GPRS tunneling protocol version 2 (GTPv2) added in Junos OS Release 11.4. Command output updated in Junos OS Release 15.1X49-D40.

Description Display all existing GTP tunnels.

Options
- brief—Display a short listing of all GTP tunnels.
- summary—Display a summary of all GTP tunnels.
- detail–Display detailed information about all the GTP tunnels.

Required Privilege Level view

Related Documentation
- Security Configuration Statement Hierarchy on page 603

List of Sample Output
show security gprs gtp tunnels on page 2416
show security gprs gtp tunnels summary on page 2416
show security gprs gtp tunnels detail on page 2417

Sample Output

show security gprs gtp tunnels

```
user@host> show security gprs gtp tunnels

FPC 7 PIC 0:

 Index: 72000002, EBI/LBI: 5/5(V2)to sgw, Timeout: 1440m
 User: 61.0.0.102, 12345678 --> 62.0.0.102, 00000021
 Control: 61.0.0.101, 00325ac1 --> 62.0.0.101, 00000001

FPC 8 PIC 0:

 Index: 0x02000040 Tunnel ID: 0x50502410121507f5(V0), Timeout: 59m
 User: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001
 Ctrl: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001

 3 tunnels active in total
```

show security gprs gtp tunnels summary

```
user@host> show security gprs gtp tunnels summary

FPC 1 PIC 0:

FPC 1 PIC 1:

FPC 2 PIC 0:

FPC 2 PIC 1:
```


Copyright © 2016, Juniper Networks, Inc.

```
2 tunnels active in total
```

show security gprs gtp tunnels detail

```
user@host> show security gprs gtp tunnels detail
node0:
--------------------------------------------------------------------------
 FPC 0 PIC 0:

 FPC 0 PIC 1:

 FPC 0 PIC 2:

 FPC 0 PIC 3:

Index: 0x02000040 Tunnel ID: 0x50502410121507f5(V0), Timeout: 59m
User: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001
Ctrl: 20.1.0.1, 00000001 -> 20.0.2.1, 00000001

 1 tunnels active in total
```

show security gprs gtp counters

Syntax show security gprs gtp counters <*all* | *error* | *ha* | *message* <*message-name*> | *packet* | *request* | *tunnel* | *path-rate-limit*>

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display counters that can be used to indicate the number of GTP tunnel counters (Allocated and Freed), GTP packet counters (Received, Passed, and Dropped), Brief message counters (Receive, Forward, and Drop), Error counters, Request counters, HA counters, and Path-rate-limit counters (Drop and Alarm).

Options

- **all**—Show all GTP counters.
- **error**—Show GTP error counters.
- **ha**—Show GTP HA counters.
- **message** ***message-name***—Show GTP message counters.
- **packet**—Show GTP packet counters.
- **request**—Show GTP request counters.
- **tunnel**—Show GTP tunnel counters.
- **path-rate-limit**—Show path-rate-limit counters.

Required Privilege Level view

Related Documentation

- clear security gprs gtp counters on page 2411

List of Sample Output

show security gprs gtp counters all on page 2420
show security gprs gtp counters error on page 2422
show security gprs gtp counters ha on page 2423
show security gprs gtp counters message v0-create-aa-pdp-req on page 2423
show security gprs gtp counters packet on page 2423
show security gprs gtp counters request on page 2424
show security gprs gtp counters tunnel on page 2424
show security gprs gtp counters path-rate-limit on page 2424

Output Fields Table 255 lists the output fields for the **show security gprs gtp counters** command.

Copyright © 2016, Juniper Networks, Inc.

Table 255: show security gprs gtp counters all Output Fields

Field Name	Field Description
Tunnel counters	Tunnel counters are used to track the number of tunnels that are created on the device. There are two entries: • Allocated • Freed Active tunnel number = Number of Allocated counters - Number of Freed counters
Packet counters	Packet counters indicate the number of GTP packets that are received and processed on the device. There are three entries: • Received – Number of GTP packet messages received. • Passed – Number of GTP packet messages passed. • Dropped – Number of GTP packet messages dropped because of an error. Number of Received counters = Number of Dropped counters + Number of Passed counters
Brief message counters	GTP messages counters indicate the number of GTP messages that are received and processed on the device. There are three entries: • Receive – Number of GTP messages received. • Forward – Number of GTP messages forwarded. • Drop – Number of GTP messages dropped because of an error. Number of Received counters = Number of Dropped counters + Number of Forward counters
Error counters	Drop reason and drop counters indicate the number of GTP packets that are dropped as a result of an error. Total error = Sum of all the following errors (see Sample Output)
Request counters	Request counters indicate the number of GTP request messages that are received and processed on the device. This information can be used for debugging purpose.
HA counters	HA counters indicate the number of messages that are received or sent by the device.

Copyright © 2016, Juniper Networks, Inc.

Table 255: show security gprs gtp counters all Output Fields *(continued)*

Field Name	Field Description
Path-rate-limit counters	Path-rate-limit counters indicate the number of PDP Create, Delete, Echo, and Other messages that are received and processed on the device after drop-threshold and alarm-threshold are reached. • Create Request – Number of create PDP messages. • Delete Request – Number of delete PDP messages. • Echo Request– Number of PDP Echo messages. • Others – Control messages other than the above messages. Drop - Indicate the number of packets dropped. Alarm – Indicate the number of packets transferred after the alarm-threshold is reached.

Sample Output

show security gprs gtp counters all

```
user@host> show security gprs gtp counters all


 Tunnel counters:
                 Total      GTPv0    GTPv1-c    GTPv1-u    GTPv2-c    GTPv2-u
  Allocated          0          0          0          0          0          0
  Freed              0          0          0          0          0          0

Packet counters:
                 Total      GTPv0      GTPv1      GTPv2       GTP'
  Received           0          0          0          0          0
  Passed             0          0          0          0          0
  Dropped            0          0          0          0          0

Brief message counters:
                                Receive    Forward       Drop
GTPv0
  Create PDP Request                  0          0          0
  Create PDP Response                 0          0          0
  Update PDP Request                  0          0          0
  Update PDP Response                 0          0          0
  Delete PDP Request                  0          0          0
  Delete PDP Response                 0          0          0
  Create AA PDP Request               0          0          0
  Create AA PDP Response              0          0          0
  Delete AA PDP Request               0          0          0
  Delete AA PDP Response              0          0          0
  Others                              0          0          0
GTPv1
  Create PDP Request                  0          0          0
  Create PDP Response                 0          0          0
  Update PDP Request                  0          0          0
  Update PDP Response                 0          0          0
  Delete PDP Request                  0          0          0
  Delete PDP Response                 0          0          0
  Others                              0          0          0
GTPv2
  Create Session Request              0          0          0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Create Session Response          0          0          0
  Delete Session Request           0          0          0
  Delete Session Response          0          0          0
  Create Bearer Request            0          0          0
  Create Bearer Response           0          0          0
  Modify Bearer Request            0          0          0
  Modify Bearer Response           0          0          0
  Delete Bearer Request            0          0          0
  Delete Bearer Response           0          0          0
  Others                           0          0          0

Error counters:
        Total error            : 0
        Exception              : 0
        Gate failed            : 0
        Invalid header         : 0
        Message length         : 0
        Zero IMSI              : 0
        Zero charge ID         : 0
        Sequence               : 0
        APN filter             : 0
        Port not match         : 0
        GTP-in-GTP             : 0
        Message too short      : 0
        Message too long       : 0
        GSN not exist          : 0
        Over GSN rate limit    : 0
        Request not found      : 0
        Retransmit response    : 0
        Missing IE             : 0
        Unexpected IE          : 0
        Unknown IE type        : 0
        IE order               : 0
        IE length              : 0
        Duplicate IE           : 0
        Non-digit TID/TEID     : 0
        Non-zero TID/TEID      : 0
        Zero TID/TEID          : 0
        Control TID/TEID       : 0
        Data TID/TEID          : 0
        Control GSN IE         : 0
        Data GSN IE            : 0
        End user IE            : 0
        GGSN IP for handover   : 0
        Disallowed v0 message  : 0
        Disallowed v1 message  : 0
        Disallowed v2 message  : 0
        Invalid message type   : 0
        No tunnel0             : 0
        No control tunnel      : 0
        No user tunnel         : 0
        Invalid tunnel0        : 0
        Invalid control tunnel : 0
        Invalid user tunnel    : 0
        Create tunnel0         : 0
        Create control tunnel  : 0
        Create user tunnel     : 0
        No request             : 0
        Out of request         : 0
        No action              : 0
        Out of action          : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
        GTPv2 TEID not exist   : 0
        GTPv2 Missing TEID     : 0
        GTPv2 Non-zero EBI     : 0
        GTPv2 EBI not found    : 0
        GTPv2 IE context       : 0

HA counters:
        Total message received              : 0
        Message received success            : 0
        Bad message received                : 0
        Unknown message type received       : 0
        Unknown message version received    : 0
        Total message send                  : 0
        Message send success                : 0
        Message send failed                 : 0
        Memory allocate failed              : 0

Request counters:
        Request allocated        : 0
        Request freed            : 0
        Request activated        : 0
        Request died             : 0
        Request action allocated : 0
        Request action freed     : 0

Path-rate-limit counters:
                             Drop           Alarm
  Create Request              0              0
  Delete Request              0              0
  Echo Request                0              0
  Others                      0              0
```

show security gprs gtp counters error

```
user@host> show security gprs gtp counters error
Error counters:
        Total error           : 0
        Exception             : 0
        Gate failed           : 0
        Invalid header        : 0
        Message length        : 0
        Zero IMSI             : 0
        Zero charge ID        : 0
        Sequence              : 0
        APN filter            : 0
        Port not match        : 0
        GTP-in-GTP            : 0
        Message too short     : 0
        Message too long      : 0
        GSN not exist         : 0
        Over GSN rate limit   : 0
        Request not found     : 0
        Retransmit response   : 0
        Missing IE            : 0
        Unexpected IE         : 0
        Unknown IE type       : 0
        IE order              : 0
        IE length             : 0
        Duplicate IE          : 0
        Non-digit TID/TEID    : 0
        Non-zero TID/TEID     : 0
```


Copyright © 2016, Juniper Networks, Inc.

```
        Zero TID/TEID          : 0
        Control TID/TEID       : 0
        Data TID/TEID          : 0
        Control GSN IE         : 0
        Data GSN IE            : 0
        End user IE            : 0
        GGSN IP for handover   : 0
        Disallowed v0 message  : 0
        Disallowed v1 message  : 0
        Disallowed v2 message  : 0
        Invalid message type   : 0
        No tunnel0             : 0
        No control tunnel      : 0
        No user tunnel         : 0
        Invalid tunnel0        : 0
        Invalid control tunnel : 0
        Invalid user tunnel    : 0
        Create tunnel0         : 0
        Create control tunnel  : 0
        Create user tunnel     : 0
        No request             : 0
        Out of request         : 0
        No action              : 0
        Out of action          : 0
        GTPv2 TEID not exist   : 0
        GTPv2 Missing TEID     : 0
        GTPv2 Non-zero EBI     : 0
        GTPv2 EBI not found    : 0
        GTPv2 IE context       : 0
```

show security gprs gtp counters ha

```
user@host> show security gprs gtp counters ha
HA counters:
        Total message received              : 0
        Message received success            : 0
        Bad message received                : 0
        Unknown message type received       : 0
        Unknown message version received    : 0
        Total message send                  : 0
        Message send success                : 0
        Message send failed                 : 0
        Memory allocate failed              : 0
```

show security gprs gtp counters message v0-create-aa-pdp-req

```
user@host> show security gprs gtp counters message v0-create-aa-pdp-req
Message counters:
  Received              0
  Forwarded             0
  Dropped               0
```

show security gprs gtp counters packet

```
user@host> show security gprs gtp counters packet
Packet counters:
                 Total      GTPv0      GTPv1      GTPv2       GTP'
  Received           0          0          0          0          0
  Passed             0          0          0          0          0
  Dropped            0          0          0          0          0
```

Copyright © 2016, Juniper Networks, Inc.

show security gprs gtp counters request

```
user@host> show security gprs gtp counters request

Request counters:
        Request allocated           : 0
        Request freed               : 0
        Request activated           : 0
        Request died                : 0
        Request action allocated    : 0
        Request action freed        : 0
```

show security gprs gtp counters tunnel

```
user@host> show security gprs gtp counters tunnel
Tunnel counters:
                 Total      GTPv0    GTPv1-c    GTPv1-u    GTPv2-c    GTPv2-u
  Allocated          0          0          0          0          0          0
  Freed              0          0          0          0          0          0
```

show security gprs gtp counters path-rate-limit

```
user@host> show security gprs gtp counters path-rate-limit
Path-rate-limit counters:
                           Drop          Alarm
  Create Request              0              0
  Delete Request              0              0
  Echo Request                0              0
  Others                      0              0
```


Copyright © 2016, Juniper Networks, Inc.

show security gprs gtp counters path-rate-limit

Syntax **show security gprs gtp counters path-rate-limit**

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Display information about path-rate-limit counters.

Required Privilege Level view

Related Documentation
- Security Configuration Statement Hierarchy on page 603

List of Sample Output show security gprs gtp counters path-rate-limit on page 2425

Output Fields Table 256 lists the output fields for the **show security gprs gtp counters path-rate-limit** command.

Table 256: show security gprs gtp counters path-rate-limit Output Fields

Field Name	Field Description
Create Request	Specify the number of create request messages received in a second after the alarm-threshold or drop-threshold is reached.
Delete Request	Specify the number of delete request messages received in a second after the alarm-threshold or drop-threshold is reached.
Echo Request	Specify the number of echo request messages received in a minute after the alarm-threshold or drop-threshold is reached.
Other messages	Specify the number of other GTP control messages received in a second after the alarm-threshold or drop-threshold is reached.
Drop	Display the number of packets dropped after the drop-threshold is reached.
Alarm	Display the number of packets received after the alarm-threshold is reached.

Sample Output

show security gprs gtp counters path-rate-limit

```
user@host> show security gprs gtp counters path-rate-limit

Path-rate-limit counters:
                                    Drop                    Alarm
Create Request                      200                      100
Delete Request                      300                      200
Echo Request                        600                      400
Others                              900                      800
```

show security gprs gtp gsn statistics

Syntax show security gprs gtp gsn statistics

Release Information Command introduced in Junos OS Release 12.1X46-D25.

Description Display a brief summary of GPRS support node (GSN) statistics, including active GSNs, obsolete GSNs, and the usage rate of each SPU.

Required Privilege Level view

Related Documentation
- show security gprs gtp counters path-rate-limit on page 2425

List of Sample Output show security gprs gtp gsn statistics on page 2426

Sample Output

show security gprs gtp gsn statistics

```
user@host> show security gprs gtp gsn statistics
FPC 1 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%

FPC 2 PIC 0:

Active GSNs: 0 Obsolete GSNs: 0 Use rate: 0%
```


Copyright © 2016, Juniper Networks, Inc.

show security gprs sctp association

Syntax

show security gprs sctp association
<all>
<destination-ip>
<destination-port>
<guid>
<init-state>
<source-ip>
<source-port>
<summary>

Release Information

Command introduced in Junos OS Release 12.1X44-D10. The **all**, **destination-ip**, **destination-port**, **guid**, **init**, **source-ip**, **source-port**, and **summary** options introduced in Junos OS Release 12.1X45-D10.

Description

Display the Stream Control Transmission Protocol (SCTP) association information.

Options

none—Display the live security SCTP association.

all—Display information about all the SCTP associations, both initiated and live.

destination-ip—Display information about the destination IP address associations.

destination-port—Display information about the destination port associations.

guid—Display information about the globally unique identifier associations.

init—Display information about initiated associations.

source-ip—Display information about the source IP address associations.

source-port—Display information about the source port associations.

summary—Display the output summary.

Required Privilege Level

view

Related Documentation

- clear security gprs sctp counters on page 2415
- clear security gprs sctp association on page 2413

List of Sample Output

show security gprs sctp association on page 2428

Output Fields

Table 257 lists the output fields for the **show security gprs sctp association** command. Output fields are listed in the approximate order in which they appear.

Table 257: show security gprs sctp association

Field Name	Field Description
Association Information	Association information of FPC and PIC.

Copyright © 2016, Juniper Networks, Inc.

Table 257: show security gprs sctp association *(continued)*

SCTP association numbers • **Total association**	Number of established SCTP associations. The SCTP association numbers field contains the total number of associations.
Association GUID	Globally unique association identifier information.

Sample Output

show security gprs sctp association

```
user@host>show security gprs sctp association
Association Information for FPC: 1 PIC: 0
SCTP association numbers:
Output association 0
Association Information for FPC: 1 PIC: 1

Association ID: 56b1d6a4-fc4b1872-0105000000000002 :
source:
2.0.0.8 (2.0.0.8)
2.0.0.9 (2.0.0.9)
2.0.0.10 (2.0.0.10)
2.0.0.11 (2.0.0.11)
2.0.0.12 (2.0.0.12)
2.0.0.13 (2.0.0.13)
2.0.0.14 (2.0.0.14)
2.0.0.15 (2.0.0.15)
port: 2906, state: SCTP_ESTABLISHED, tag: 0x000186a0;
rate limits to source:
destination:
3.0.0.8 (3.0.0.8)
3.0.0.9 (3.0.0.9)
3.0.0.10 (3.0.0.10)
3.0.0.11 (3.0.0.11)
3.0.0.12 (3.0.0.12)
3.0.0.13 (3.0.0.13)
3.0.0.14 (3.0.0.14)
3.0.0.15 (3.0.0.15)
port: 2906, state: SCTP_ESTABLISHED, tag: 0x00030d40;
rate limits to destination:
time left: 59995 s, access time: 4234 s;
policy name/id: sctp_policy/4, cfg live timeout: 1000 min, handshake timeout: 30
 s.
redirect SPU id for source: 0, redirect SPU id for destination: 0.

SCTP association numbers:
Output association 1
```


Copyright © 2016, Juniper Networks, Inc.

show security gprs sctp counters

Syntax **show security gprs sctp counters <detail>**

Release Information Command introduced in Junos OS Release 10.2. Support for the **detail** option added in Junos OS Release 12.1X45-D10. Support for SCTP payload protocols chunk counters added in Junos OS Release 12.1X47-D10.

Description Display the statistics of the received and dropped Stream Control Transmission Protocol (SCTP) chunks.

Options **none**—Display the statistics of all received and dropped SCTP chunks.

detail—Display detailed debugging counters for SCTP chunks.

Required Privilege Level view

Related Documentation
- clear security gprs sctp counters on page 2415

List of Sample Output
show security gprs sctp counters on page 2430
show security gprs sctp counters detail on page 2433

Output Fields Table 258 lists the output fields for the **show security gprs sctp counters** command. Output fields are listed in the approximate order in which they appear.

Table 258: show security gprs sctp counters

Field Name	Field Description
Name	Name of the SCTP payload protocol identifier.
Received Counter	Number of SCTP chunk counters received.
Drop Counter	Number of SCTP chunk counters dropped due to error.
Counter Information	Association information of FPC and PIC.
Association detail counters	(detail output only) Number of total and dying associations.
Dbg records	(detail output only) Number and type of debugging records.
Packet error	(detail output only) Number and type of packet errors.
Association matching error	(detail output only) Number of association matching errors.
Association state error	(detail output only) Number of state errors.
Over rate drop	(detail output only) Number of messages over the rate limit.

Copyright © 2016, Juniper Networks, Inc.

Table 258: show security gprs sctp counters *(continued)*

Field Name	Field Description
Memory counters	(detail output only) Number and type of memory counters.
Other error	(detail output only) Number and type of other errors.

Sample Output

show security gprs sctp counters

```
user@host> show security gprs sctp counters
Counter Information for FPC: 1 PIC: 0
Association detail counters:
Total association: 0
Dying association: 0
Ready wrap: 0

Dbg records:
pak-without-profile : 0
pak-nat-only : 0
pak-inspection : 0
drop-at-clearing-all : 0
src-pnat : 0
dst-pnat : 0
hostname : 0
dup-init : 0
dup-initack : 0
tag-null-abort : 0
error-chunk : 0
bad-interest : 0
wing-attach : 0
wing-detach : 0
wrap-with-assoc : 0
unwrap-from-assoc : 0
conflict-assoc : 0
conflict-redr : 0
wrong-distribution : 0

Packet error:
chunk-unsupport : 0
cookie-invalid : 0
pkt-len : 0
chunk-len : 0
tag-error : 0
bad-len : 0
bad-chk-hdr : 0

Association matching error:
ha-assoc : 0
data-assoc : 0
initack-assoc : 0
sack-assoc : 0
hb-assoc : 0
hb-ack-assoc : 0
abort-assoc : 0
shutdown-assoc : 0
shutdown-ack-assoc : 0
```


Copyright © 2016, Juniper Networks, Inc.

```
err-assoc : 0
cookie-echo-assoc : 0
cookie-ack-assoc : 0
shutdown-complete-assoc : 0
lookup-no-assoc : 0
dup-init-diff-ip-list : 0
dup-init-diff-dst-ip : 0
dup-initack-src-ip-invalid : 0
dup-initack-diff-ip-lis : 0

Associaiton state error:
data-state : 0
init-state : 0
initack-state : 0
sack-state : 0
shutdown-state : 0
shutdown-ack-state : 0
cookie-echo-state : 0
cookie-ack-state : 0
shutdown-complete-state : 0
cookie-echo-retrans-timeout : 0
cookie-ack-retrans-timeout : 0

Association LoadBalance counter:
redirect-assoc-request-send : 0
redirect-assoc-request-ack-recv : 0
redirect-assoc-request-nack-recv : 0
redirect-assoc-request-ack-timeoute : 0
redirect-assoc-request-recv : 0
redirect-assoc-request-ack-send : 0
redirect-assoc-request-nack-send : 0

Over rate drop:
sccp : 0
ssp : 0
sst : 0

Memory counters:
alloc-assoc : 0
free-assoc : 0
alloc-redr : 0
free-redr : 0
alloc-assoc-wrap : 0
free-assoc-wrap : 0
alloc-cookie : 0
free-cookie : 0
alloc-addr : 0
free-addr : 0

HA counters:
invalid-type : 0
bad-msg : 0
no-assoc-info : 0
send-fail : 0
dup-create : 0
no-policy : 0
no-profile : 0
alloc-fail : 0
non-established-issu: 0

Other error:
```

Copyright © 2016, Juniper Networks, Inc.

```
over-max : 0
over-min : 0
del-error : 0
sess-cookie-set-fail : 0
sess-cookie-get-fail : 0
no-assoc-install-redr-cb : 0
wrap-alloced-failure : 0
wrap-null-assocp : 0
assoc-alloced-failure : 0
redr-assoc-alloced-failure : 0
invalid-pkt-pointer : 0
nat-jbuf-alloc-fail : 0
evt-cookie-alloc-fail : 0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security gprs sctp counters detail

```
user@host> show security gprs sctp counters detail
Counter Information for FPC: 1 PIC: 0
Association detail counters:
Total association: 0
Dying association: 0
Ready wrap: 0

Dbg records:
pak-without-profile : 0
pak-nat-only : 0
pak-inspection : 0
drop-at-clearing-all : 0
src-pnat : 0
dst-pnat : 0
hostname : 0
dup-init : 0
dup-initack : 0
tag-null-abort : 0
error-chunk : 0
bad-interest : 0
wing-attach : 0
wing-detach : 0
wrap-with-assoc : 0
unwrap-from-assoc : 0
conflict-assoc : 0
conflict-redr : 0
wrong-distribution : 0

Packet error:
chunk-unsupport : 0
cookie-invalid : 0
pkt-len : 0
chunk-len : 0
tag-error : 0
bad-len : 0
bad-chk-hdr : 0

Association matching error:
ha-assoc : 0
data-assoc : 0
initack-assoc : 0
sack-assoc : 0
hb-assoc : 0
hb-ack-assoc : 0
abort-assoc : 0
shutdown-assoc : 0
shutdown-ack-assoc : 0
err-assoc : 0
cookie-echo-assoc : 0
cookie-ack-assoc : 0
shutdown-complete-assoc : 0
lookup-no-assoc : 0
dup-init-diff-ip-list : 0
dup-init-diff-dst-ip : 0
dup-initack-src-ip-invalid : 0
dup-initack-diff-ip-lis : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Associaiton state error:
data-state : 0
init-state : 0
initack-state : 0
sack-state : 0
shutdown-state : 0
shutdown-ack-state : 0
cookie-echo-state : 0
cookie-ack-state : 0
shutdown-complete-state : 0
cookie-echo-retrans-timeout : 0
cookie-ack-retrans-timeout : 0

Association LoadBalance counter:
redirect-assoc-request-send : 0
redirect-assoc-request-ack-recv : 0
redirect-assoc-request-nack-recv : 0
redirect-assoc-request-ack-timeoute : 0
redirect-assoc-request-recv : 0
redirect-assoc-request-ack-send : 0
redirect-assoc-request-nack-send : 0

Over rate drop:
sccp : 0
ssp : 0
sst : 0

Memory counters:
alloc-assoc : 0
free-assoc : 0
alloc-redr : 0
free-redr : 0
alloc-assoc-wrap : 0
free-assoc-wrap : 0
alloc-cookie : 0
free-cookie : 0
alloc-addr : 0
free-addr : 0

HA counters:
invalid-type : 0
bad-msg : 0
no-assoc-info : 0
send-fail : 0
dup-create : 0
no-policy : 0
no-profile : 0
alloc-fail : 0
non-established-issu: 0

Other error:
over-max : 0
over-min : 0
del-error : 0
sess-cookie-set-fail : 0
sess-cookie-get-fail : 0
no-assoc-install-redr-cb : 0
wrap-alloced-failure : 0
wrap-null-assocp : 0
assoc-alloced-failure : 0
redr-assoc-alloced-failure : 0
```


Copyright © 2016, Juniper Networks, Inc.

```
invalid-pkt-pointer : 0
nat-jbuf-alloc-fail : 0
evt-cookie-alloc-fail : 0
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Interfaces Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 34

Overview

- Introduction to Interfaces on page 2441
- Configuring Interface Logical Properties on page 2451
- Understanding Interface Physical Properties on page 2475
- Configuring VLAN Tagging on page 2483

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 108

Introduction to Interfaces

- Understanding Interfaces on page 2441
- Network Interfaces on page 2442
- Services Interfaces on page 2443
- Special Interfaces on page 2445
- Interface Naming Conventions on page 2446
- Understanding the Data Link Layer on page 2447
- Monitoring Interfaces on page 2449

Understanding Interfaces

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

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

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

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

Related Documentation

- Interface Naming Conventions on page 2446
- Understanding Interface Logical Properties on page 2451
- Understanding Interface Physical Properties on page 2475

Copyright © 2016, Juniper Networks, Inc.

- Understanding the Data Link Layer on page 2447

Network Interfaces

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

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

Table 259 describes network interfaces that are available on SRX Series devices.

Table 259: Network Interfaces

Interface Name	Description
ae	Aggregated Ethernet interface. See "Understanding Aggregated Ethernet Interfaces" on page 2557.
at	ATM-over-ADSL or ATM-over-SHDSL WAN interface. See *ADSL and SHDSL Interfaces Configuration Overview*.
dl	Dialer interface for initiating USB modem connections. See *USB Modem Interface Overview*.
e1	E1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 2491.
e3	E3 (also called DS3) WAN interface. See *Understanding T3 and E3 Interfaces*.
fe	Fast Ethernet interface. See "Understanding Ethernet Interfaces" on page 2541.
ge	Gigabit Ethernet interface. See "Understanding Ethernet Interfaces" on page 2541.
pt	VDSL2 interface. See "Example: Configuring VDSL2 Interfaces (Detail)" on page 2513.
reth	For chassis cluster configurations only, redundant Ethernet interface. See "Understanding Ethernet Interfaces" on page 2541.
se	Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21). See *Serial Interfaces Overview*.
t1	T1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 2491.
t3	T3 (also called DS3) WAN interface. See *Understanding T3 and E3 Interfaces*.
wx	WXC Integrated Services Module (ISM 200) interface for WAN acceleration. See the WXC Integrated Services Module Installation and Configuration.
xe	10-Gigabit Ethernet interface. See "Understanding the 2-Port 10-Gigabit Ethernet XPIM" on page 2595.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Interfaces on page 2441
- Services Interfaces on page 2443
- Special Interfaces on page 2445

Services Interfaces

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

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

Table 260 describes services interfaces that you can configure on SRX Series Services Gateways.

Table 260: Configurable Services Interfaces

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

Copyright © 2016, Juniper Networks, Inc.

Table 260: Configurable Services Interfaces *(continued)*

Interface Name	Description
lsq-0/0/0	Configurable link services queuing interface. Link services include the multilink services MLPPP, MLFR, and Compressed Real-Time Transport Protocol (CRTP). Packets are routed to this internal interface for link bundling or compression. The link services interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform multilink services. NOTE: The ls-0/0/0 interface has been deprecated. All multiclass multilink features supported by ls-0/0/0 are now supported by lsq-0/0/0. See *Link Services Interfaces Overview*.
lt-0/0/0	Configurable logical tunnel interface that interconnects logical systems on SRX Series devices. See the *Logical Systems Feature Guide for Security Devices*.
pp0	Configurable PPPoE encapsulation interface. PPP packets being routed in an Ethernet network use PPPoE encapsulation. Packets are routed to this internal interface for PPPoE encapsulation. The PPPoE encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to forward PPPoE traffic. See "Understanding Point-to-Point Protocol over Ethernet" on page 2631.
ppd0	Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a device, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the **[edit protocol pim]** hierarchy to perform PIM de-encapsulation. Use the **show pim interfaces** command to check the status of ppd0 interface.
ppe0	Protocol Independent Multicast (PIM) encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a device, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the **[edit protocol pim]** hierarchy to perform PIM encapsulation.
st0	Secure tunnel interface used for IPSec VPNs. See the *VPN Feature Guide for Security Devices*.
umd0	Configurable USB modem physical interface. This interface is detected when a USB modem is connected to the USB port on the device. See *USB Modem Configuration Overview*.

Copyright © 2016, Juniper Networks, Inc.

Table 261 describes non-configurable services interfaces for SRX Series Services Gateways.

Table 261: Non-Configurable Services Interfaces

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

Related Documentation

- Junos Services Interfaces Configuration
- Understanding Interfaces on page 2441
- Network Interfaces on page 2442
- Special Interfaces on page 2445

Special Interfaces

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

Table 262 describes special interfaces for SRX Series Services Gateways.

Table 262: Special Interfaces

Interface Name	Description
fxp0, fxp1	On SRX Series devices, the fxp0 management interface is a dedicated port located on the Routing Engine.

Copyright © 2016, Juniper Networks, Inc.

Table 262: Special Interfaces *(continued)*

Interface Name	Description
lo0	Loopback address. The loopback address has several uses, depending on the particular Junos feature being configured.
dsc	Discard interface.

Related Documentation

- Understanding Interfaces on page 2441
- Network Interfaces on page 2442
- Services Interfaces on page 2443

Interface Naming Conventions

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

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

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

  type-slot/pim-or-ioc/port

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

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

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

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

The parts of an interface name are summarized in Table 263.

Table 263: Network Interface Names

Name Part	Meaning	Possible Values
type	Type of network medium that can connect to this interface.	ae, at, ei, e3, fe, fxp0, fxp1, ge, lo0, lsq, lt, ppo, pt, sto, t1, t3, xe, and so on.
slot	Number of the chassis slot in which a PIM or IOC is installed.	SRX5600 and SRX5800 devices: The slot number begins at **0** and increases as follows from left to right, bottom to top: • SRX5600 device—Slots 0 to 5 • SRX5800 device—Slots 0 to 5, 7 to 11

Copyright © 2016, Juniper Networks, Inc.

Table 263: Network Interface Names *(continued)*

Name Part	Meaning	Possible Values
pim-or-ioc	Number of the PIM or IOC on which the physical interface is located.	SRX5600 and SRX5800 devices: For 40-port Gigabit Ethernet IOCs or 4-port 10-Gigabit Ethernet IOCs, this number can be **0**, **1**, **2**, or **3**.
port	Number of the port on a PIM or IOC on which the physical interface is located.	On SRX5600 and SRX5800 devices: • For 40-port Gigabit Ethernet IOCs, this number begins at **0** and increases from left to right to a maximum of **9**. • For 4-port 10-Gigabit Ethernet IOCs, this number is always **0**. Port numbers appear on the PIM or IOC faceplate.
channel	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	• On an E1 interface, a value from **1** through **31**. The **1** time slot is reserved. • On a T1 interface, a value from **1** through **24**.
unit	Number of the logical interface created on a physical interface.	A value from **0** through **16384**. If no logical interface number is specified, unit **0** is the default, but must be explicitly configured. In addition to user-configured interfaces, there are some logical interfaces that are created dynamically. Hence, for Junos OS, the maximum limit for configuring logical interfaces is 2,62,143 (user configured and dynamically created). Based on performance, for each platform, the maximum number of logical interfaces supported can vary.

Related Documentation

- Understanding Interfaces on page 2441

Understanding the Data Link Layer

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

- Physical Addressing on page 2448
- Network Topology on page 2448
- Error Notification on page 2448
- Frame Sequencing on page 2448
- Flow Control on page 2448
- Data Link Sublayers on page 2448
- MAC Addressing on page 2449

Copyright © 2016, Juniper Networks, Inc.

Physical Addressing

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

Network Topology

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

Error Notification

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

Frame Sequencing

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

Flow Control

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

Data Link Sublayers

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

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

Copyright © 2016, Juniper Networks, Inc.

MAC Addressing

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

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

Most Layer 2 networks use one of three primary numbering spaces—MAC-48, EUI-48 (extended unique identifier), and EUI-64—which are all globally unique. MAC-48 and EUI-48 spaces each use 48-bit addresses, and EUI-64 spaces use a 64-bit addresses, but all three use the same numbering format. MAC-48 addresses identify network hardware, and EUI-48 addresses identify other devices and software.

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

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

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

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

Related Documentation

- Understanding Interfaces on page 2441

Monitoring Interfaces

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

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

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

Copyright © 2016, Juniper Networks, Inc.

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

To change the interface display, use the following options:

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

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

- **show interfaces terse**

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

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

Related Documentation

- *Monitoring Overview*
- *Monitoring Address Pools*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 109

Configuring Interface Logical Properties

- Understanding Interface Logical Properties on page 2451
- Understanding Protocol Families on page 2452
- Understanding IPv4 Addressing on page 2453
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Configuring the inet6 IPv6 Protocol Family on page 2459
- Enabling Flow-Based Processing for IPv6 Traffic on page 2460
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461
- Understanding IPv6 Support VDSL2 Interfaces on page 2470
- Understanding MAC Limiting on Layer 3 Routing Interfaces on page 2471

Understanding Interface Logical Properties

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

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

Related Documentation

- Understanding Interfaces on page 2441
- Understanding Protocol Families on page 2452
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Understanding Virtual LANs on page 2483

Copyright © 2016, Juniper Networks, Inc.

Understanding Protocol Families

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

This topic contains the following sections:

- Common Protocol Suites on page 2452
- Other Protocol Suites on page 2452

Common Protocol Suites

Junos OS protocol families include the following common protocol suites:

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

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

Other Protocol Suites

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

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

Related Documentation

- Understanding Interface Logical Properties on page 2451

Copyright © 2016, Juniper Networks, Inc.

Understanding IPv4 Addressing

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

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

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

This topic contains the following sections:

- IPv4 Classful Addressing on page 2453
- IPv4 Dotted Decimal Notation on page 2454
- IPv4 Subnetting on page 2454
- IPv4 Variable-Length Subnet Masks on page 2455

IPv4 Classful Addressing

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

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

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

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

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

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

Copyright © 2016, Juniper Networks, Inc.

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

IPv4 Dotted Decimal Notation

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

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

IPv4 Subnetting

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

Figure 112 shows two subnets in a network.

Figure 112: Subnets in a Network

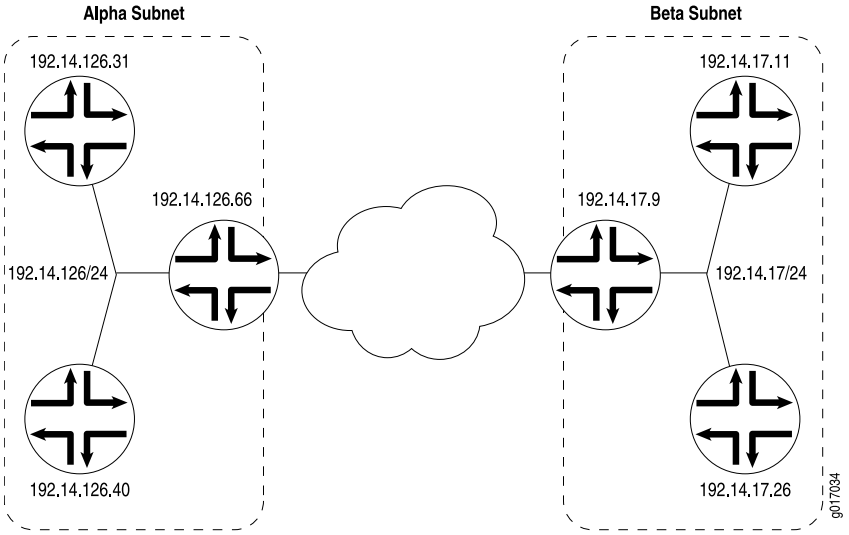

Figure 112 shows three devices connected to one subnet and three more devices connected to a second subnet. Collectively, the six devices and two subnets make up the larger network. In this example, the network is assigned the network prefix **192.14.0.0**, a class B address. Each device has an IP address that falls within this network prefix.

In addition to sharing a network prefix (the first two octets), the devices on each subnet share a third octet. The third octet identifies the subnet. All devices on a subnet must have the same subnet address. In this case, the alpha subnet has the IP address **192.14.126.0** and the beta subnet has the IP address **192.14.17.0**.

Copyright © 2016, Juniper Networks, Inc.

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

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

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

IPv4 Variable-Length Subnet Masks

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

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

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

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

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

The subnet mask includes 27 significant digits.

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

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

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

Related Documentation

- Understanding Interface Logical Properties on page 2451
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455

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

Understanding IP Version 6 (IPv6)

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

Copyright © 2016, Juniper Networks, Inc.

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

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

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

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

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

- Unicast

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

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

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

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

- Multicast

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

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

Copyright © 2016, Juniper Networks, Inc.

- Anycast

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

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

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

IPv6 Address Scope

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

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

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

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

IPv6 Address Structure

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

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

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

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

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

Copyright © 2016, Juniper Networks, Inc.

Understanding IPv6 Address Space, Addressing, and Address Types

Addressing is the area where most of the differences between IP version 4 (IPv4) and IPv6 exist, but the changes are largely about the ways in which addresses are implemented and used. IPv6 has a vastly larger address space than the impending exhausted IPv4 address space. IPv6 increases the size of the IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of the address space.

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

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

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

Understanding IPv6 Address Format

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

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

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

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

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

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

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

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

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

An IPv6 address prefix is a combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form *ipv6-prefix/prefix-length* and represents a block of address space (or a network). The *ipv6-prefix* variable follows general IPv6 addressing rules. The

Copyright © 2016, Juniper Networks, Inc.

/prefix-length variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 10FA:6604:8136:6502::/64 is a possible IPv6 prefix.

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

Limitations

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

Related Documentation

- About the IPv6 Basic Packet Header on page 1764
- Understanding IPv6 Packet Header Extensions on page 1766

Configuring the inet6 IPv6 Protocol Family

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

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

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

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

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

Copyright © 2016, Juniper Networks, Inc.

```
> tcc                   Translational cross-connect parameters
> vpls                  Virtual private LAN service parameters

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

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Enabling Flow-Based Processing for IPv6 Traffic on page 2460

Enabling Flow-Based Processing for IPv6 Traffic

You have the following options for handling IPv6 traffic:

- Drop—Do not forward IPv6 packets. This is the default behavior.
- Packet-based forwarding—Do not create a session and process according to packet-based features only (includes firewall filters and class of service).
- Flow-based forwarding—Create a session and process according to packet-based features (including firewall filters and class of service) but also flow-based security features, such as screens and firewall security policy.

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

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

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

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


Copyright © 2016, Juniper Networks, Inc.

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

If you change the forwarding option mode for IPv6, you might need to perform a reboot to initialize the configuration change. Table 172 summarizes device status upon configuration change.

Table 264: Device Status Upon Configuration Change

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

For details on packet-based and flow-based processing, see the *Flow-Based and Packet-Based Processing Feature Guide for Security Devices*.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Configuring the inet6 IPv6 Protocol Family on page 2459

Configuring Flow Aggregation to Use Version 9 Flow Templates

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

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

Copyright © 2016, Juniper Networks, Inc.

The following sections contain additional information:

- Configuring the Traffic to Be Sampled on page 2462
- Configuring the Version 9 Template Properties on page 2462
- Restrictions on page 2464
- Fields Included in Each Template Type on page 2465
- inet Sampling Behavior on page 2466
- Verification on page 2467
- Examples: Configuring Version 9 Flow Templates on page 2467

Configuring the Traffic to Be Sampled

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

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

You can include **family inet**.

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

Configuring the Version 9 Template Properties

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

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

The following details apply to the configuration statements:

Copyright © 2016, Juniper Networks, Inc.

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

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

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

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

Copyright © 2016, Juniper Networks, Inc.

Restrictions

The following restrictions apply to version 9 templates:

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

  inet | Layer 2 Header | IPv4

  In this case, **inet-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the inet header. Packets are dropped on the PIC and are accounted as parser errors.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.

On all branch SRX Series devices, flow monitoring IPv6 version 9 has the following limitations:

- MPLS in not supported.
- User-defined version 9 templates are not supported.
- Routing Engine based flow monitoring version 9 is not supported.
- Flow monitoring and accounting are not supported in chassis cluster mode.
- Flow monitoring and accounting are not supported on an ae interface.
- J-Web for IPv6 sampled packets is not supported.
- SNMP queries for IPv6 sampled packets are not supported
- Flow monitoring can be configured in version 5, version 8, or version 9 export mode. Up to eight version 9 collectors are supported in export mode.
- Scope of accounting of IPv6 flow monitoring version 9 packets associated with pseudointerfaces (such as IRB, ML, LAG, VLAN, and GRE) is not supported.
- Creation of an SCTP session (parallel to TCP) between an exporter and a collector for gathering flow monitoring information is not supported.
- Maximum flow sessions that might be supported include:

Copyright © 2016, Juniper Networks, Inc.

    - A device with 1-GB RAM, such as an SRX320 device, might support up to 15,000 flow monitoring sessions at a time.

- Routing Engine based flow monitoring V5 or V8 mode is mutually exclusive with inline flow monitoring V9.
- High-end SRX Series devices do not support multiple collectors like branch SRX Series devices. Only one V9 collector per IPv4 or IPv6 is supported
- Flow aggregation for V9 export is not supported.
- Only UDP over IPv4 or IPv6 protocol can be used as the transport protocol.
- Only the standard IPv4 or IPv6 template is supported for exporting flow monitoring records.
- User-defined or special templates are not supported for exporting flow monitoring records.
- Chassis cluster is supported without flow monitoring session synchronization.

Fields Included in Each Template Type

The following fields are common to all template types:

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

The IPv4 template includes the following specific fields:

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

The IPv6 template includes the following specific fields:

Copyright © 2016, Juniper Networks, Inc.

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

The inet template includes the following specific fields:

- inet Label #1
- inet Label #2
- inet Label #3
- inet EXP Information
- FEC IP Address

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

The peer AS billing template includes the following specific fields:

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

inet Sampling Behavior

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

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

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

Copyright © 2016, Juniper Networks, Inc.

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

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

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

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

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

Verification

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

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

Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

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

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

Copyright © 2016, Juniper Networks, Inc.

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

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

```
inline-jflows {
    at-0/1/1 {
        unit 0 {
            family inet {
                filter {
                    input inet_sample;
                }
            }
        }
    }
    sp-7/0/0 {
        unit 0 {
            family inet;
            family inet;
        }
    }
}
```

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

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


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
```

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

```
firewall {
  family inet {
    filter peer-as-filter {
      term 0 {
          from {
            destination-class dcu-1;
            inline-jflow ge-2/1/0;
            forwarding-class class-1;
          }
          then count count_team_0;
      }
      }
      term 1 {
          from {
            destination-class dcu-2;
            inline-jflow ge-2/1/0;
            forwarding-class class-1;
          }
          then count count_team_1;
      }
      term 2 {
          from {
            destination-class dcu-3;
            inline-jflow ge-2/1/0;
            forwarding-class class-1;
          }
          then count count_team_2;
      }
    }
  }
}
```

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

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

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

```
forwarding-options {
  sampling {
    }
    input {
      rate 1;
    }
    family inet {
```

Copyright © 2016, Juniper Networks, Inc.

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

Related Documentation

- Understanding Interface Logical Properties on page 2451

Understanding IPv6 Support VDSL2 Interfaces

The branch SRX Series devices support IPv6 on the following DSL encapsulations:

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

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

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

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Understanding Interface Logical Properties on page 2451

Understanding MAC Limiting on Layer 3 Routing Interfaces

- Overview on page 2471
- Limitations on page 2473

Overview

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

Both the physical interface level **source-address-filter** and logical interface level **accept-source-mac** configurations are supported on SRX300, SRX320, and SRX340 devices. The following considerations apply when you configure the **source-address-filter** and **accept-source-mac** statements:

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

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

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

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

Copyright © 2016, Juniper Networks, Inc.

- Physical level MAC filter configuration on an untagged interface

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

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

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

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

Copyright © 2016, Juniper Networks, Inc.

Limitations

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

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

  MAC filtering is not supported on chassis clusters.

Related Documentation

- Understanding Interface Logical Properties on page 2451

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 110

Understanding Interface Physical Properties

- Understanding Interface Physical Properties on page 2475
- Understanding Bit Error Rate Testing on page 2476
- Understanding Interface Clocking on page 2477
- Understanding Frame Check Sequences on page 2478
- MTU Default and Maximum Values on page 2479
- Understanding Jumbo Frames Support for Ethernet Interfaces on page 2481

Understanding Interface Physical Properties

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

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

Table 265 summarizes some key physical properties of device interfaces.

Table 265: Interface Physical Properties

Physical Property	Description
bert-error-rate	Bit error rate (BER). The error rate specifies the number of bit errors in a particular bit error rate test (BERT) period required to generate a BERT error condition. See "Understanding Bit Error Rate Testing" on page 2476.
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See "Understanding Bit Error Rate Testing" on page 2476.
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying **chap** enables CHAP authentication on the interface. See "Understanding CHAP Authentication on a PPPoE Interface" on page 2642.

Copyright © 2016, Juniper Networks, Inc.

Table 265: Interface Physical Properties *(continued)*

Physical Property	Description
clocking	Clock source for the link. Clocking can be provided by the local system (internal) or a remote endpoint on the link (external). By default, all interfaces use the internal clocking mode. If an interface is configured to accept an external clock source, one adjacent interface must be configured to act as a clock source. Under this configuration, the interface operates in a loop timing mode, in which the clocking signal is unique for that individual network segment or loop. See "Understanding Interface Clocking" on page 2477.
description	A user-defined text description of the interface, often used to describe the interface's purpose.
disable	Administratively disables the interface.
encapsulation	Type of encapsulation on the interface. Common encapsulation types include PPP, Frame Relay, Cisco HDLC, and PPP over Ethernet (PPPoE). See "Understanding Physical Encapsulation on an Interface" on page 2623.
fcs	Frame check sequence (FCS). FCS is an error-detection scheme that appends parity bits to a digital signal and uses decoding algorithms that detect errors in the received digital signal.
mtu	Maximum transmission unit (MTU) size. MTU is the largest size packet or frame, specified in bytes or octets, that can be sent in a packet-based or frame-based network. The TCP uses MTU to determine the maximum size of each packet in any transmission. See "MTU Default and Maximum Values" on page 2479.
no-keepalives	Disabling of keepalive messages across a physical link. A keepalive message is sent between network devices to indicate that they are still active. Keepalives help determine whether the interface is operating correctly. Except for ATM-over-ADSL interfaces, all interfaces use keepalives by default.
pap	Password Authentication Protocol (PAP). Specifying **pap** enables PAP authentication on the interface. See "Understanding CHAP Authentication on a PPPoE Interface" on page 2642.
payload-scrambler	Scrambling of traffic transmitted out the interface. Payload scrambling randomizes the data payload of transmitted packets. Scrambling eliminates nonvariable bit patterns (strings of all 1s or all 0s) that generate link-layer errors across some physical links.

Related Documentation

- Understanding Interfaces on page 2441
- Understanding Bit Error Rate Testing on page 2476
- Understanding Interface Clocking on page 2477
- Understanding Frame Check Sequences on page 2478
- MTU Default and Maximum Values on page 2479

Understanding Bit Error Rate Testing

In telecommunication transmission, the bit error rate (BER) is the percentage of bits that have errors compared to the total number of bits received in a transmission, usually expressed as 10 to a negative power. For example, a transmission with a BER of 10^{-6} received 1 errored bit in 1,000,000 bits transmitted. The BER indicates how often a packet

Copyright © 2016, Juniper Networks, Inc.

or other data unit must be retransmitted because of an error. If the BER is too high, a slower data rate might improve the overall transmission time for a given amount of data if it reduces the BER and thereby lowers the number of resent packets.

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

Related Documentation

- Understanding Interface Physical Properties on page 2475

Understanding Interface Clocking

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

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

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

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

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

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

This topic contains the following sections:

- Data Stream Clocking on page 2478
- Explicit Clocking Signal Transmission on page 2478

Copyright © 2016, Juniper Networks, Inc.

Data Stream Clocking

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

Explicit Clocking Signal Transmission

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

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

Related Documentation

- Understanding Interface Physical Properties on page 2475

Understanding Frame Check Sequences

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

This topic contains the following sections:

- Cyclic Redundancy Checks and Checksums on page 2478
- Two-Dimensional Parity on page 2479

Cyclic Redundancy Checks and Checksums

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

If the result of the second calculation matches the contents of the FCS field, the packet was sent and received without bit errors. If the values do not match, an FCS error is generated, the frame is discarded and the originating host is notified of the error.

Checksums function similarly to CRCs, but use a different algorithm.

Copyright © 2016, Juniper Networks, Inc.

Two-Dimensional Parity

On a link that uses two-dimensional parity bits for frame checking, the sending and receiving hosts examine each frame in the total packet transmission and create a parity byte that is evaluated to detect transmission errors.

For example, a host can create the parity byte for the following frame sequence by summing up each column (each bit position in the frame) and keeping only the least-significant bit:

```
Frame 1       0 1 0 1 0 0 1
Frame 2       1 1 0 1 0 0 1
Frame 3       1 0 1 1 1 1 0
Frame 4       0 0 0 1 1 1 0
Frame 5       0 1 1 0 1 0 0
Frame 6       1 0 1 1 1 1 1

Parity Byte   1 1 1 1 0 1 1
```

If the sum of the bit values in a bit position is even, the parity bit for the position is 0. If the sum is odd, the parity bit is 1. This method is called even parity. Matching parity bytes on the originating and receiving hosts indicate that the packet was received without error.

Related Documentation

- Understanding Interface Physical Properties on page 2475

MTU Default and Maximum Values

The MTU values are by default without any MTU configurations. If the MTU value is set, then the formula **IFF MTU (IP MTU) = IFD MTU (Media MTU) – L2 Overhead** is applicable. See Table 266 for default MTU values.

NOTE: For ATM MLPPP irrespective of UIFD MTU, the IP MTU is always 1500 because the IP MTU calculation is based on the LSQ interface. Even if you configure the LSQ family MTU, the IP MTU value cannot exceed 1504.

Table 266 lists MTU values for the SRX Series Services Gateways Physical Interface Modules (PIMs).

Table 266: MTU Values for the SRX Series Services Gateways PIMs

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
1-Port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM	1514	9010	1500
1-Port Small Form-Factor Pluggable (SFP) Mini-PIM	1514	1518	1500
DOCSIS Mini-PIM	1504	1504	1500

Copyright © 2016, Juniper Networks, Inc.

Table 266: MTU Values for the SRX Series Services Gateways PIMs *(continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
Serial Mini-PIM	1504	2000	1500
T1/E1 Mini-PIM	1504	2000	1500
Dual CT1/E1 GPIM	1504	9000	1500
Quad CT1/E1 GPIM	1504	9000	1500
2-Port 10- Gigabit Ethernet XPIM	1514	9192	1500
16-Port Gigabit Ethernet XPIM	1514	9192	1500
24-Port Gigabit Ethernet XPIM	1514	9192	1500
ADSL2+ Mini-PIM (Encapsulation)			
atm-snap	1512	1512	1504
atm-vcmux	1512	1512	1512
atm-nlpid	1512	1512	1508
atm-cisco-nlpid	1512	1512	1510
ether-over-atm-llc	1512	1512	1488
atm-ppp-llc	1512	1512	1506
atm-ppp-vcmux	1512	1512	1510
atm-mlppp-llc	1512	1512	1500
ppp-over-ether-over-atm-llc	1512	1512	1480
VDSL- Mini-PIM AT mode (Encapsulation)			
atm-snap	1514	1514	1506
atm-vcmux	1514	1514	1514
atm-nlpid	1514	1514	1510
atm-cisco-nlpid	1514	1514	1512

Copyright © 2016, Juniper Networks, Inc.

Table 266: MTU Values for the SRX Series Services Gateways PIMs *(continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
ether-over-atm-llc	1514	1524	1490
atm-ppp-llc	1514	1514	1508
atm-ppp-vcmux	1514	1514	1512
atm-mlppp-llc	1514	1514	1500
ppp-over-ether-over-atm-llc	1514	1514	1482
VDSL- Mini-PIM PT mode	1514	1514	1500
G.SHDSL Mini-PIM AT mode (Encapsulation)			
atm-snap	4482	4482	4470
atm-vcmux	4482	4482	4470
atm-nlpid	4482	4482	4470
atm-cisco-nlpid	4482	4482	4470
ether-over-atm-llc	4482	4482	1500
atm-ppp-llc	4482	4482	4476
atm-ppp-vcmux	4482	4482	4480
atm-mlppp-llc	4482	4482	1500
ppp-over-ether-over-atm-llc	4482	4482	1492
G.SHDSL Mini-PIM PT mode	1514	1514	1500

Related Documentation

- Understanding Interface Physical Properties on page 2475

Understanding Jumbo Frames Support for Ethernet Interfaces

SRX Series devices support jumbo frames up to 9192 bytes.

Jumbo frames are Ethernet frames with more than 1500 bytes of payload (maximum transmission unit [MTU]). Jumbo frames can carry up to 9000 bytes of payload.

Copyright © 2016, Juniper Networks, Inc.

You configure jumbo frames at the physical interface by using the following command:

set interface *interface-name* mtu *mtu-value*

Example:

```
user@host# set interfaces ge-0/0/0 mtu 9192
```

The supported range for configuring an MTU packet size is 256 through 9192.

Related Documentation

- MTU Default and Maximum Values on page 2479

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 111

Configuring VLAN Tagging

- Understanding Virtual LANs on page 2483
- VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices on page 2485
- Configuring VLAN Tagging on page 2485

Understanding Virtual LANs

A LAN is a single broadcast domain. When traffic is broadcast, all hosts within the LAN receive the broadcast traffic. A LAN is determined by the physical connectivity of devices within the domain.

Within a traditional LAN, hosts are connected by a hub or repeater that propagates any incoming traffic throughout the network. Each host and its connecting hubs or repeaters make up a LAN segment. LAN segments are connected through switches and bridges to form the broadcast domain of the LAN. Figure 113 shows a typical LAN topology.

Copyright © 2016, Juniper Networks, Inc.

Figure 113: Typical LAN

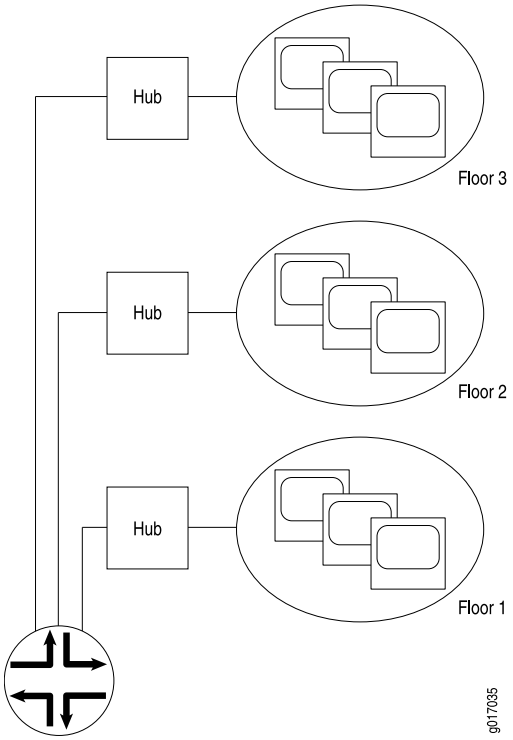

Virtual LANs (VLANs) allow network architects to segment LANs into different broadcast domains based on logical groupings. Because the groupings are logical, the broadcast domains are not determined by the physical connectivity of the devices in the network. Hosts can be grouped according to a logical function, to limit the traffic broadcast within the VLAN to only the devices for which the traffic is intended.

Suppose a corporate network has three major organizations: engineering, sales, and support. Using VLAN tagging, hosts within each organization can be tagged with a different VLAN identifier. Traffic sent to the broadcast domain is then checked against the VLAN identifier and broadcast to only the devices in the appropriate VLAN. Figure 114 shows a typical VLAN topology.

Figure 114: Typical VLAN

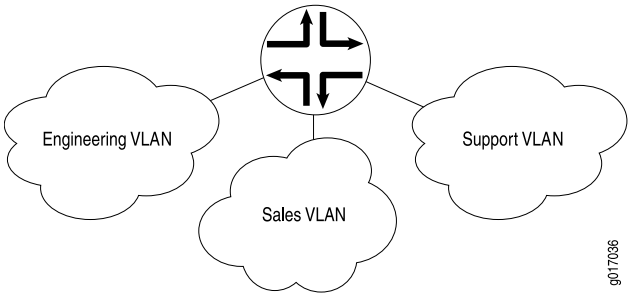

Related Documentation

- Understanding Interface Logical Properties on page 2451
- *MPLS Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices

Table 267 lists VLAN ID range by interface type supported on SRX Series devices:

Table 267: VLAN ID Range by Interface Type Supported on the SRX Series Devices

Interface Type	Interface Type VLAN ID Range
2-Port 10-Gigabit Ethernet	1 through 4094
10-Gigabit Ethernet	1 through 4094
16-Port Gigabit Ethernet	1 through 4094
24-Port Gigabit Ethernet	1 through 4094
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023

NOTE: On SRX320, and SRX340 devices, on 1-GE SFP Mini-PIM, the VLAN ID 4093 falls under the reserved VLAN address range. Because of this, you will not be able to configure VLAN ID from this range.

Related Documentation

- Understanding Interface Physical Properties on page 2475

Configuring VLAN Tagging

You can configure the branch SRX Series devices to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

See Table 268 for flexible VLANs.

Table 268: Flexible VLANs

Number of Tags	VLAN ID
0 (Untagged)	Native
1 (Tagged)	Single
2 (Dual tagged)	Dual

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Configuring Single-Tag Framing on page 2486
- Configuring Dual Tagging on page 2486
- Configuring Mixed Tagging on page 2486
- Configuring Mixed Tagging Support for Untagged Packets on page 2487

Configuring Single-Tag Framing

To configure a device to receive and forward single-tag frames with 802.1Q VLAN tags, include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
vlan-tagging;
```

NOTE: SRX Series high-end devices only support single-tag framing.

Configuring Dual Tagging

To configure the device to receive and forward dual-tag frames with 802.1Q VLAN tags, include the **flexible-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
flexible-vlan-tagging;
```

Configuring Mixed Tagging

Mixed tagging is supported on ethernet interfaces of all branch SRX Series devices. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
   vlan-id number;
   family family {
      address address;
   }
}
unit logical-unit-number {
   vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
   family family {
      address address;
   }
}
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
unit 0 {
   vlan-id 232;
   family inet {
      address 10.66.1.2/30;
   }
}
unit 1 {
   vlan-tags outer 0x8100.222 inner 0x8100.221;
   family inet {
      address 10.66.1.2/30;
   }
}
```

Configuring Mixed Tagging Support for Untagged Packets

You can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the **native-vlan-id** statement and the **flexible-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;
```

The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
   vlan-id 232;
```

Copyright © 2016, Juniper Networks, Inc.

```
    family inet {
      address 10.66.1.2/30;
    }
  }
  unit 1 {
    vlan-tags outer 0x8100.222 inner 0x8100.221;
    family inet {
      address 10.66.1.2/30;
    }
  }
```

Related Documentation

- Understanding Virtual LANs on page 2483

Copyright © 2016, Juniper Networks, Inc.

PART 35

Configuring DS1 Interfaces

- Configuring DS1 Interfaces on page 2491

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 112

Configuring DS1 Interfaces

- Understanding T1 and E1 Interfaces on page 2491
- Example: Configuring a T1 Interface on page 2494
- Example: Deleting a T1 Interface on page 2496

Understanding T1 and E1 Interfaces

T1 and E1 are equivalent digital data transmission formats that carry DS1 signals. T1 and E1 lines can be interconnected for international use.

This topic contains the following sections:

- T1 Overview on page 2491
- E1 Overview on page 2492
- T1 and E1 Signals on page 2492
- Encoding on page 2492
- T1 and E1 Framing on page 2493
- T1 and E1 Loopback Signals on page 2493

T1 Overview

T1 is a digital data transmission medium capable of handling 24 simultaneous connections running at a combined 1.544 Mbps. T1 combines these 24 separate connections, called channels or time slots, onto a single link. T1 is also called DS1.

The T1 data stream is broken into frames. Each frame consists of a single framing bit and 24 8-bit channels, totaling 193 bits per T1 frame. Frames are transmitted 8,000 times per second, at a data transmission rate of 1.544 Mbps (8,000 x 193 = 1.544 Mbps).

As each frame is received and processed, the data in each 8-bit channel is maintained with the channel data from previous frames, enabling T1 traffic to be separated into 24 separate flows across a single medium. For example, in the following set of 4-channel frames (without a framing bit), the data in channel 1 consists of the first octet of each frame, the data in channel 2 consists of the second octet of each frame, and so on:

```
           Chan. 1   Chan. 2   Chan. 3   Chan. 4
Frame 1  [10001100][00110001][11111000][10101010]
Frame 2  [11100101][01110110][10001000][11001010]
Frame 3  [00010100][00101111][11000001][00000001]
```

Copyright © 2016, Juniper Networks, Inc.

E1 Overview

E1 is the European format for DS1 digital transmission. E1 links are similar to T1 links except that they carry signals at 2.048 Mbps. Each signal has 32 channels, and each channel transmits at 64 Kbps. E1 links have higher bandwidth than T1 links because they use all 8 bits of a channel. T1 links use 1 bit in each channel for overhead.

T1 and E1 Signals

T1 and E1 interfaces consist of two pairs of wires—a transmit data pair and a receive data pair. Clock signals, which determine when the transmitted data is sampled, are embedded in the T1 and E1 transmissions.

Typical digital signals operate by sending either zeros (0s) or ones (1s), which are usually represented by the absence or presence of a voltage on the line. The receiving device need only detect the presence of the voltage on the line at the particular sampling edge to determine whether the signal is 0 or 1. T1 and E1, however, use bipolar electrical pulses. Signals are represented by no voltage (0), positive voltage (1), or negative voltage (1). The bipolar signal allows T1 and E1 receivers to detect error conditions in the line, depending on the type of encoding that is being used.

Encoding

The following are common T1 and E1 encoding techniques:

- Alternate mark inversion (AMI)—T1 and E1
- Bipolar with 8-zero substitution (B8ZS)—T1 only
- High-density bipolar 3 code (HDB3)—E1 only

AMI Encoding

AMI encoding forces the 1s signals on a T1 or E1 line to alternate between positive and negative voltages for each successive 1 transmission, as in this sample data transmission:

```
1 1 0 1 0 1 0 1
+ - 0 + 0 - 0 +
```

When AMI encoding is used, a data transmission with a long sequence of 0s has no voltage transitions on the line. In this situation, devices have difficulty maintaining clock synchronization, because they rely on the voltage fluctuations to constantly synchronize with the transmitting clock. To counter this effect, the number of consecutive 0s in a data stream is restricted to 15. This restriction is called the 1s density requirement, because it requires a certain number of 1s for every 15 0s that are transmitted.

On an AMI-encoded line, two consecutive pulses of the same polarity—either positive or negative—are called a bipolar violation (BPV), which is generally flagged as an error.

B8ZS and HDB3 Encoding

Neither B8ZS nor HDB3 encoding restricts the number of 0s that can be transmitted on a line. Instead, these encoding methods detect sequences of 0s and substitute bit patterns

Copyright © 2016, Juniper Networks, Inc.

for the sequences to provide the signal oscillations required to maintain timing on the link.

The B8ZS encoding method for T1 lines detects sequences of eight consecutive 0 transmissions and substitutes a pattern of two consecutive BPVs (11110000). Because the receiving end uses the same encoding, it detects the BPVs as 0s substitutions, and no BPV error is flagged. A single BPV, which does not match the 11110000 substitution bit sequence is likely to generate an error, depending on the configuration of the device.

The HDB3 encoding method for E1 lines detects sequences of four consecutive 0 transmissions and substitutes a single BPV (1100). Similar to B8ZS encoding, the receiving device detects the 0s substitutions and does not generate a BPV error.

T1 and E1 Framing

T1 interfaces uses extended superframe (ESF). E1 interfaces use G.704 framing or G.704 with no CRC4 framing, or can be in unframed mode.

ESF Framing for T1

ESF extends the D4 superframe from 12 frames to 24 frames. By expanding the size of the superframe, ESF increases the number of bits in the superframe framing pattern from 12 to 24. The extra bits are used for frame synchronization, error detection, and maintenance communications through the facilities data link (FDL).

The ESF pattern for synchronization bits is 001011. Only the framing bits from frames 4, 8, 12, 16, 20, and 24 in the superframe sequence are used to create the synchronization pattern.

The framing bits from frames 2, 6, 10, 14, 18, and 22 are used to pass a CRC code for each superframe block. The CRC code verifies the integrity of the received superframe and detects bit errors with a CRC6 algorithm.

The framing bits for frames 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, and 23 are used for the data link channel. These 12 bits enable the operators at the network control center to query the remote equipment for information about the performance of the link.

T1 and E1 Loopback Signals

The control signal on a T1 or E1 link is the loopback signal. Using the loopback signal, the operators at the network control center can force the device at the remote end of a link to retransmit its received signals back onto the transmit path. The transmitting device can then verify that the received signals match the transmitted signals, to perform end-to-end checking on the link.

Two loopback signals are used to perform the end-to-end testing:

- The loop-up command signal sets the link into loopback mode, with the following command pattern:

  ```
  ...100001000010000100...
  ```

- The loop-down signal returns the link to its normal mode, with the following command pattern:

Copyright © 2016, Juniper Networks, Inc.

```
...100100100100100100...
```

While the link is in loopback mode, the operator can insert test equipment onto the line to test its operation.

Related Documentation

- Example: Configuring a T1 Interface on page 2494

Example: Configuring a T1 Interface

This example shows how to complete the initial configuration on a T1 interface.

- Requirements on page 2494
- Overview on page 2494
- Configuration on page 2494
- Verification on page 2495

Requirements

Before you begin, install a PIM, connect the interface cables to the ports, and power on the device. See the *Getting Started Guide* for your device.

Overview

This example describes the initial configuration that you must complete on each network interface. In this example, you configure the t1-1/0/0 interface as follows:

- You create the basic configuration for the new interface by setting the encapsulation type to ppp. You can enter additional values for physical interface properties as needed.
- You set the logical interface to 0. Note that the logical unit number can range from 0 through 16,384. You can enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

```
set interfaces t1-1/0/0 encapsulation ppp unit 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a T1 interface:

1. Create the interface.

```
[edit]
user@host# edit interfaces t1-1/0/0
```


Copyright © 2016, Juniper Networks, Inc.

2. Create the basic configuration for the new interface.

```
[edit interfaces t1-1/0/0]
user@host# set encapsulation ppp
```

3. Add logical interfaces.

```
[edit interfaces t1-1/0/0]
user@host# set unit 0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show interfaces** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  ...
  t1-1/0/0 {
  encapsulation ppp;
  unit 0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying the Link State of All Interfaces on page 2495
- Verifying Interface Properties on page 2496

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the device are operational.

Action For each interface on the device:

1. In the J-Web interface, select **Troubleshoot>Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time, in milliseconds, is listed in the time field.

Copyright © 2016, Juniper Networks, Inc.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the operational mode, enter the **show interfaces detail** command.

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
  - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces t1-1/0/0] level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces> t1-1/0/0 page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. It indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of input and output bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics t1-1/0/0** command.

Related Documentation
- Understanding T1 and E1 Interfaces on page 2491
- Example: Deleting a T1 Interface on page 2496

Example: Deleting a T1 Interface

This example shows how to delete a T1 interface.

- Requirements on page 2496
- Overview on page 2496
- Configuration on page 2497
- Verification on page 2497

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the t1-1/0/0 interface.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web pages.

Configuration

Step-by-Step Procedure

To delete a T1 interface:

1. Specify the interface you want to delete.

   ```
   [edit interfaces]
   user@host# delete t1-1/0/0
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit interfaces]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- Understanding T1 and E1 Interfaces on page 2491
- Example: Configuring a T1 Interface on page 2494

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 36

Configuring DSL Interfaces

- Configuring VDSL2 Interfaces on page 2501

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 113

Configuring VDSL2 Interfaces

- VDSL2 Interface Technology Overview on page 2501
- VDSL2 Network Deployment Topology on page 2502
- VDSL2 Interface Support on SRX Series Devices on page 2503
- Example: Configuring VDSL2 Interfaces (Basic) on page 2507
- Example: Configuring VDSL2 Interfaces (Detail) on page 2513

VDSL2 Interface Technology Overview

Very-high-bit-rate digital subscriber line (VDSL) technology is part of the xDSL family of modem technologies that provide faster data transmission over a single flat untwisted or twisted pair of copper wires. The VDSL lines connect service provider networks and customer sites to provide high bandwidth applications (Triple Play services) such as high-speed Internet access, telephone services like voice over IP (VoIP), high-definition TV (HDTV), and interactive gaming services over a single connection.

VDSL2 is an enhancement to G.993.1 (VDSL) and permits the transmission of asymmetric (half-duplex) and symmetric (full-duplex) aggregate data rates up to 100 Mbps on short copper loops using a bandwidth up to 30 MHz. The VDSL2 technology is based on the ITU-T G.993.2 (VDSL2) standard which is the International Telecommunication Union standard describing a data transmission method for VDSL2 transceivers.

The VDSL2 uses discrete multitone (DMT) modulation. DMT is a method of separating a digital subscriber line signal so that the usable frequency range is separated into 256 frequency bands (or channels) of 4.3125 KHz each. The DMT uses the Fast Fourier Transform (FFT) algorithm for demodulation or modulation for increased speed.

VDSL2 interface supports Packet Transfer Mode (PTM). The PTM mode transports packets (IP, PPP, Ethernet, MPLS, and so on) over DSL links as an alternative to using Asynchronous Transfer Mode (ATM). PTM is based on the Ethernet in the First Mile (EFM) IEEE802.3ah standard.

VDSL2 provides backward compatibility with ADSL, ADSL2, and ADSL2+ because this technology is based on both the VDSL1-DMT and ADSL2/ADSL2+ recommendations.

Related Documentation

- VDSL2 Network Deployment Topology on page 2502
- VDSL2 Interface Support on SRX Series Devices on page 2503

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring VDSL2 Interfaces (Basic) on page 2507
- Example: Configuring VDSL2 Interfaces (Detail) on page 2513

VDSL2 Network Deployment Topology

In standard telephone cables of copper wires, voice signals use only a fraction of the available bandwidth. Like any other DSL technology, the VDSL2 technology utilizes the remaining capacity to carry the data and multimedia on the wire without interrupting the line's ability to carry voice signals.

This example depicts the typical VDSL2 network topology deployed using SRX Series Services Gateways.

A VDSL2 link between network devices is set up as follows:

1. Connect an end-user device such as a LAN, hub, or PC through an Ethernet interface to the customer premises equipment (CPE) (for example, an SRX Series device).
2. Connect the CPE to a DSLAM.
3. The VDSL2 interface uses either Gigabit Ethernet or fiber as second mile to connect to the Broadband Remote Access Server (B-RAS) as shown in Figure 115.
4. The ADSL interface uses either Gigabit Ethernet (in case of IP DSLAM] as the “second mile” to connect to the B-RAS or OC3/DS3 ATM as the second mile to connect the B-RAS as shown in Figure 116.

   NOTE: **The VDSL2 technology is backward compatible with ADSL. VDSL2 provides an ADSL interface in an ATM DSLAM topology and provides a VDSL2 interface in an IP or VDSL DSLAM topology.**

   The DSLAM accepts connections from many customers and aggregates them to a single, high-capacity connection to the Internet.

Figure 115 shows a typical VDSL2 network topology.

Copyright © 2016, Juniper Networks, Inc.

Figure 115: Typical VDSL2 End-to-End Connectivity and Topology Diagram

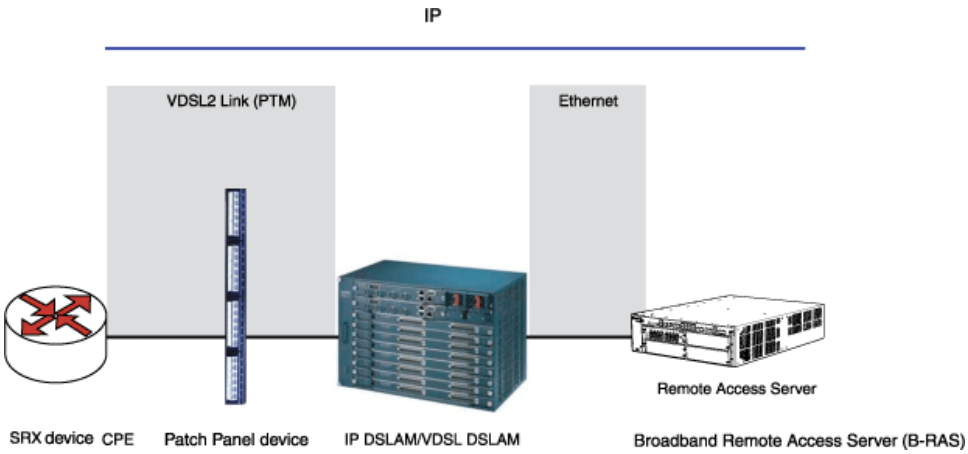

Figure 116 shows a backward-compatible ADSL topology using ATM DSLAM.

Figure 116: Backward-Compatible ADSL Topology (ATM DSLAM)

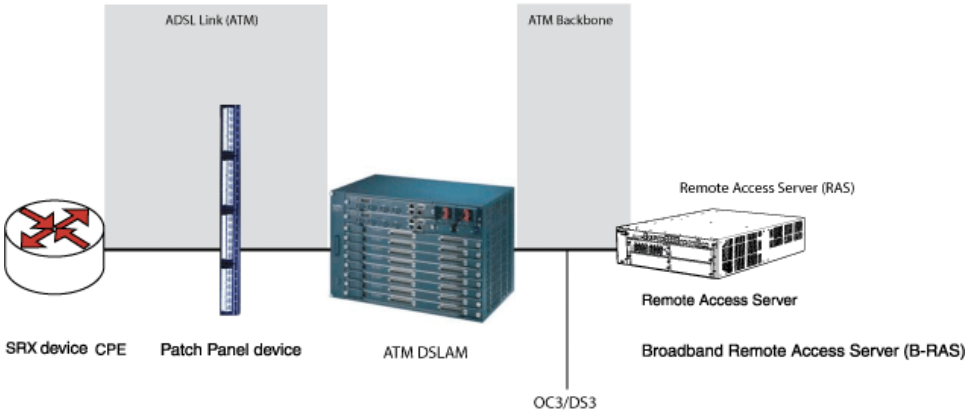

Related Documentation

- VDSL2 Interface Technology Overview on page 2501
- VDSL2 Interface Support on SRX Series Devices on page 2503
- Example: Configuring VDSL2 Interfaces (Basic) on page 2507
- Example: Configuring VDSL2 Interfaces (Detail) on page 2513

VDSL2 Interface Support on SRX Series Devices

The VDSL2 interface is supported on the following SRX Series devices as shown in Table 269.

Copyright © 2016, Juniper Networks, Inc.

Table 269: VDSL2 Annex A and Annex B Features

Features	POTS	ISDN
Devices	Integrated VDSL Module (SRX320-POTS) VDSL Mini-PIM (SRX320, SRX340)	Integrated VDSL Module (SRX320-ISDN)
Supported Annex Operating Modes	Annex A and Annex B*	Annex B
Supported Bandplans	997/998	998
Supported standards	ITU-T G.993.2 (VDSL2)	ITU-T G.993.2 (VDSL2)
Used in	North American network implementations	European network implementations
ADSL backward compatibility	ADSL G992.5-A (ADSL Annex A)	ADSL G992.5-B (ADSL Annex B)

* Annex B support is not available on VDSL2 Mini-PIMs.

VDSL2 Interface Compatibility with ADSL Interfaces

VDSL2 interfaces on SRX Series devices are backward compatible with most ADSL interface standards. The VDSL2 interface uses Ethernet in the First Mile (EFM) mode or Packet Transfer Mode (PTM) and uses the named interface **pt-1/0/0**. In ADSL fallback mode, VDSL2 operates on the ATM encapsulation interface in the first mile and uses the named interface **at-1/0/0**.

NOTE:

- The VDSL2 interface has backward compatibility with ADSL/ADSL2/ADSL2+. The VDSL2 interface is represented by the pt interface when configured to function as VDSL2, and the ADSL interface is represented by the at interface when configured to function as ADSL.
- On VDSL2 interfaces, by default the pt-1/0/0 interface is created when there is no configuration already created for either the pt-1/0/0 or the at-1/0/0 interface.

NOTE: It requires around 60 seconds to switch from VDSL2 to ADSL or from ADSL to VDSL2 operating modes.

Table 270 lists VDSL2 operating modes and their backward compatibility with ADSL interface standards.

Copyright © 2016, Juniper Networks, Inc.

Table 270: VDLS2 Operating Mode Backward Compatibility with ADSL

VDSL2 Annex Type	Operating Modes	Description
VDSL2 Annex A interface (ADSL modes for Annex A only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface uses either ANSI T1.413 Issue II mode or ITU G.992.1 mode. NOTE: Auto operating mode does not work when DSLAM located at the central office is operating at ADSL2+ Annex M mode.
	ansi-dmt	Configures the ADSL interface to use ANSI T1.413 Issue II mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
VDSL2 Annex B interface (ADSL modes for Annex B only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex B, the ADSL interface trains in ITU G.992.1 mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
	itu-annexb-ur2	Configures the ADSL line to use G.992.1 Deutsche Telekom UR-2 mode.

VDSL2 Interfaces Supported Profiles

A profile is a table that contains a list of preconfigured VDSL2 settings. Table 271 lists the different profiles supported on the VDSL2 interfaces and their properties.

Copyright © 2016, Juniper Networks, Inc.

Table 271: Supported Profiles on the VDSL2 Interfaces

Profiles	Data Rate
8a	50
8b	50
8c	50
8d	50
12a	68
12b	68
17a	100
Auto	Negotiated (based on operating mode)

VDSL2 Interfaces Supported Features

The following features are supported on the VDSL2 interfaces:

- ADSL/ADSL2/ADSL2+ backward compatibility with Annex A, Annex M support
- PTM or EFM (802.3ah) support
- Operation, Administration, and Maintenance (OAM) support for ADSL/ADSL2/ADSL2+ mode
- ATM quality of service (QoS) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- Multilink Point-to-Point Protocol (MLPPP) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- MTU size of 1514 bytes (maximum) in VDSL2 mode and 1496 bytes in ADSL mode.
- Support for maximum of 10 permanent virtual connections (PVCs) (only in ADSL/ADSL2/ADSL2+ mode)
- Dying gasp support (ADSL and VDSL2 mode)

NOTE: **On SRX320 devices with VDLS2, ATM COS VBR-related functionality cannot be tested.**

Related Documentation

- VDSL2 Interface Technology Overview on page 2501
- VDSL2 Network Deployment Topology on page 2502
- Example: Configuring VDSL2 Interfaces (Basic) on page 2507

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring VDSL2 Interfaces (Detail) on page 2513

Example: Configuring VDSL2 Interfaces (Basic)

This example shows how to configure the VDSL2 interfaces for SRX320, and SRX340 devices.

- Requirements on page 2507
- Overview on page 2507
- Configuration on page 2507
- Verifying the Configuration on page 2509

Requirements

Before you begin:

- Establish basic connectivity. See the *Quick Start Guide* for your device for factory default settings.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 2546.

Overview

In this example, you create a VDSL2 interface called **pt-1/0/0**, specify the type of encapsulation, and set the VDSL2 profile to auto.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pt-1/0/0 vdsl-options vdsl-profile auto
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
set interfaces pt-1/0/0 unit 0 family inet dhcp
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure

To configure the VDSL2 interfaces for the SRX320, and SRX340 devices and enable VLAN tagging:

1. Create an interface.

   ```
   [edit]
   user@host# edit interfaces pt-1/0/0
   ```

2. Set the type of VDSL2 profile.

   ```
   [edit interfaces pt-1/0/0]
   user@host# set vdsl-options vdsl-profile auto
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Specify the logical unit to connect to this physical VDSL2 interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0
```

4. Specify the family protocol type.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 family inet
```

5. To enable the DHCP client on the interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 family inet dhcp
```

6. Specify the type of encapsulation on the VDSL2 logical interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 encapsulation ppp-over-ether
```

NOTE: The VDSL2 interface supports PPPoE. You can also set no encapsulation for the VDSL2 interface.

NOTE: To configure VLAN tagging, continue the configuration with the next step.

7. To enable VLAN tagging on the pt interface.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 vlan-tagging
```

8. Specify the value of the VLAN ID to be configured.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 unit 0 vlan-id 100
```

NOTE: This feature is supported only on the pt interface, and the range of VLANs that can be configured is 0 to 4093.

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
   user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile auto;
}
unit 0 {
   encapsulation ppp-over-ether;
```


Copyright © 2016, Juniper Networks, Inc.

```
    Family inet {
      address 100.100.100.1/24;
      dhcp;
  }
```

NOTE: When VLAN tagging is configured, the intended output is:

```
[edit]
  user@host# show interfaces pt-1/0/0
  vlan-tagging;
vdsl-options {
vdsl-profile auto;
}
unit 0 {
  encapsulation ppp-over-ether;
  vlan-id 100;
    Family inet {
    address 100.100.100.1/24;
    dhcp;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verifying the Configuration

Confirm that the configuration is working properly.

- Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex A Mode) on page 2509
- Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex B Mode) on page 2512

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex A Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 524, Generation: 149
 Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed:    45440kbps

 Speed: VDSL2
 Device flags : Present Running
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:b1:7e:85:84:ff
 Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
 Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)
 Traffic statistics:
  Input bytes : 22438962 97070256 bps
```

Copyright © 2016, Juniper Networks, Inc.

```
  Output bytes : 10866024 43334088 bps
  Input packets: 15141 8187 pps
  Output packets: 7332 3655 pps
Input errors:
   Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
   L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
 Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU  errors: 0,
  Resource errors: 0
 Egress queues: 8 supported, 4 in use
 Queue counters: Queued packets Transmitted packets Dropped packets
    0 best-effort 6759 6760 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
    3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
   LOF 0 0 OK
   LOS 0 0 OK
   LOM 0 0 OK
   LOP 0 0 OK
   LOCDI 0 0 OK
   LOCDNI 0 0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Profile-17a Annex A
  Last fail code: None
  Subfunction : 0x00
  Seconds in showtime : 45171
VDSL Chipset Information: VTU-R VTU-C
 Vendor Country : 0xb5 0xb5
  Vendor ID : BDCM BDCM
  Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
  Attenuation (dB) : 0.0 0.0
  Capacity used (%) : 0 0
  Noise margin (dB) : 20.0 20.0
  Output power (dBm) : 6.0 12.0
                     Interleave Fast Interleave Fast
  Bit rate (kbps) :     100004     0    45440       0
  CRC :                 0          0    0           0
 FEC :                  0          0    0            0
 HEC :                  0          0    0            0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
   Direction : Output
   CoS transmit queue Bandwidth Buffer Priority
Limit
                        %        bps        %       usec
   0 best-effort        95       43168000   95       0      low
none
   3 network-control    5        2272000    5        0       low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
      Input bytes :  23789064
      Output bytes :  10866024
      Input packets:   16052
```


Copyright © 2016, Juniper Networks, Inc.

```
    Output packets:   7332
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes :  23789064  97070256 bps
    Output bytes :  10866024  43334088 bps
    Input packets:  16052  8187 pps
     Output packets:  7332   3655 pps
   Security: Zone: Null
   Flow Statistics :
  Flow Input statistics :
    Self packets :                     0
    ICMP packets :                     0
    VPN packets :                      0
    Multicast packets :                0
    Bytes permitted by policy :        0
    Connections established :          0
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1482, Generation: 169, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary

   Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast:  10.10.10.255,
Generation: 158
```

The output shows a summary of VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).

Copyright © 2016, Juniper Networks, Inc.

- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex A indicates VDSL profile of the DSLAM connected at other end.

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex B Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 536, Generation: 238
  Type: PTM, Link-level type: Ethernet, MTU: 1514, VDSL mode, Speed: 45439kbps
  Speed: VDSL2
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:df:20
  Last flapped   : 2011-05-13 07:34:33 PDT (00:46:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors:
 0, L2 mismatch timeouts: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
 Resource errors: 0
  VDSL  alarms   : None
  VDSL  defects  : None
  VDSL media:            Seconds        Count  State
    LOF                      177            0  OK
    LOS                      177            0  OK
    LOM                        0            0  OK
    LOP                        0            0  OK
    LOCDI                      0            0  OK
    LOCDNI                   177            0  OK
  VDSL status:
    Modem status  : Showtime (Profile-17a)
    VDSL profile  :     Auto  Annex B
    Last fail code: None
    Subfunction   : 0x00
   Seconds in showtime : 2794  VDSL Chipset Information:              VTU-R
             VTU-C
    Vendor Country :                    0xb5                  0xb5
    Vendor ID      :                    BDCM                  BDCM
    Vendor Specific:                  0x9385                0x9395
  VDSL Statistics:                     VTU-R                 VTU-C
    Attenuation (dB)         :          0.0                   0.0
    Capacity used (%)        :            0                     0
    Noise margin (dB)        :         18.5                   9.5
    Output power (dBm)       :         14.5                   3.0
```


Copyright © 2016, Juniper Networks, Inc.

```
                                    Interleave       Fast  Interleave       Fast
    Bit rate (kbps)   :                 100015          0       45439          0
    CRC               :                      0          0           0          0
    FEC               :                      0          0           0          0
    HEC               :                      0          0           0          0
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %           bps     %           usec
    0 best-effort            95      43167050    95              0      low
none
    3 network-control         5       2271950     5              0      low
none
```

The output shows a summary of the VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).
- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex B indicates the VDSL profile of the DSLAM connected at other end.

Related Documentation

- Understanding Interfaces on page 2441
- VDSL2 Interface Technology Overview on page 2501
- Example: Configuring VDSL2 Interfaces (Detail) on page 2513
- *Example: Configuring VDSL2 Interfaces in ADSL Mode (Detail)*
- *Example: Configuring the Device as a DHCP Client*

Example: Configuring VDSL2 Interfaces (Detail)

This example shows how to configure VDSL2 interfaces on SRX Series Services Gateways.

This example uses VDSL2 Mini-PIM installed on SRX320 devices. The information is also applicable to SRX320 devices (with VDSL2 Mini-PIMs).

- Requirements on page 2513
- Overview on page 2514
- Configuration on page 2515
- Verification on page 2526

Requirements

Before you begin:

- Install Junos OS Release 10.1 or later on the SRX Series devices.
- Establish basic connectivity and set up and perform initial configuration. See the *Quick Start Guide* for your device for factory default settings.
- Install the VDSL2 Mini-PIM on the SRX320 device chassis.
- Connect the SRX320 device to a DSLAM.
- On VDSL2 Mini-PIMs, by default the **pt-1/0/0** interface is created when there is no configuration already created for either the **pt-1/0/0** or the **at-1/0/0** interface. You can switch to ADSL mode by just configuring **at-1/0/0**. If the configurations are already created for **pt-1/0/0** or **at-1/0/0**, then you need to deactivate **pt-1/0/0** before you create **at-1/0/0** or deactivate **at-1/0/0** to create **pt-1/0/0**.
- Make sure that you have deleted the previous configurations on **pt-1/0/0** and **pp0**.

Overview

This example uses SRX320 devices. The information is also applicable to SRX340 devices.

Figure 117 shows typical SRX Series devices with VDSL2 Mini-PIM network connections.

Figure 117: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario

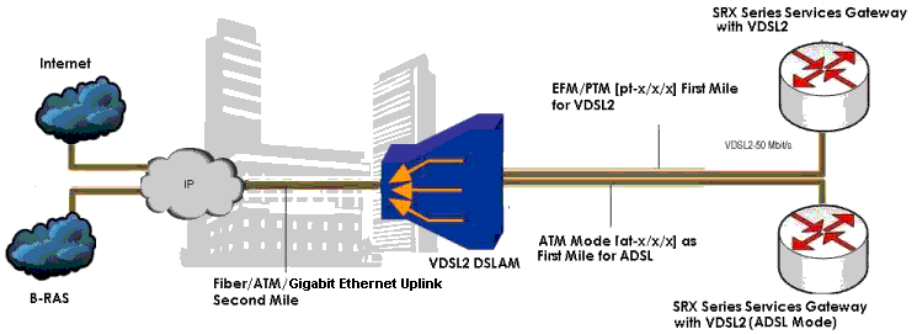

In this example, you begin a new configuration on a VDSL2 Mini-PIM. You first deactivate previous interfaces and delete any old configuration from the device. Then you set the interfaces with the VDSL profile and the Layer 3 configuration for the end-to-end data path.

You then configure the PPPoE on the pt-1/0/0 interface with a static IP address or CHAP authentication. You configure PPPoE on the pt-1/0/0 interface with unnumbered IP address (PAP authentication or CHAP authentication).

Finally, you configure PPPoE on the pt-1/0/0 interface with negotiated IP address (PAP authentication or CHAP authentication).

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Beginning a New Configuration on a VDSL2 Mini-PIM on page 2515
- Configuring the VDSL2 Mini-PIM for End-to-End Data Path on page 2516
- Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 2517
- Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication) on page 2519
- Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication) on page 2520
- Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication) on page 2522
- Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication) on page 2523
- Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication) on page 2525

Beginning a New Configuration on a VDSL2 Mini-PIM

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
deactivate interface pt-1/0/0
deactivate interface at-1/0/0
delete interface pt-1/0/0
delete interface pp0
```

Step-by-Step Procedure

To begin a new configuration on a VDSL2 Mini-PIM:

1. Deactivate any previous interfaces.

   ```
   [edit]
   user@host# deactivate interface pt-1/0/0
   user@host# deactivate interface at-1/0/0
   ```

2. Delete any old configurations.

   ```
   [edit]
   user@host# delete interface pt-1/0/0
   user@host# delete interface pp0
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show chassis fpc** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# run show chassis fpc
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online -------------------- CPU less FPC --------------------
1 Online -------------------- CPU less FPC --------------------
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the VDSL2 Mini-PIM for End-to-End Data Path

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24
```

Step-by-Step Procedure To configure the VDSL2 Mini-PIM for end-to-end data path:

1. Configure the interfaces with the VDSL profile and the Layer 3 configuration for end-to-end data path.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
family inet {
  address 11.11.11.1/24;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name
  locky local-password india passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
user@host# set access profile pap_prof authentication-order password client cuttack
  pap-password india
```

NOTE: To configure VLAN tagging while configuring PPPoE on the pt-1/0/0 interface with

- Static IP address
- Static IP address (CHAP authentication)
- Unnumbered IP address (PAP Authentication)
- Unnumbered IP address (CHAP Authentication)
- Negotiated IP address (PAP Authentication)
- Negotiated IP address (CHAP Authentication)

the following commands must be included at **[edit]** hierarchy level:

```
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure

To configure the PPPoE on the pt-1/0/0 interface with a static IP address:

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the PPP options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
   user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
   user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
   user@host# set interfaces pp0 unit 0 ppp-options pap passive
   ```

3. Configure the PPPoE options for the interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

5. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0** and **show access profile pap_prof** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pp0
unit 0 {
ppp-options {
  pap {
access-profile pap_prof;
local-name locky;
  local-password "$ABC123"; ## SECRET-DATA
passive;
}
  }
pppoe-options {
underlying-interface pt-1/0/0.0;
  auto-reconnect 120;
client;
}
  family inet {
address 10.1.1.6/24;
}
}
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
  user@host# show access profile pap_prof
authentication-order password;
client cuttack pap-password "$ABC123"; ## SECRET-DATA
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
  local-name locky passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

Step-by-Step Procedure

To configure the PPPoE on the pt-1/0/0 interface with a static IP address (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the PPP options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
   user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
   user@host# set interfaces pp0 unit 0 ppp-options chap passive
   ```

3. Configure the PPPoE options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
   user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
   user@host# set interfaces pp0 unit 0 pppoe-options client
   ```

4. Configure the IP address for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** and **show interfaces pp0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
  user@host# show interfaces pp0
unit 0 {
ppp-options {
  chap {
default-chap-secret "$ABC123"; ## SECRET-DATA
local-name locky;
  passive;
}
}
  pppoe-options {
underlying-interface pt-1/0/0.0;
auto-reconnect 120;
  client;
}
family inet {
  address 10.1.1.6/24;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name
  locky local-password india passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination
  10.1.1.1
user@host# set access profile pap_prof authentication-order password client cuttack
  pap-password india
```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the IP address for the interface.

   ```
   [edit]
   user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
   ```

3. Configure the PPP options for the interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

4. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

5. Configure the unnumbered address and destination for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination
  10.1.1.1
```

6. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lo0**, **show interfaces pt-1/0/0**, and **show interfaces pp0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces lo0
unit 0 {
family inet {
  address 10.1.1.24/32;
  }
}
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
  user@host# show interfaces pp0
unit 0 {
ppp-options {
  pap {
access-profile pap_prof;
local-name locky;
  local-password "$ABC123"; ## SECRET-DATA
passive;
}
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
pppoe-options {
underlying-interface pt-1/0/0.0;
  auto-reconnect 120;
client;
}
  family inet {
unnumbered-address lo0.0 destination 10.1.1.1;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
  local-name locky passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination
  10.1.1.1
```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the IP address for the interface.

   ```
   [edit]
   user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
   ```

3. Configure the PPP options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
   user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
   user@host# set interfaces pp0 unit 0 ppp-options chap passive
   ```

4. Configure the PPPoE options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
   user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
   user@host# set interfaces pp0 unit 0 pppoe-options client
   ```

5. Configure the unnumbered address and destination for the interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination
  10.1.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0**, and **show interfaces lo0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pp0
unit 0 {
ppp-options {
  chap {
default-chap-secret "$ABC123"; ## SECRET-DATA
local-name locky;
  passive;
}
}
  pppoe-options {
underlying-interface pt-1/0/0.0;
auto-reconnect 120;
  client;
}
family inet {
  unnumbered-address lo0.0 destination 10.1.1.1;
  }
}
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
  user@host# show interfaces lo0
unit 0 {
family inet {
  address 10.1.1.24/32;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile my_prf local-name
  purple local-password <password> passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet negotiate-address
user@host# set access profile my_prf authentication-order password
user@host# set access profile my_prf
```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the PPP options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 ppp-options pap access-profile my_prf
   user@host# set interfaces pp0 unit 0 ppp-options pap local-name purple
   user@host# set interfaces pp0 unit 0 ppp-options pap local-password <password>
   user@host# set interfaces pp0 unit 0 ppp-options pap passive
   ```

3. Configure the PPPoE options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
   user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
   user@host# set interfaces pp0 unit 0 pppoe-options client
   ```

4. Configure the negotiated IP address for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 family inet negotiate-address
   ```

5. Configure the access profile for the interface.

   ```
   [edit]
   user@host# set access profile my_prf authentication-order password
   user@host# set access profile my_prf
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0**, **show interfaces pp0**, and **show access profile my_prf** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
```


Copyright © 2016, Juniper Networks, Inc.

```
    user@host# show interfaces pp0
unit 0 {
ppp-options {
    pap {
access-profile my_prf;
local-name purple;
    local-password "$ABC123"; ## SECRET-DATA
passive;
}
    }
pppoe-options {
underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
client;
}
    family inet {
negotiate-address;
}
}
[edit]
    user@host# show access profile my_prf
authentication-order password;
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret <password>
    local-name purple passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
    auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet negotiate-address
```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

   ```
   [edit]
   user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
   user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
   ```

2. Configure the PPP options for the interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret
     <password>
   user@host# set interfaces pp0 unit 0 ppp-options chap local-name purple
   user@host# set interfaces pp0 unit 0 ppp-options chap passive
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the negotiated IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet negotiate-address
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** and **show interfaces pt-1/0/0**commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pp0
unit 0 {
ppp-options {
  chap {
default-chap-secret "$ABC123"; ## SECRET-DATA
local-name purple;
  passive;
}
}
  pppoe-options {
underlying-interface pt-1/0/0.0;
auto-reconnect 120;
  client;
}
family inet {
  negotiate-address;
  }
}
[edit]
  user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
  }
unit 0 {
encapsulation ppp-over-ether;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying the Configuration on page 2527
- Verifying the VDSL2 Mini-PIM for End-to-End Data Path on page 2529
- Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 2532

Copyright © 2016, Juniper Networks, Inc.

- Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication) on page 2533
- Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (PAP Authentication) on page 2534
- Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication) on page 2535
- Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication) on page 2536
- Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication) on page 2537

Verifying the Configuration

Purpose Verify the FPC status and the command output.

Action

1. Verify the FPC status by entering the **show chassis fpc** command. The output should display FPC status as online.

```
user@host# run show chassis fpc
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online -------------------- CPU less FPC --------------------
1 Online -------------------- CPU less FPC --------------------
```

NOTE: The VDSL2 Mini-PIM is installed in the first slot of the SRX320 device chassis; therefore, the FPC used here is fpc 1. For SRX340 devices, the FPC used will be fpc 1, fpc 2, fpc 3, or fpc 4.

2. Enter **run show interface pt-1/0/0** and verify the following information in the command output:
   - Status of interface pt-1/0/0 is displayed as physical link is up.
   - Modem status is displayed as Showtime (Profile-17a).
   - Time in seconds during which the interface stayed up is displayed as Seconds in Showtime.
   - VDSL profile of DSLAM is displayed as Auto Annex A.

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 524, Generation: 149
 Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed:    45440kbps

 Speed: VDSL2
 Device flags : Present Running
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:b1:7e:85:84:ff
 Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
```

Copyright © 2016, Juniper Networks, Inc.

```
 Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)
 Traffic statistics:
  Input bytes : 22438962 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 15141 8187 pps
  Output packets: 7332 3655 pps
Input errors:
   Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
   L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
 Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU  errors:
0,
  Resource errors: 0
 Egress queues: 8 supported, 4 in use
 Queue counters: Queued packets Transmitted packets Dropped packets
    0 best-effort 6759 6760 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
    3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
   LOF 0 0 OK
   LOS 0 0 OK
   LOM 0 0 OK
   LOP 0 0 OK
   LOCDI 0 0 OK
   LOCDNI 0 0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Profile-17a Annex A
  Last fail code: None
  Subfunction : 0x00
  Seconds in showtime : 45171
VDSL Chipset Information: VTU-R VTU-C
 Vendor Country : 0xb5 0xb5
  Vendor ID : BDCM BDCM
  Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
  Attenuation (dB) : 0.0 0.0
  Capacity used (%) : 0 0
  Noise margin (dB) : 20.0 20.0
  Output power (dBm) : 6.0 12.0
                     Interleave Fast Interleave Fast
  Bit rate (kbps) :     100004     0    45440       0
  CRC :                 0          0    0           0
 FEC :                  0          0    0            0
 HEC :                  0          0    0            0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
   Direction : Output
   CoS transmit queue Bandwidth Buffer Priority
Limit
                       %        bps        %       usec
   0 best-effort        95       43168000   95       0      low
none
   3 network-control    5        2272000    5        0       low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
    Flags: SNMP-Traps Encapsulation: ENET2
```


Copyright © 2016, Juniper Networks, Inc.

```
  Traffic statistics:
    Input bytes :  23789064
    Output bytes :  10866024
    Input packets:   16052
    Output packets:   7332
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes :  23789064  97070256 bps
    Output bytes :  10866024  43334088 bps
    Input packets:  16052  8187 pps
     Output packets:  7332   3655 pps
   Security: Zone: Null
   Flow Statistics :
  Flow Input statistics :
    Self packets :                     0
    ICMP packets :                     0
    VPN packets :                      0
    Multicast packets :                0
    Bytes permitted by policy :        0
    Connections established :          0
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1482, Generation: 169, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary

     Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast:  10.10.10.255,
Generation: 158
```

Verifying the VDSL2 Mini-PIM for End-to-End Data Path

Purpose Verify the interface status and check traffic statistics.

Copyright © 2016, Juniper Networks, Inc.

Action 1. Verify interface status by using the **show interface terse** command and test end-to-end data path connectivity by sending the ping packets to the remote end IP address.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up   inet     11.11.11.1/24

[edit]
user@host# run ping 11.11.11.2 count 1000 rapid
PING 11.11.11.2 (11.11.11.2): 56 data bytes
!!!
```

```
  VDSL status:
    Modem status  : Showtime (Profile-17a)
    VDSL profile  : Profile-17a Annex A
    Last fail code: None
    Subfunction   : 0x00
    Seconds in showtime : 723
  VDSL Chipset Information:              VTU-R                  VTU-C
    Vendor Country :                      0xb5                   0xb5
    Vendor ID      :                      BDCM                   BDCM
    Vendor Specific:                    0x9385                 0x9385
  VDSL Statistics:                      VTU-R                  VTU-C
    Attenuation (dB)         :           0.0                    0.0
    Capacity used (%)        :             0                      0
    Noise margin (dB)        :          16.0                   20.0
    Output power (dBm)       :           5.0                   13.0

                                Interleave       Fast  Interleave       Fast

    Bit rate (kbps)   :            100004          0       45440
0
    CRC               :                 0          0           0
0
    FEC               :                 0          0           0
0
    HEC               :                 0          0           0
0
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth              Buffer Priority
  Limit
                              %           bps     %           usec
    0 best-effort            95       43168000    95              0      low
   none
    3 network-control         5        2272000     5              0      low
   none

  Logical interface pt-1/0/0.0 (Index 72) (SNMP ifIndex 521) (Generation 158)

    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               84000
     Output bytes  :               98000
     Input  packets:                1000
     Output packets:                1000
    Local statistics:
     Input  bytes  :               84000
     Output bytes  :               98000
     Input  packets:                1000
     Output packets:                1000
    Transit statistics:
     Input  bytes  :                   0                  0 bps
     Output bytes  :                   0                  0 bps
     Input  packets:                   0                  0 pps
     Output packets:                   0                  0 pps
    Security: Zone: Null
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
```

```
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        0
      Connections established :          0
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        0
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1482, Generation: 169, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 11.11.11/24, Local: 11.11.11.1, Broadcast: 11.11.11.255,
 Generation: 189
```

Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address

Purpose Verify the interface output and the end-to-end data path.

Action 1. Verify the interface output.

```
user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 510
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
    Input packets : 0
    Output packets: 0

  Logical interface pp0.0 (Index 71) (SNMP ifIndex 522)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 57
```

 Copyright © 2016, Juniper Networks, Inc.

```
    Output packets: 56
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 22 (00:00:40 ago), Output: 25 (00:00:04 ago)
  LCP state: Down
  NCP state: inet: Down, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.6
```

2. Verify the end-to-end data path on the interface.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.6/24

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 12
    Output packets: 10
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1 (00:00:08 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Success
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.6
```

2. Verify the interface and check the end-to-end data path connectivity.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.6/24

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes

!!!
```

```
      State: SessionUp, Session ID: 33,
      Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 22
    Output packets: 20
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1 (00:00:08 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Success
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1.1, Local: 10.1.1.24
```

2. Verify the end-to-end data path testing.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.24           --> 10.1.1.1

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 35,
      Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 25
    Output packets: 22
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 2 (00:00:10 ago), Output: 2 (00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Success
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1.1, Local: 10.1.1.24
```

2. Verify the end-to-end data path testing on the PPPoE interface.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.24           --> 10.1.1.1

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
  Logical interface pp0.0 (Index 72) (SNMP ifIndex 522)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 4,
      Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 18
    Output packets: 18
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 11 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Success
    Security: Zone: Null
    Protocol inet, MTU: 1474
      Flags: Negotiate-Address
      Addresses, Flags: Kernel Is-Preferred Is-Primary
        Destination: 12.12.12.1, Local: 12.12.12.11
```

2. Verify the end-to-end data path connectivity.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     12.12.12.11         --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!
```

```
    Input packets : 0
    Output packets: 0

  Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 8,
      Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 12
    Output packets: 11
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 4 (00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Success
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1474
      Flags: Negotiate-Address
      Addresses, Flags: Kernel Is-Preferred Is-Primary
        Destination: 12.12.12.1, Local: 12.12.12.12
```

2. Verify the end-to-end data path connectivity.

```
user@host# run show interfaces pt-1/0/0 terse
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     12.12.12.12         --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!
```

PART 37

Configuring Ethernet Interfaces

- Performing Initial Configuration on Ethernet Interfaces on page 2541
- Configuring Aggregated Ethernet Interfaces on page 2557
- Configuring Link Aggregation Control Protocol on page 2571
- Configuring Gigabit Ethernet Physical Interface Modules on page 2587
- Configuring Ethernet OAM Link Fault Management on page 2603
- Configuring Power over Ethernet on page 2611

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 114

Performing Initial Configuration on Ethernet Interfaces

- Understanding Ethernet Interfaces on page 2541
- Understanding Static ARP Entries on Ethernet Interfaces on page 2545
- Understanding Promiscuous Mode on Ethernet Interface on page 2545
- Example: Creating an Ethernet Interface on page 2546
- Example: Deleting an Ethernet Interface on page 2547
- Example: Configuring Static ARP Entries on Ethernet Interfaces on page 2548
- Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 2551
- Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 2552

Understanding Ethernet Interfaces

Ethernet is a Layer 2 technology that operates in a shared bus topology. Ethernet supports broadcast transmission, uses best-effort delivery, and has distributed access control. Ethernet is a point-to-multipoint technology.

In a shared bus topology, all devices connect to a single, shared physical link through which all data transmissions are sent. All traffic is broadcast so that all devices within the topology receive every transmission. The devices within a single Ethernet topology make up a broadcast domain.

Ethernet uses best-effort delivery to broadcast traffic. The physical hardware provides no information to the sender about whether the traffic was received. If the receiving host is offline, traffic to the host is lost. Although the Ethernet data link protocol does not inform the sender about lost packets, higher layer protocols such as TCP/IP might provide this type of notification.

This topic contains the following sections:

- Ethernet Access Control and Transmission on page 2542
- Collisions and Detection on page 2542
- Collision Domains and LAN Segments on page 2543

Copyright © 2016, Juniper Networks, Inc.

- Broadcast Domains on page 2544
- Ethernet Frames on page 2544

Ethernet Access Control and Transmission

Ethernet's access control is distributed because Ethernet has no central mechanism that grants access to the physical medium within the network. Instead, Ethernet uses carrier-sense multiple access with collision detection (CSMA/CD). Because multiple devices on an Ethernet network can access the physical medium, or wire, simultaneously, each device must determine whether the physical medium is in use. Each host listens on the wire to determine if a message is being transmitted. If it detects no transmission, the host begins transmitting its own data.

The length of each transmission is determined by fixed Ethernet packet sizes. By fixing the length of each transmission and enforcing a minimum idle time between transmissions, Ethernet ensures that no pair of communicating devices on the network can monopolize the wire and block others from sending and receiving traffic.

Collisions and Detection

When a device on an Ethernet network begins transmitting data, the data takes a finite amount of time to reach all hosts on the network. Because of this delay, or latency, in transmitting traffic, a device might detect an idle state on the wire just as another device initially begins its transmission. As a result, two devices might send traffic across a single wire at the same time. When the two electrical signals collide, they become scrambled so that both transmissions are effectively lost.

Collision Detection

To handle collisions, Ethernet devices monitor the link while they are transmitting data. The monitoring process is known as collision detection. If a device detects a foreign signal while it is transmitting, it terminates the transmission and attempts to transmit again only after detecting an idle state on the wire. Collisions continue to occur if two colliding devices both wait the same amount of time before retransmitting. To avoid this condition, Ethernet devices use a binary exponential backoff algorithm.

Backoff Algorithm

With the binary exponential backoff algorithm, each device that sends a colliding transmission randomly selects a value within a range. The value represents the number of transmission times that the device must wait before retransmitting its data. If another collision occurs, the range of values is doubled and retransmission takes place again. Each time a collision occurs, the range of values doubles, to reduce the likelihood that two hosts on the same network can select the same retransmission time. Table 272 shows collision rounds up to round 10.

Table 272: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}

Copyright © 2016, Juniper Networks, Inc.

Table 272: Collision Backoff Algorithm Rounds *(continued)*

Round	Size of Set	Elements in the Set
2	4	{0,1,2,3}
3	8	{0,1,2,3,...,7}
4	16	{0,1,2,3,4,...,15}
5	32	{0,1,2,3,4,5,...,31}
6	64	{0,1,2,3,4,5,6,...,63}
7	128	{0,1,2,3,4,5,6,7,...,127}
8	256	{0,1,2,3,4,5,6,7,8,...,255}
9	512	{0,1,2,3,4,5,6,7,8,9,...,511}
10	1024	{0,1,2,3,4,5,6,7,8,9,10,...,1023}

Collision Domains and LAN Segments

Collisions are confined to a physical wire over which data is broadcast. Because the physical wires are subject to signal collisions, individual LAN segments are known as *collision domains*. Although the physical limitations on the length of an Ethernet cable restrict the length of a LAN segment, multiple collision domains can be interconnected by repeaters, bridges, and switches.

Repeaters

Repeaters are electronic devices that act on analog signals. Repeaters relay all electronic signals from one wire to another. A single repeater can double the distance between two devices on an Ethernet network. However, the Ethernet specification restricts the number of repeaters between any two devices on an Ethernet network to two, because collision detection with latencies increases in complexity as the wire length and number of repeaters increase.

Bridges and Switches

Bridges and switches combine LAN segments into a single Ethernet network by using multiple ports to connect the physical wires in each segment. Although bridges and switches are fundamentally the same, bridges generally provide more management and more interface ports. As Ethernet packets flow through a bridge, the bridge tracks the source MAC address of the packets and stores the addresses and their associated input ports in an interface table. As it receives subsequent packets, the bridge examines its interface table and takes one of the following actions:

- If the destination address does not match an address in the interface table, the bridge transmits the packet to all hosts on the network using the Ethernet broadcast address.

Copyright © 2016, Juniper Networks, Inc.

- If the destination address maps to the port through which the packet was received, the bridge or switch discards the packet. Because the other devices on the LAN segment also received the packet, the bridge does not need to retransmit it.
- If the destination address maps to a port other than the one through which the packet was received, the bridge transmits the packet through the appropriate port to the corresponding LAN segment.

Broadcast Domains

The combination of all the LAN segments within an Ethernet network is called a *broadcast domain*. In the absence of any signaling devices such as a repeater, bridge, or switch, the broadcast domain is simply the physical wire that makes up the connections in the network. If a bridge or switch is used, the broadcast domain consists of the entire LAN.

NOTE: **On all branch SRX Series devices, the subnet directed broadcast feature is not supported.**

Ethernet Frames

Data is transmitted through an Ethernet network in frames. The frames are of variable length, ranging from 64 octets to 1518 octets, including the header, payload, and cyclic redundancy check (CRC) value. Figure 118 shows the Ethernet frame format.

Figure 118: Ethernet Frame Format

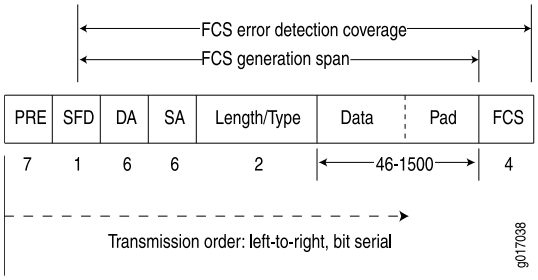

Ethernet frames have the following fields:

- The preamble (PRE) field is 7 octets of alternating 0s and 1s. The predictable format in the preamble allows receiving interfaces to synchronize themselves to the data being sent. The preamble is followed by a 1-octet start-of-frame delimiter (SFD).
- The destination address (DA) and source address (SA) fields contain the 6-octet (48-bit) MAC addresses for the destination and source ports on the network. These Layer 2 addresses uniquely identify the devices on the LAN.
- The Length/Type field is a 2-octet field that either indicates the length of the frame's data field or identifies the protocol stack associated with the frame. Here are some common frame types:
  - AppleTalk—**0x809B**
  - AppleTalk ARP—**0x80F3**

Copyright © 2016, Juniper Networks, Inc.

  - DECnet—**0x6003**
  - IP—**0x0800**
  - IPX—**0x8137**
  - Loopback—**0x9000**
  - XNS—**0x0600**
- The Data field contains the packet payload.
- The frame check sequence (FCS) is a 4-octet field that contains the calculated CRC value. This value is calculated by the originating host and appended to the frame. When it receives the frames, the receiving host calculates the CRC and checks it against this appended value to verify the integrity of the received frame.

Related Documentation

- Understanding Interfaces on page 2441
- Example: Creating an Ethernet Interface on page 2546
- Example: Deleting an Ethernet Interface on page 2547
- Understanding Static ARP Entries on Ethernet Interfaces on page 2545
- Understanding Promiscuous Mode on Ethernet Interface on page 2545

Understanding Static ARP Entries on Ethernet Interfaces

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Configuring Static ARP Entries on Ethernet Interfaces on page 2548

Understanding Promiscuous Mode on Ethernet Interface

When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet. You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Copyright © 2016, Juniper Networks, Inc.

Understanding Promiscuous Mode on the SRX5K-MPC

The promiscuous mode function is supported on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on the I/O cards (IOCs) and the SRX5000 line Module Port Concentrator (SRX5K-MPC).

When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or to the Services Processing Unit (SPU) regardless of the destination MAC address of the packet.

By default, an interface enables MAC filtering. You can configure promiscuous mode on the interface to disable MAC filtering. When you delete the promiscuous mode configuration, the interface will perform MAC filtering again.

You can change the MAC address of an interface even when the interface is operating in promiscuous mode. When the interface is operating in normal mode again, the MAC filtering function on the IOC uses the new MAC address to filter the packets.

You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 2551
- Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 2552

Example: Creating an Ethernet Interface

This example shows how to create an Ethernet interface.

- Requirements on page 2546
- Overview on page 2547
- Configuration on page 2547

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you create the ge-1/0/0 Ethernet interface and set the logical interface to 0. The logical unit number can range from 0 to 16,384. You can also add values for properties that you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

Step-by-Step Procedure

To configure an Ethernet interface:

1. Create the Ethernet interface and set the logical interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/0/0 unit 0
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

Purpose Verify if the configuration is working properly after creating the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Deleting an Ethernet Interface on page 2547

Example: Deleting an Ethernet Interface

This example shows how to delete an Ethernet interface.

- Requirements on page 2547
- Overview on page 2547
- Configuration on page 2548

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the ge-1/0/0 interface.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on J-Web pages.

Configuration

Step-by-Step Procedure

To delete an Ethernet interface:

1. Specify the interface you want to delete.

```
[edit]
user@host# delete interfaces ge-1/0/0
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

Purpose Verify if the configuration is working properly after deleting the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Creating an Ethernet Interface on page 2546

Example: Configuring Static ARP Entries on Ethernet Interfaces

- Requirements on page 2548
- Overview on page 2548
- Configuration on page 2549
- Verification on page 2549

Requirements

No special configuration beyond device initialization is required before creating an interface.

Overview

In this example, you configure a static ARP entry on the logical unit 0 of the ge-0/0/3 Gigabit Ethernet interface. The entry consists of the interface's IP address (10.1.1.1/24) and the corresponding MAC address of a node on the same Ethernet subnet (00:ff:85:7f:78:03). The example also configures the device to reply to ARP requests from the node using the publish option.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 mac
  00:ff:85:7f:78:03
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 publish
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static ARP entry on an Ethernet interface:

1. Create the Gigabit Ethernet interface.

   ```
   [edit]
   user@host# edit interfaces ge-0/0/3
   ```

2. Configure a static ARP entry.

   ```
   [edit interfaces ge-0/0/3]
   user@host# edit unit 0 family inet address 10.1.1.1/24
   ```

3. Set the IP address of the subnet node and the corresponding MAC address.

   ```
   [edit interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24]
   user@host# set arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-0/0/3** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host#show interfaces ge-0/0/3
unit 0 {
   family inet {
      address 10.1.1.1/24 {
         arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Static ARP Configurations on page 2550
- Verifying the Link State of All Interfaces on page 2550
- Verifying Interface Properties on page 2550

Copyright © 2016, Juniper Networks, Inc.

Verifying Static ARP Configurations

Purpose Verify the IP address and MAC (hardware) address of the node.

Action From operational mode, enter the **show interfaces ge-0/0/3** command.

Verifying the Link State of All Interfaces

Purpose Verify that all interfaces on the device are operational using the ping tool on each peer address in the network.

Action For each interface on the device:

1. In the J-Web interface, select **Troubleshoot>Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the time field..

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From operational mode, enter the **show interfaces detail** command.

```
user@host> show interfaces detail
Physical interface: ge-0/0/3, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 27, Generation: 17
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
  Last flapped   : 2004-08-25 15:42:30 PDT (4w5d 22:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0
```


Copyright © 2016, Juniper Networks, Inc.

```
  1 expedited-fo              0                0                0
  2 assured-forw              0                0                0
  3 network-cont              0                0                0
Active alarms  : None
Active defects : None
```

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
  - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces ge-0/0/3] level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces> ge-0/0/3 page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. The Last Flapped time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics ge-0/0/3** command.

Related Documentation

- Understanding Static ARP Entries on Ethernet Interfaces on page 2545

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

To enable promiscuous mode on an interface:

user@host# **set interfaces** *interface-name* **promiscuous-mode**

To disable promiscuous mode on an interface:

user@host# **delete interfaces** *interface-name* **promiscuous-mode**

Related Documentation

- Understanding Promiscuous Mode on Ethernet Interface on page 2545
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Promiscuous Mode on the SRX5K-MPC

This example shows how to configure promiscuous mode on an SRX5K-MPC interface in an SRX5600 to disable MAC address filtering.

- Requirements on page 2552
- Overview on page 2552
- Configuration on page 2552
- Verification on page 2553

Requirements

This example uses the following hardware and software components:

- An SRX5600 with an SRX5K-MPC that includes a 100-Gigabit Ethernet CFP transceiver
- Junos OS Release 12.1X47-D10 or later

No special configuration beyond device initialization is required before configuring this feature.

Overview

By default, the interfaces on an SRX5K-MPC have MAC address filtering enabled. In this example, you configure promiscuous mode on an interface to disable MAC address filtering. Then you delete promiscuous mode to reenable MAC address filtering on the interface.

Configuration

Configuring Promiscuous Mode on an Interface

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces et-4/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces et-4/0/0 promiscuous-mode
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure promiscuous mode:

1. Configure the ingress interface.

   ```
   [edit interfaces]
   user@host# set et-4/0/0 unit 0 family inet address 10.1.1.1/24
   ```

2. Enable promiscuous mode on the interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set et-4/0/0 promiscuous-mode
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
et-4/0/0 {
  promiscuous-mode;
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Disabling Promiscuous Mode on an Interface

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Step-by-Step Procedure To disable promiscuous mode:

1. Disable promiscuous mode on the interface.

```
[edit]
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Verification

Confirm that the configuration is working properly.

- Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC on page 2553
- Verifying the Status of Promiscuous Mode on page 2554
- Verifying That Promiscuous Mode Is Disabled on page 2555

Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC

Purpose Verify that promiscuous mode is enabled on the interface.

Action From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic
```

```
Physical interface: et-4/0/0, Enabled, Physical link is Up
  Interface index: 137, SNMP ifIndex: 511
```

Copyright © 2016, Juniper Networks, Inc.

```
Link-level type: Ethernet, MTU: 1518, Speed: 100Gbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: Promiscuous SNMP-Traps Internal: 0x4000
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 2c:21:72:3a:05:28, Hardware address: 2c:21:72:3a:05:28
Last flapped   : 2014-01-17 14:44:53 PST (5d 06:30 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0

Logical interface et-4/0/0.0 (Index 71) (SNMP ifIndex 513)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1351 ]  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: HOST
  Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
  ospf pgm pim rip router-discovery rsvp sap vrrp
  Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 122.122.122/24, Local: 122.122.122.1,
      Broadcast: 122.122.122.255
  Protocol multiservice, MTU: Unlimited
    Flags: Is-Primary

Logical interface et-4/0/0.32767 (Index 72) (SNMP ifIndex 517)
  Flags: SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: HOST
  Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
  ospf pgm pim rip router-discovery rsvp sap vrrp
  Protocol multiservice, MTU: Unlimited
    Flags: None
```

Meaning The **Interface flags: Promiscuous** field shows that promiscuous mode is enabled on the interface.

Verifying the Status of Promiscuous Mode

Purpose Verify that promiscuous mode works on the **et-4/0/0** interface.

Action Send traffic into the **et-4/0/0** interface with a MAC address that is different from the interface MAC address and turn on promiscuous mode.

From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic

Interface    Link  Input packets        (pps)     Output packets      (pps)
gr-0/0/0      Up              0          (0)               0          (0)
ip-0/0/0      Up              0          (0)               0          (0)
lt-0/0/0      Up              0          (0)               0          (0)
xe-1/2/0    Down              0          (0)               0          (0)
```


Copyright © 2016, Juniper Networks, Inc.

```
xe-1/2/1     Down              0          (0)              0          (0)
xe-1/2/2     Down              0          (0)              0          (0)
xe-1/2/3     Down              0          (0)              0          (0)
xe-1/2/4     Down              0          (0)              0          (0)
xe-1/2/5     Down              0          (0)              0          (0)
xe-1/2/6     Down              0          (0)              0          (0)
xe-1/2/7     Down              0          (0)              0          (0)
xe-1/2/8     Down              0          (0)              0          (0)
xe-1/2/9     Down              0          (0)              0          (0)
et-4/0/0       Up        4403996     (100002)              0          (0)
et-4/2/0     Up                3          (0)        4403924      (99997)
avs0         Up                0          (0)              0          (0)
avs1         Up                0          (0)              0          (0)
dsc          Up                0                           0
em0          Up            15965                       14056
```

Meaning The **input packets** and **pps** fields show that traffic is passing through the **et-4/0/0** interface as expected after promiscuous mode is enabled.

Verifying That Promiscuous Mode Is Disabled

Purpose Verify that disabled promiscuous mode works on the **et-4/0/0** interface.

Action Send traffic into the **et-4/0/0** interface with a MAC address that is different from the interface MAC address and turn off promiscuous mode.

From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic
```

```
Interface    Link  Input packets        (pps)     Output packets      (pps)
gr-0/0/0       Up              0          (0)                  0          (0)
ip-0/0/0       Up              0          (0)                  0          (0)
lt-0/0/0       Up              0          (0)                  0          (0)
xe-1/2/0     Down              0          (0)                  0          (0)
xe-1/2/1     Down              0          (0)                  0          (0)
xe-1/2/2     Down              0          (0)                  0          (0)
xe-1/2/3     Down              0          (0)                  0          (0)
xe-1/2/4     Down              0          (0)                  0          (0)
xe-1/2/5     Down              0          (0)                  0          (0)
xe-1/2/6     Down              0          (0)                  0          (0)
xe-1/2/7     Down              0          (0)                  0          (0)
xe-1/2/8     Down              0          (0)                  0          (0)
xe-1/2/9     Down              0          (0)                  0          (0)
et-4/0/0       Up       11505495          (0)                  0          (0)
et-4/2/0     Up                6          (0)           11505425          (0)
avs0         Up                0          (0)                  0          (0)
avs1         Up                0          (0)                  0          (0)
dsc          Up                0                               0
em0          Up            37964                           31739
```

Meaning The **pps** field shows that the traffic is not passing through the **et-4/0/0** interface after promiscuous mode is disabled.

Related Documentation

- Understanding Promiscuous Mode on Ethernet Interface on page 2545

- Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 2551

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 115

Configuring Aggregated Ethernet Interfaces

- Understanding Aggregated Ethernet Interfaces on page 2557
- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Understanding the Aggregated Ethernet Interfaces Device Count on page 2560
- Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 2561
- Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 2562
- Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 2562
- Understanding Aggregated Ethernet Interface Link Speed on page 2563
- Example: Configuring Aggregated Ethernet Link Speed on page 2564
- Understanding Minimum Links for Aggregated Ethernet Interfaces on page 2565
- Example: Configuring Aggregated Ethernet Minimum Links on page 2565
- Understanding Aggregated Ethernet Interface Removal on page 2566
- Example: Deleting Aggregated Ethernet Interfaces on page 2566
- Example: Deleting Aggregated Ethernet Interface Contents on page 2567
- Verifying Aggregated Ethernet Interfaces on page 2568
- Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 2570
- Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 2570

Understanding Aggregated Ethernet Interfaces

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on Layer 3 information carried in the packet, Layer 4 information carried in the packet, or both, or based on session ID data. (The session ID data has higher precedence than the Layer 3 or 4 information.) This implementation uses the same load-balancing algorithm used for per-packet load balancing.

Aggregated Ethernet interfaces can be Layer 3 interfaces (VLAN-tagged or untagged) and Layer 2 interfaces.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This topic is specific to the SRX3000 and SRX5000 line devices. For information about link aggregation for other SRX Series devices, see the *Layer 2 Bridging and Transparent Mode for Security Devices*.

This topic contains the following sections:

- LAGs on page 2558
- LACP on page 2558

LAGs

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle, such that a media access control (MAC) client can treat the LAG as if it were a single link. Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface. For the LAG to operate correctly, it is necessary to coordinate the two end systems connected by the LAG, either manually or automatically.

Internally, a LAG is a virtual interface presented on SRX3000 and SRX5000 line devices or on any system (consisting of devices such as routers and switches) supporting 802.3ad link aggregation. Externally, a LAG corresponds to a bundle of physical Ethernet links connected between an SRX3000 or SRX5000 line device and another system capable of link aggregation. This bundle of physical links is a virtual link.

Follow these guidelines for aggregated Ethernet support for the SRX3000 and SRX5000 lines:

- The devices support a maximum of 16 physical interfaces per single aggregated Ethernet bundle.
- Aggregated Ethernet interfaces can use interfaces from the same or different Flexible PIC Concentrators (FPCs) and PICs.
- On the aggregated bundle, capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available.

LACP

Junos OS supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs.

Starting with Junos OS Release 15.1X49-D40, LACP is supported on Layer 2 transparent mode in addition to existing support on Layer 3 mode. For information about link aggregation for other SRX Series devices, see the *Layer 2 Bridging and Transparent Mode for Security Devices*.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation

Copyright © 2016, Juniper Networks, Inc.

control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

For example, when LACP is not enabled, a local LAG might attempt to transmit packets to a remote individual interface, which causes the communication to fail. (An individual interface is a nonaggregatable interface.) When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device. Then you associate a set of ports that have the same speed and are in full-duplex mode. The physical ports can be 100-megabit Ethernet, 1-Gigabit Ethernet, and 10-Gigabit Ethernet.

When configuring LACP, follow these guidelines:

- LACP does not support automatic configuration on SRX3000 and SRX5000 line devices, but partner systems are allowed to perform automatic configuration. When an SRX3000 or SRX5000 line device is connected to a fully 802.3ad-compliant partner system, static configuration of LAGs is initiated on the SRX3000 and SRX5000 line device side, and static configuration is not needed on the partner side.
- When an SRX3000 or SRX5000 line device is connected to a Juniper Networks MX Series router, static configuration of LAGs is needed at both the actor (local or near-end of the link) and partner systems.
- Although the LACP functions on the SRX3000 and SRX5000 line devices are similar to the LACP features on Juniper Networks MX Series routers, the following LACP features on MX Series routers are not supported on SRX3000 and SRX5000 line devices: link protection, system priority, and port priority for aggregated Ethernet interfaces. Instead, SRX3000 and SRX5000 line devices provide active/standby support with redundant Ethernet interface LAGs in chassis cluster deployments.

LACP is supported in standalone deployments, where aggregated Ethernet interfaces are supported, and in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Understanding LACP on Standalone Devices on page 2571
- Understanding LACP on Chassis Clusters on page 2575
- Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 2570
- Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 2570

Copyright © 2016, Juniper Networks, Inc.

Aggregated Ethernet Interfaces Configuration Overview

NOTE: This topic is specific to the SRX3000 and SRX5000 line devices.

To configure an aggregated Ethernet interface:

1. Set the number of aggregated Ethernet interfaces on the device. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
2. Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.
3. (Optional) Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 2564.
4. (Optional) Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 2565.
5. (Optional) Enable or disable VLAN tagging. See "Understanding VLAN Tagging for Aggregated Ethernet Interfaces" on page 2570.
6. (Optional) Enable promiscuous mode. See "Understanding Promiscuous Mode for Aggregated Ethernet Interfaces" on page 2570.

Related Documentation

- *Layer 2 Bridging and Transparent Mode for Security Devices*
- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Configuring LACP on Standalone Devices on page 2572
- Example: Configuring LACP on Chassis Clusters on page 2577

Understanding the Aggregated Ethernet Interfaces Device Count

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them. Once you set the device count, the system creates that number of empty aggregated Ethernet interfaces. A globally unique MAC address is assigned to every aggregated Ethernet interface. More aggregated Ethernet interfaces can be created by increasing the parameter.

The maximum number of aggregated devices you can configure is 128. The aggregated interfaces are numbered from ae0 through ae127.

Similarly, you can permanently remove an aggregated Ethernet interface from the device configuration by deleting it from the device count. When you reduce the device count, only the aggregated Ethernet interface objects at the end of the list are removed, leaving the newly specified number of interfaces. That is, if you set the device count to 10 and then reduce it to 6, the system removes the last 4 interface objects from the list.

Copyright © 2016, Juniper Networks, Inc.

WARNING: Be aware that this approach deletes the aggregated Ethernet interface and *all* of its objects from the device configuration.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 2561
- Example: Deleting Aggregated Ethernet Interfaces on page 2566

Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device

This example shows how to configure the number of aggregated Ethernet interfaces on a device.

- Requirements on page 2561
- Overview on page 2561
- Configuration on page 2561
- Verification on page 2561

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you create two aggregate Ethernet interfaces, thereby enabling all the interfaces that you need for your configuration in one step.

Configuration

Step-by-Step Procedure

To configure the number of aggregated Ethernet interfaces on a device:

1. Set the number of aggregated Ethernet interfaces.

   ```
   [edit]
   user@host# set chassis aggregated-devices ethernet device-count 2
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show chassis aggregated-devices** command.

Related Documentation

- Understanding the Aggregated Ethernet Interfaces Device Count on page 2560

Copyright © 2016, Juniper Networks, Inc.

- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Example: Deleting Aggregated Ethernet Interfaces on page 2566
- Verifying Aggregated Ethernet Interfaces on page 2568

Understanding Physical Interfaces for Aggregated Ethernet Interfaces

You associate a physical interface with an aggregated Ethernet interface. Doing so associates the physical child links with the logical aggregated parent interface to form a link aggregation group (LAG). You must also specify the constituent physical links by including the **802.3ad** configuration statement.

A physical interface can be added to any aggregated Ethernet interface as long as all member links have the same link speed and the maximum number of member links does not exceed 16. The aggregated Ethernet interface instance number ae*x* can be from 0 through 127, for a total of 128 aggregated interfaces.

NOTE: If you specify (on purpose or accidentally) that a link already associated with an aggregated Ethernet interface be associated with another aggregated Ethernet interface, the link is removed from the previous interface (there is no need for you to explicitly delete it) and it is added to the other one.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 2562

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces

This example shows how to associate physical interfaces with aggregated Ethernet interfaces.

- Requirements on page 2562
- Overview on page 2563
- Configuration on page 2563
- Verification on page 2563

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See “Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device” on page 2561.

 Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you associate the physical child link of the ge-1/0/0 and ge-2/0/0 physical interfaces with the logical aggregate parent, ae0, thereby creating a LAG. Similarly, you create a LAG that associate the ge-3/0/0, ge-3/0/1, and ge-4/0/1 physical interfaces with the ae1 aggregated Ethernet interface.

Configuration

Step-by-Step Procedure

To associate physical interfaces with aggregated Ethernet interfaces:

1. Create the first LAG.

```
[edit]
user@host# set interfaces ge-1/0/0 gigether-options 802.3ad ae0
user@host# set interfaces ge-2/0/0 gigether-options 802.3ad ae0
```

2. Create the second LAG.

```
[edit]
user@host# set interfaces ge-3/0/0 gigether-options 802.3ad ae1
user@host# set interfaces ge-3/0/1 gigether-options 802.3ad ae1
user@host# sset interfaces ge-4/0/0 gigether-options 802.3ad ae1
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 2562
- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Verifying Aggregated Ethernet Interfaces on page 2568

Understanding Aggregated Ethernet Interface Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed you specify in the **link-speed** parameter, an error message will be logged.

The speed value is specified in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Copyright © 2016, Juniper Networks, Inc.

Aggregated Ethernet interfaces on SRX3000 and SRX5000 line devices can have one of the following speed values:

- 100m—Links are 100 Mbps.
- 10g—Links are 10 Gbps.
- 1g—Links are 1 Gbps.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Configuring Aggregated Ethernet Link Speed on page 2564
- Understanding Minimum Links for Aggregated Ethernet Interfaces on page 2565

Example: Configuring Aggregated Ethernet Link Speed

This example shows how to configure the aggregated Ethernet link speed.

- Requirements on page 2564
- Overview on page 2564
- Configuration on page 2564
- Verification on page 2565

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.

Overview

In this example, you set the required link speed for all interfaces included in the bundle to 10 Gbps. All interfaces that make up a bundle must be the same speed.

Configuration

Step-by-Step Procedure

To configure the aggregated Ethernet link speed:

1. Set the link speed.

```
[edit]
user@host# set interfaces ae0 aggregated-ether-options link-speed 10g
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- Understanding Aggregated Ethernet Interface Link Speed on page 2563
- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Verifying Aggregated Ethernet Interfaces on page 2568

Understanding Minimum Links for Aggregated Ethernet Interfaces

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. By default, only one link must be up for the bundle to be labeled as up.

On SRX3000 and SRX5000 line devices, the valid range for the minimum links number is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled as up.

If the number of links configured in an aggregated Ethernet interface is less than the **minimum-links** value configured in the **minimum-links** statement, the configuration commit fails and an error message is displayed.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Configuring Aggregated Ethernet Minimum Links on page 2565
- Understanding Aggregated Ethernet Interface Link Speed on page 2563

Example: Configuring Aggregated Ethernet Minimum Links

This example shows how to configure the minimum number of links on an aggregated Ethernet interface that must be up for the bundle as a whole to be labeled as up.

- Requirements on page 2565
- Overview on page 2566
- Configuration on page 2566
- Verification on page 2566

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.

Copyright © 2016, Juniper Networks, Inc.

- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 2564.

Overview

In this example, you specify that on interface ae0 at least eight links must be up for the bundle as a whole to be labeled as up.

Configuration

Step-by-Step Procedure

To configure the minimum number of links on an aggregated Ethernet interface:

1. Set the minimum number of links.

   ```
   [edit]
   user@host# set interfaces ae0 aggregated-ether-options minimum-links 8
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- Understanding Aggregated Ethernet Interface Link Speed on page 2563
- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Verifying Aggregated Ethernet Interfaces on page 2568

Understanding Aggregated Ethernet Interface Removal

You can delete an aggregated Ethernet interface from the interface configuration. Junos OS removes the configuration statements related to **aex** and sets this interface to the down state. The deleted aggregated Ethernet interface still exists, but it becomes an empty interface.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Example: Deleting Aggregated Ethernet Interfaces on page 2566
- Example: Deleting Aggregated Ethernet Interface Contents on page 2567

Example: Deleting Aggregated Ethernet Interfaces

This example shows how to delete aggregated Ethernet interfaces using the device count.

- Requirements on page 2567
- Overview on page 2567

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 2567
- Verification on page 2567

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.

Overview

This example shows how to clean up unused aggregated Ethernet interfaces. In this example, you reduce the number of interfaces from 10 to 6, thereby removing the last 4 interfaces from the interface object list.

Configuration

Step-by-Step Procedure

To delete an interface:

1. Set the number of aggregated Ethernet interfaces.

   ```
   [edit]
   user@host# delete chassis aggregated-devices ethernet device-count 6
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show chassis aggregated-devices** command.

Related Documentation

- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Example: Deleting Aggregated Ethernet Interface Contents on page 2567
- Verifying Aggregated Ethernet Interfaces on page 2568

Example: Deleting Aggregated Ethernet Interface Contents

This example shows how to delete the contents of an aggregated Ethernet interface.

- Requirements on page 2567
- Overview on page 2568
- Configuration on page 2568
- Verification on page 2568

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Set the number of aggregated Ethernet interfaces on the device. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
- Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.
- Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 2564.
- Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 2565.

Overview

In this example, you delete the contents of the ae4 aggregated Ethernet interface, which sets it to the down state.

Configuration

Step-by-Step Procedure

To delete the contents of an aggregated Ethernet interface:

1. Delete the interface.

   ```
   [edit]
   user@host# delete interfaces ae4
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Example: Deleting Aggregated Ethernet Interfaces on page 2566
- Verifying Aggregated Ethernet Interfaces on page 2568

Verifying Aggregated Ethernet Interfaces

- Verifying Aggregated Ethernet Interfaces (terse) on page 2568
- Verifying Aggregated Ethernet Interfaces (extensive) on page 2569

Verifying Aggregated Ethernet Interfaces (terse)

Purpose Display status information in terse (concise) format for aggregated Ethernet interfaces.

Action From operational mode, enter the **show interfaces ae0 terse** command.

```
user@host> show interfaces ae0 terse
```


Copyright © 2016, Juniper Networks, Inc.

```
ge-2/0/0.0              up    up   aenet    --> ae0.0
ge-2/0/0.32767          up    up   aenet    --> ae0.32767
ge-2/0/1.0              up    up   aenet    --> ae0.0
ge-2/0/1.32767          up    up   aenet    --> ae0.32767
ae0                     up    up
ae0.0                   up    up   bridge
ae0.32767               up    up   multiservice
```

The output shows the bundle relationship for the aggregated Ethernet interface and the overall status of the interface, including the following information:

- The link aggregation control PDUs run on the .0 child logical interfaces for the untagged aggregated Ethernet interface.
- The link aggregation control PDUs run on the .32767 child logical interfaces for the VLAN-tagged aggregated Ethernet interface.
- The .32767 logical interface is created for the parent link and all child links.

Verifying Aggregated Ethernet Interfaces (extensive)

Purpose Display status information and statistics in extensive (detailed) format for aggregated Ethernet interfaces.

Action From operational mode, enter the **show interfaces ae0 extensive** command.

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  ...
  Logical interface ae0.0 (Index 67) (SNMP ifIndex 628) (Generation 134)
    ...
    LACP info:        Role     System             System      Port    Port  Port

                              priority         identifier  priority  number   key

      ge-5/0/0.0     Actor        127  00:1f:12:8c:af:c0       127     832     1

      ge-5/0/0.0   Partner        127  00:1f:12:8f:d7:c0       127     640     1

      ge-5/0/1.0     Actor        127  00:1f:12:8c:af:c0       127     833     1

      ge-5/0/1.0   Partner        127  00:1f:12:8f:d7:c0       127     641     1

    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-5/0/0.0             12830        7090            0            0
      ge-5/0/1.0             10304        4786            0            0
    ...
  Logical interface ae0.32767 (Index 70) (SNMP ifIndex 630) (Generation 135)
    ...
    LACP info:        Role     System             System      Port    Port  Port

                              priority         identifier  priority  number   key

      ge-5/0/0.32767    Actor      127 00:1f:12:8c:af:c0       127     832     1

      ge-5/0/0.32767  Partner      127 00:1f:12:8f:d7:c0       127     640     1

      ge-5/0/1.32767    Actor      127 00:1f:12:8c:af:c0       127     833     1
```

Copyright © 2016, Juniper Networks, Inc.

```
    ge-5/0/1.32767  Partner      127 00:1f:12:8f:d7:c0       127     641     1

  LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
    ge-5/0/0.32767         12830        7090            0            0
    ge-5/0/1.32767         10304        4786            0            0
...
```

The output shows detailed aggregated Ethernet interface information. This portion of the output shows LACP information and LACP statistics for each logical aggregated Ethernet interface.

Related Documentation

- Aggregated Ethernet Interfaces Configuration Overview on page 2560

Understanding VLAN Tagging for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can be either VLAN-tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on the SRX3000 and SRX5000 lines support the configuration of **native-vlan-id**, which consists of the following configuration statements:

- **inner-tag-protocol-id**
- **inner-vlan-id**
- **pop-pop**
- **pop-swap**
- **push-push**
- **swap-push**
- **swap-swap**

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Aggregated Ethernet Interfaces Configuration Overview on page 2560

Understanding Promiscuous Mode for Aggregated Ethernet Interfaces

You can enable promiscuous mode on aggregated Ethernet interfaces. When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Aggregated Ethernet Interfaces Configuration Overview on page 2560

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 116

Configuring Link Aggregation Control Protocol

- Understanding LACP on Standalone Devices on page 2571
- Example: Configuring LACP on Standalone Devices on page 2572
- Verifying LACP on Standalone Devices on page 2573
- Understanding LACP on Chassis Clusters on page 2575
- Example: Configuring LACP on Chassis Clusters on page 2577
- Verifying LACP on Redundant Ethernet Interfaces on page 2579
- LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs) on page 2580
- Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3 on page 2582

Understanding LACP on Standalone Devices

Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between partner systems on a link. Within LACP, the local end of a child link is known as the actor and the remote end of the link is known as the partner.

LACP is enabled on an aggregated Ethernet interface by setting the mode to either passive or active. However, to initiate the transmission of link aggregation control protocol data units (PDUs) and response link aggregation control PDUs, you must enable LACP at both the local and remote ends of the links, and one end must be active:

- Active mode—If either the actor or partner is active, they exchange link aggregation control PDUs. The actor sends link aggregation control PDUs to its protocol partner that convey what the actor knows about its own state and that of the partner's state.
- Passive mode—If the actor and partner are both in passive mode, they do not exchange link aggregation control PDUs. As a result, the aggregated Ethernet links do not come up. In passive transmission mode, links send out link aggregation control PDUs only when they receive them from the remote end of the same link.

By default, the actor and partner transmit link aggregation control PDUs every second. You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

Copyright © 2016, Juniper Networks, Inc.

You configure the interval at which the interfaces on the remote side of the link transmit link aggregation control PDUs by configuring the **periodic** statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval. The interval can be **fast** (every second) or **slow** (every 30 seconds).

NOTE: **On all high-end SRX Series devices, the LACP is not supported on Layer 2 interfaces.**

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 2557
- Understanding LACP on Chassis Clusters on page 2575
- Example: Configuring LACP on Standalone Devices on page 2572

Example: Configuring LACP on Standalone Devices

This example shows how to configure LACP on standalone devices.

- Requirements on page 2572
- Overview on page 2572
- Configuration on page 2573
- Verification on page 2573

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.
- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 2564.
- Configure the aggregated Ethernet minimum links speed. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 2565.

Overview

In this example, you set LACP to passive mode for the ae0 interface. You set the LACP mode for the ae1 interface to active and set the link aggregation control PDU transmit interval to slow, which is every 30 seconds.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure LACP on standalone devices:

1. Set the first LACP.

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options lacp passive
   ```

2. Set the second LACP.

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-options lacp active
   user@host# set ae1 aggregated-ether-options lacp periodic slow
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit interfaces]
   user@host# commit
   ```

Verification

To confirm that the configuration is working properly, perform these tasks:

Related Documentation

- Understanding LACP on Standalone Devices on page 2571
- Verifying LACP on Standalone Devices on page 2573

Verifying LACP on Standalone Devices

- Verifying LACP Statistics on page 2573
- Verifying LACP Aggregated Ethernet Interfaces on page 2574

Verifying LACP Statistics

Purpose Display LACP statistics for aggregated Ethernet interfaces.

Action From operational mode, enter the **show lacp statistics interfaces ae0** command.

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-2/0/0                1352        2035            0            0
      ge-2/0/1                1352        2056            0            0
      ge-2/2/0                1352        2045            0            0
      ge-2/2/1                1352        2043            0            0
```

The output shows LACP statistics for each physical interface associated with the aggregated Ethernet interface, such as the following:

Copyright © 2016, Juniper Networks, Inc.

- The LACP received counter that increments for each normal hello
- The number of LACP transmit packet errors logged
- The number of unrecognized packet errors logged
- The number of invalid packets received

Use the following command to clear the statistics and see only new changes:

```
user@host# clear lacp statistics interfaces ae0
```

Verifying LACP Aggregated Ethernet Interfaces

Purpose Display LACP status information for aggregated Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces ae0** command.

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/0                  Current   Fast periodic Collecting distributing
      ge-2/0/1                  Current   Fast periodic Collecting distributing
      ge-2/2/0                  Current   Fast periodic Collecting distributing
      ge-2/2/1                  Current   Fast periodic Collecting distributing
```

The output shows aggregated Ethernet interface information, including the following information:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up ifit is collecting and distributing packets.

Related Documentation

- Example: Configuring LACP on Standalone Devices on page 2572
- Verifying LACP on Redundant Ethernet Interfaces on page 2579

Copyright © 2016, Juniper Networks, Inc.

Understanding LACP on Chassis Clusters

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle, such that a media access control (MAC) client can treat the LAG as if it were a single link.

LAGs can be established across nodes in a chassis cluster to provide increased interface bandwidth and link availability.

The Link Aggregation Control Protocol (LACP) provides additional functionality for LAGs. LACP is supported in standalone deployments, where aggregated Ethernet interfaces are supported, and in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

You configure LACP on a redundant Ethernet interface by setting the LACP mode for the parent link with the **lacp** statement. The LACP mode can be off (the default), active, or passive.

This topic contains the following sections:

- Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups on page 2575
- Sub-LAGs on page 2576
- Supporting Hitless Failover on page 2576
- Managing Link Aggregation Control PDUs on page 2577

Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups

A redundant Ethernet interface has active and standby links located on two nodes in a chassis cluster. All active links are located on one node, and all standby links are located on the other node. You can configure up to eight active links and eight standby links per node.

When at least two physical child interface links from each node are included in a redundant Ethernet interface configuration, the interfaces are combined within the redundant Ethernet interface to form a redundant Ethernet interface LAG.

Having multiple active redundant Ethernet interface links reduces the possibility of failover. For example, when an active link is out of service, all traffic on this link is distributed to other active redundant Ethernet interface links, instead of triggering a redundant Ethernet active/standby failover.

Aggregated Ethernet interfaces, known as local LAGs, are also supported on either node of a chassis cluster but cannot be added to redundant Ethernet interfaces. Likewise, any child interface of an existing local LAG cannot be added to a redundant Ethernet interface, and vice versa. The total maximum number of combined individual node LAG interfaces (ae) and redundant Ethernet (reth) interfaces per cluster is 128.

However, aggregated Ethernet interfaces and redundant Ethernet interfaces can coexist, because the functionality of a redundant Ethernet interface relies on the Junos OS aggregated Ethernet framework.

Copyright © 2016, Juniper Networks, Inc.

For more information, see *Understanding Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups for Branch SRX Series Devices* or *Understanding Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups for High-End SRX Series Devices*.

Minimum Links

Redundant Ethernet interface configuration includes a **minimum-links** setting that allows you to set a minimum number of physical child links in a redundant Ethernet interface LAG that must be working on the primary node for the interface to be up. The default **minimum-links** value is 1. When the number of physical links on the primary node in a redundant Ethernet interface falls below the **minimum-links** value, the interface will be down even if some links are still working. For more information, see *Example: Configuring Chassis Cluster Minimum Links*.

Sub-LAGs

LACP maintains a point-to-point LAG. Any port connected to the third point is denied. However, a redundant Ethernet interface does connect to two different systems or two remote aggregated Ethernet interfaces by design.

To support LACP on both redundant Ethernet interface active and standby links, a redundant Ethernet interface can be modeled to consist of two sub-LAGs, where all active links form an active sub-LAG and all standby links form a standby sub-LAG.

In this model, LACP selection logic is applied and limited to one sub-LAG at a time. In this way, two redundant Ethernet interface sub-LAGs are maintained simultaneously while all the LACP advantages are preserved for each sub-LAG.

It is necessary for the switches used to connect the nodes in the cluster to have a LAG link configured and 802.3ad enabled for each LAG on both nodes so that the aggregate links will be recognized as such and correctly pass traffic.

NOTE: The redundant Ethernet interface LAG child links from each node in the chassis cluster must be connected to a different LAG at the peer devices. If a single peer switch is used to terminate the redundant Ethernet interface LAG, two separate LAGs must be used in the switch.

Supporting Hitless Failover

With LACP, the redundant Ethernet interface supports hitless failover between the active and standby links in normal operation. The term *hitless* means that the redundant Ethernet interface state remains up during a failover.

The lacpd process manages both the active and standby links of the redundant Ethernet interfaces. A redundant Ethernet interface state remains up when the number of active up links is more than the number of minimum links configured. Therefore, to support hitless failover, the LACP state on the redundant Ethernet interface standby links must be collected and distributed before failover occurs.

Copyright © 2016, Juniper Networks, Inc.

Managing Link Aggregation Control PDUs

The protocol data units (PDUs) contain information about the state of the link. By default, aggregated and redundant Ethernet links do not exchange link aggregation control PDUs.

You can configure PDUs exchange in the following ways:

- Configure Ethernet links to actively transmit link aggregation control PDUs
- Configure Ethernet links to passively transmit PDUs, sending out link aggregation control PDUs only when they are received from the remote end of the same link

The local end of a child link is known as the actor and the remote end of the link is known as the partner. That is, the actor sends link aggregation control PDUs to its protocol partner that convey what the actor knows about its own state and that of the partner's state.

You configure the interval at which the interfaces on the remote side of the link transmit link aggregation control PDUs by configuring the **periodic** statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval. The interval can be **fast** (every second) or **slow** (every 30 seconds).

For more information, see "Example: Configuring LACP on Chassis Clusters" on page 2577.

By default, the actor and partner transmit link aggregation control PDUs every second. You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

Related Documentation

- Example: Configuring LACP on Chassis Clusters on page 2577

Example: Configuring LACP on Chassis Clusters

This example shows how to configure LACP on chassis clusters.

- Requirements on page 2577
- Overview on page 2578
- Configuration on page 2578
- Verification on page 2578

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See "Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device" on page 2561.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 2562.

Copyright © 2016, Juniper Networks, Inc.

- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 2564.
- Configure the aggregated Ethernet minimum links speed. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 2565.
- Configure the LACP on standalone devices. See "Example: Configuring LACP on Standalone Devices" on page 2572.

Overview

In this example, you set LACP to passive mode for the reth0 interface. You set the LACP mode for the reth1 interface to active and set the link aggregation control PDU transmit interval to slow, which is every 30 seconds.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To configure LACP on chassis clusters:

1. Set the first LACP on primary node1.

   ```
   [edit interfaces]
   user@host# set reth0 redundant-ether-options lacp passive
   ```

2. Set the second LACP.

   ```
   [edit interfaces]
   user@host# set reth1 redundant-ether-options lacp active
   user@host# set reth1 redundant-ether-options lacp periodic slow
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit interfaces]
   user@host# commit
   ```

Verification

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth0** command.

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-11/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
```


Copyright © 2016, Juniper Networks, Inc.

```
      ge-11/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-11/0/0               Current   Fast periodic Collecting distributing
      ge-11/0/1               Current   Fast periodic Collecting distributing
      ge-11/0/2               Current   Fast periodic Collecting distributing
      ge-11/0/3               Current   Fast periodic Collecting distributing
      ge-3/0/0                Current   Fast periodic Collecting distributing
      ge-3/0/1                Current   Fast periodic Collecting distributing
      ge-3/0/2                Current   Fast periodic Collecting distributing
      ge-3/0/3                Current   Fast periodic Collecting distributing
{primary:node1}
```

The output shows redundant Ethernet interface information, such as the following:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

Related Documentation

- Understanding LACP on Chassis Clusters on page 2575
- Verifying LACP on Redundant Ethernet Interfaces on page 2579

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth0** command.

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-11/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
```

```
      ge-3/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-11/0/0                Current   Fast periodic Collecting distributing
      ge-11/0/1                Current   Fast periodic Collecting distributing
      ge-11/0/2                Current   Fast periodic Collecting distributing
      ge-11/0/3                Current   Fast periodic Collecting distributing
      ge-3/0/0                 Current   Fast periodic Collecting distributing
      ge-3/0/1                 Current   Fast periodic Collecting distributing
      ge-3/0/2                 Current   Fast periodic Collecting distributing
      ge-3/0/3                 Current   Fast periodic Collecting distributing
{primary:node1}
```

The output shows redundant Ethernet interface information, such as the following:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

Related Documentation

- Example: Configuring LACP on Chassis Clusters on page 2577
- Verifying LACP on Standalone Devices on page 2573

LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs)

LAG and LACP Support on the SRX5000 Module Port Concentrator

The SRX5000 Module Port Concentrator (SRX5K-MPC) on SRX5400, SRX5600, and SRX5800 devices supports link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP).

Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

The following LAG and LACP features are supported on the SRX5K-MPC:

- Bandwidth aggregation—Increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

Copyright © 2016, Juniper Networks, Inc.

- Link redundancy and load balancing (within chassis cluster)—Provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.
- Dynamic link management—Enables automatic addition and deletion of individual links to the aggregate bundle without user intervention.

LACP supports the following features:

- LACP bundles several physical interfaces to form one logical interface by exchanging LACP packets between the local interface and the remote interface. LACP monitors the link for changes in interface state by exchanging a periodic LACP heartbeat between two sides. Any changes in interface state are reflected in the LACP packet.
- Normally after an LACP is configured and committed, two sides start to exchange interface and port information. Once they identify each other and match the LACP state machine criteria, the LACP is declared as up. You can deactivate or delete the LACP configuration.
- By default, the LACP packets are exchanged in every second. You can configure the LACP interval as fast (every second) or slow (every 30 seconds) to ensure the health of the interfaces.
- LACP supports distributed and centralized modes. Chassis cluster setup is recommended to operate with LACP distributed mode, which handles chassis cluster failover better. The centralized mode might experience traffic loss during failover.

SRX5K-MPCs on SRX5000 line devices provide active and standby support with redundant Ethernet interface LAGs in chassis cluster deployments.

LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode

The IOC2 and IOC3 cards on SRX5400, SRX5600, and SRX5800 devices support link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP) in Express path mode.

You can use the links in a LAG as ingress or egress interfaces in Express Path mode. The LAG links can include links from cards such as IOC2 or IOC3. For a LAG link to qualify for Express Path, all its member links should be connected to Express Path-enabled network processors. If Express Path is disabled on any of the member links in a LAG, a regular session (non-Express Path session) is created.

NOTE:

- Cross-IOC LAG interfaces do not support Layer 2 transparent mode.
- Mixed interface speeds are not supported on the same aggregated bundle.
- A redundant Ethernet interface or aggregated Ethernet interface must contain child interfaces from the same IOC type. For example, if one child link is from 10-Gigabit Ethernet on IOC2, the second child link should also be from IOC2. Similarly, both child interfaces can be from IOC3. Configuring child interfaces by mixing links from both IOC2 and IOC3 is not supported.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Aggregated Ethernet Interfaces Configuration Overview on page 2560
- Example: Configuring LACP on Standalone Devices on page 2572
- Example: Configuring LACP on Chassis Clusters on page 2577

Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a LAG or bundle. The LACP provides additional functionality for LAGs.

This example shows how to configure LAG on an SRX Series device using the links from either IOC2 or IOC3 in Express Path mode.

- Requirements on page 2582
- Overview on page 2582
- Configuration on page 2583
- Verification on page 2585

Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1X49-D40 or later for SRX Series devices.
- SRX5800 with IOC2 or IOC3 with Express Path enabled on IOC2 and IOC3. For details, see "Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path" on page 1739.

Overview

In this example, you create a logical aggregated Ethernet interface and define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and LACP. Next, define the member links to be contained within the aggregated Ethernet interface—for example, four 10-Gigabit Ethernet interfaces. Finally, configure an LACP for link detection.

The following member links are used in this example:

- xe-0/0/8
- xe-0/0/9
- xe-1/0/8
- xe-1/0/9
- xe-3/1/4
- xe-3/1/5

Copyright © 2016, Juniper Networks, Inc.

- xe-5/1/4
- xe-5/1/5

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, delete, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/8 gigether-options 802.3ad ae1
set interfaces xe-0/0/9 gigether-options 802.3ad ae0
set interfaces xe-1/0/8 gigether-options 802.3ad ae1
set interfaces xe-1/0/9 gigether-options 802.3ad ae0
set interfaces xe-3/1/4 gigether-options 802.3ad ae1
set interfaces xe-3/1/5 gigether-options 802.3ad ae0
set interfaces xe-5/1/4 gigether-options 802.3ad ae1
set interfaces xe-5/1/5 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 17.0.0.1/24
set interfaces ae1 unit 0 family inet address 16.0.0.1/24
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp active
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in *CLI User Guide*.

To configure LAG Interfaces:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@host# set xe-0/0/8 gigether-options 802.3ad ae1
   user@host# set xe-0/0/9 gigether-options 802.3ad ae0
   user@host# set xe-1/0/8 gigether-options 802.3ad ae1
   user@host# set xe-1/0/9 gigether-options 802.3ad ae0
   user@host# set xe-3/1/4 gigether-options 802.3ad ae1
   user@host# set xe-3/1/5 gigether-options 802.3ad ae0
   user@host# set xe-5/1/4 gigether-options 802.3ad ae1
   user@host# set xe-5/1/5 gigether-options 802.3ad ae0
   ```

3. Assign an IP address to ae0 and ae1.

   ```
   [edit interfaces]
   user@host# set ae0 unit 0 family inet address 17.0.0.1/24
   user@host# set ae1 unit 0 family inet address 16.0.0.1/24
   ```

4. Set the LACP on reth0.

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options lacp active
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ae1 aggregated-ether-options lacp active
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
xe-0/0/8 {
  gigether-options {
    802.3ad ae1;
  }
}
xe-0/0/9 {
  gigether-options {
    802.3ad ae0;
  }
}
xe-1/0/8 {
  gigether-options {
    802.3ad ae1;
  }
}
xe-1/0/9 {
  gigether-options {
    802.3ad ae0;
  }
}
xe-3/1/4 {
  gigether-options {
    802.3ad ae1;
  }
}
xe-3/1/5 {
  gigether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 17.0.0.1/24;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    unit 0 {
      family inet {
        address 16.0.0.1/24;
      }
    }
  }
```

```
[edit]
user@host# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces** command to check that LACP has been enabled as active on one end.

```
user@host> show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/9       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/9     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/9       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/9     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/5       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/5     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/5       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/5     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/9                  Current   Fast periodic Collecting distributing
      xe-1/0/9                  Current   Fast periodic Collecting distributing
      xe-3/1/5                  Current   Fast periodic Collecting distributing
      xe-5/1/5                  Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/8       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/8     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/8       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/8     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/8                  Current   Fast periodic Collecting distributing
      xe-1/0/8                  Current   Fast periodic Collecting distributing
      xe-3/1/4                  Current   Fast periodic Collecting distributing
      xe-5/1/4                  Current   Fast periodic Collecting distributing
```

Copyright © 2016, Juniper Networks, Inc.

The output indicates that LACP has been set up correctly and is active at one end.

Related Documentation

- Understanding LACP on Chassis Clusters on page 2575
- Verifying LACP on Redundant Ethernet Interfaces on page 2579

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 117

Configuring Gigabit Ethernet Physical Interface Modules

- Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 2587
- Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 2589
- Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 2595
- Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 2598

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for Gigabit and Fast Ethernet connections. Gigabit Ethernet SFP Mini-PIMs can be used in copper and optical environments to provide maximum flexibility when upgrading from an existing infrastructure to Metro Ethernet.

The 1-Port Gigabit Ethernet SFP Mini-PIM interfaces a single Gigabit Ethernet device or a network. It supports a variety of transceivers with data speeds of 10-Mbps/100-Mbps/1-Gbps with extended LAN or WAN connectivity.

Transceivers are hot-swappable.

This topic includes the following sections:

- Supported Features on page 2587
- Interface Names and Settings on page 2588
- Available Link Speeds and Modes on page 2588
- Link Settings on page 2589

Supported Features

The following features are supported on the 1-Port Gigabit Ethernet SFP Mini-PIM:

- 10-Mbps/100-Mbps/1-Gbps link speed
- Half-duplex/full-duplex support
- Autonegotiation
- Encapsulations

Copyright © 2016, Juniper Networks, Inc.

- Maximum transmission unit (MTU) size of 1514 bytes (default) and 9010 bytes (jumbo frames)
- Loopback
- Transceivers are hot-swappable

Interface Names and Settings

The following format is used to represent the 1-Port Gigabit Ethernet SFP Mini-PIM interface names:

type-fpc/pic/port

Where:

- type—Media type (ge)
- fpc—Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- pic—Number of the PIC on which the physical interface is located (0)
- port—Specific port on a PIC (0)

Examples: **ge-1/0/0** and **ge-2/0/0**

By default, the interfaces on the ports on the uplink module installed on the device are enabled. You can also specify the MTU size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9010. The default MTU size for Gigabit Ethernet interfaces is 1514.

Available Link Speeds and Modes

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link speeds:

- **10m**—Sets the link speed to 10 Mbps.
- **100m**—Sets the link speed to 100 Mbps.
- **1g**—Sets the link speed to 1 Gbps.

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link modes:

- **Full-duplex**—Allows bidirectional communication at a given point in time.
- **Half-duplex**—Allows single directional communication at a given point in time.

Copyright © 2016, Juniper Networks, Inc.

Link Settings

The 1-Port Gigabit Ethernet SFP Mini-PIM includes the following link settings:

- **auto-negotiation**—Enables autonegotiation of link mode and speed.

  NOTE: By default, autonegotiation is enabled. To disable autonegotiation, use **set gigether-options no-autonegotiation**

  We recommend enabling autonegotiation.

- **loopback**—Enables loopback.
- **no-auto-negotiation**—Disables autonegotiation of link mode and speed.
- **no-loopback**—Disables loopback.

By default a link speed of 1 Gbps in full-duplex mode is supported.

NOTE: On SRX340 High Memory devices, traffic might stop between the SRX340 device and the Cisco switch due to link mode mismatch. We recommend setting the same value to the autonegotiation parameters on both ends.

NOTE: On SRX300 devices, the link goes down when you upgrade FPGA on 1-Port Gigabit Ethernet SFP mini-PIM. As a workaround, run the **restart fpc** command and restart the FPC.

NOTE: On SRX300, SRX320, SRX340, and SRX550 devices, LACP is not supported on the 1-Port Gigabit Ethernet SFP Mini-PIM.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 2589

Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- Requirements on page 2590
- Overview on page 2590

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 2590
- Verification on page 2593

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 2546.

Overview

In this example, you configure the ge-2/0/0 interface, set the operating speed to 100 Mbps, and define a logical interface that you can connect to the 1-Port Gigabit Ethernet SFP Mini-PIM. You also set the MTU value to 9010 and set the link option to no-loopback.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-2/0/0 link-mode full-duplex speed 100m
set interface ge-2/0/0 gigether-options no-loopback
```

Configuring Physical Properties

GUI Step-by-Step Procedure

To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0** and then click **Edit**. A pop-up window appears.
3. In the Description box, type the description for the SFP Mini-PIM.
4. In the MTU box, type **9010**.
5. From the Speed list, select **100Mbps**.
6. From the Link-mode list, select **Full-duplex**.
7. Select the Enable Auto-negotiation checkbox.
8. Select the Enable Per Unit Scheduler checkbox.
9. Click **OK**

Disabling the Interface

GUI Step-by-Step Procedure

To disable the 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces** .

Copyright © 2016, Juniper Networks, Inc.

2. Under Interface, select **ge-2/0/0** and then click **Disable**.

Configuring Logical Properties

GUI Step-by-Step Procedure

To quickly configure the logical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0.0,** and then click **Add Logical Interface**. A pop-up window appears.
3. In the Unit box, type **0**.
4. In the Description box, type a description for the SFP Mini-PIM.
5. From the Zone list, select **untrust**.
6. To edit the family protocol type to the Mini-PIM interfaces, select the IPv4 tab, and then select **Enable address configuration**.
7. Click **Add,** and then type IPv4 address.
8. Click **OK**.

Editing Logical Properties

Step-by-Step Procedure

To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web:

1. Under Interface, select the logical interface added to the 1-Port Gigabit Ethernet SFP Mini-PIM and then click **Edit**. A pop-up window appears.
2. Under Interface, select **ge-2/0/0.0,** and then click **Edit Logical Interface**. A pop-up window appears.
3. From the Zone list, select **trust**.
4. To enable DHCP client on the interface, select the IPv4 tab and then select **Enable DHCP**.
5. Click **OK**.

NOTE: You cannot add or edit Description and Unit for a logical interface.

Deleting the Logical Interface

GUI Step-by-Step Procedure

To delete the logical interface of 1–Port Gigabit Ethernet SFP Mini-PIM using J-Web,

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0.0,** and then click **Delete**.

Copyright © 2016, Juniper Networks, Inc.

Configuring a 1-Port Gigabit Ethernet SFP Mini-PIM

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a 1-Port Gigabit Ethernet SFP Mini-PIM:

1. Configure the interface.

   ```
   [edit]
   user@host# edit interfaces ge-2/0/0
   ```

2. Set the operating link-mode full-duplex speed of 100 Mbps for the SFP Mini-PIM.

   ```
   [edit interfaces ge-2/0/0]
   user@host# set link-mode full-duplex speed 100m
   ```

3. Assign the MTU value.

   ```
   [edit interfaces ge-2/0/0]
   user@host# set mtu 9010
   ```

4. Add the logical interface.

   ```
   [edit interfaces ge-2/0/0]
   user@host# set unit 0 family inet address 14.1.1.1/24
   ```

5. Set the link options.

   ```
   [edit interfaces ge-2/0/0]
   user@host# set gigether-options no-loopback
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-2/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
   user@host# show interfaces ge-2/0/0
   mtu 9010;
   speed 100m;
   gigether-options {
   no-loopback;
}
unit 0 {
   family inet {
   14.1.1.1/24
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying That the Correct Hardware Is Installed on page 2593
- Verifying the FPC Status on page 2594
- Verifying the Interface Settings on page 2594

Verifying That the Correct Hardware Is Installed

Purpose Verify that the 1-Port Gigabit Ethernet SFP Mini-PIM is installed on the device.

Action From operational mode, enter the **show chassis hardware** command.

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AG0309AA0004      SRX240b
Routing Engine   REV 16   750-021792   VL3180            RE-SRX240B
  da0     999 MB  ST72682                                Nand Flash
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90  vendor 0x0409     uhub1
  usb0 (addr 3)  ST72682  High Speed Mode 64218 STMicroelectronics umass0
FPC 0                                                  FPC
  PIC 0                                                16x GE Base PIC
FPC 1                    750-023367   112009000278      FPC
  PIC 0                                                1x T1E1 mPIM
FPC 2            REV 00   750-03273   AABC5081          FPC
  PIC 0                                                1x GE High-Perf SFP mPIM

    Xcvr 0       REV 02   740-011612   9101465           SFP-T
FPC 4                    750-029145   122009000061      FPC
PIC 0                                                  1x GE SFP mPIM
    Xcvr 0       REV 01   740-011782   PBL0C3T           SFP-SX
Power Supply 0
```

Verify that the output contains the following values:

- **FPC 2, PIC 0** —1x GE High-Perf SFP mPIM
- **FPC 4, PIC 0** —1x GE SFP mPIM

NOTE: In the example shown above, the output for 1-Port SFP Mini-Physical Interface Module is displayed as 1X GE SFP mPIM and the output for 1-Port Gigabit Ethernet SFP Mini-Physical Interface Module is displayed as 1X GE High-Perf SFP mPIM.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The 1-Port GE SFP Mini-PIM is installed in the second slot of the device chassis; therefore the output displayed is 1x GE High-Perf SFP mPIM and the Flexible PIC Concentrator (FPC) used here is fpc 2.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; therefore the output displayed is 1x GE SFP mPIM and Flexible PIC Concentrator (FPC) used here is fpc 4.

Verifying the FPC Status

Purpose Verify the FPC status.

Action From operational mode, enter the **show chassis fpc** command.

```
show@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online          -------------------- CPU less FPC --------------------
  1  Online          -------------------- CPU less FPC --------------------
  2  Online          -------------------- CPU less FPC --------------------
  3  Empty
  4  Online          -------------------- CPU less FPC --------------------
```

The output should show the FPC status as online.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; the output shows the FPC status for slot 4 as online.

The 1-Port Gigabit Ethernet SFP Mini-PIM is installed in the second slot of the device chassis; the output shows the FPC status for slot 2 as online.

Verifying the Interface Settings

Purpose Verify that the interface is configured as expected.

Action From operational mode, enter the **show interface ge-2/0/0** command.

```
user@host# run show interfaces ge-2/0/0
Physical interface: ge-2/0/0, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 552
  Link-level type: Ethernet, MTU: 9010, Link-mode: Full-duplex, Speed: 100mbps,
BPDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:22:83:99:ac:f2, Hardware address: 00:22:83:99:ac:f2
  Last flapped   : 2010-08-17 12:20:33 UTC (00:00:20 ago)
  Input rate     : 0 bps (0 pps)
```


Copyright © 2016, Juniper Networks, Inc.

```
Output rate     : 0 bps (0 pps)
Active alarms   : None
Active defects  : None

Logical interface ge-2/0/0.0 (Index 88) (SNMP ifIndex 557)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 108
  Output packets: 1
  Security: Zone: Null
  Protocol inet, MTU: 8996
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 14.1.1/24, Local: 14.1.1.1, Broadcast: 14.1.1.255
```

Verify the following information in the command output:

- Physical interface—ge-2/0/0, Enabled, Physical link is Up
- MTU—9010; Link-mode—Full-duplex
- Speed—100 Mbps
- Loopback—Disabled

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 2587
- *Example: Configuring the Device as a DHCP Client*

Understanding the 2-Port 10-Gigabit Ethernet XPIM

The 10-Gigabit Ethernet (also known as 10GBASE-T or IEEE 802.3an) is a telecommunication technology that offers data speeds up to 10 billion bits per second over unshielded or shielded twisted pair cables.

The 2-Port 10-Gigabit Ethernet Physical Interface Module (XPIM) is a 2 x 10GBASE-T / SFP+ XPIM line card. (SFP+ is a fiber optic transceiver module designed for 10-Gigabit Ethernet and 8.5 Gbps-fiber channel systems.) The 2-Port 10-Gigabit Ethernet XPIM provides a front-end interface connection that includes the following ports:

- 2 X copper ports. The copper ports support 10GBASE-T running with CAT6A or CAT7 Ethernet cable for up to 100 meters.
- 2 X fiber (SFP+) ports. The fiber ports support SFP+ multiple 10G modules.

The 2-Port 10-Gigabit Ethernet XPIM provides interconnects for LANs, WANs, and metropolitan area networks (MANs). The XPIM provides multiple service levels (1-Gigabit Ethernet to 10-Gigabit Ethernet in increments) and a single connection option for a wide range of customer needs and applications.

NOTE: By default, the 2-Port 10-Gigabit Ethernet XPIM ports comes up in fiber mode, while autonegotiation is not supported.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Supported Features on page 2596
- Interface Names and Settings on page 2596
- Copper and Fiber Operating Modes on page 2597
- Link Speeds on page 2597
- Link Settings on page 2597

Supported Features

The following features are supported on the 2-Port 10-Gigabit Ethernet XPIM:

- Multiple SFP+ 10G modules and the following SFP modules:
  - SFPP-10GE-SR
  - SFPP-10GE-LR
  - SFPP-10GE-ER
  - SFPP-10GE-LRM
- Copper TWIN-AX 1M and Copper TWIN-AX 3M
- Online Insertion and Removal (OIR) functionality
- Link speeds of up to 10-Gbps
- Full-duplex and half-duplex modes
- Flow control
- Autonegotiation and autosensing
- Quality of service (QoS)

Interface Names and Settings

The following format is used to represent the 2-Port 10-Gigabit Ethernet XPIM interface names:

type-fpc/pic/port

Where:

- type — Media type (xe)
- fpc — Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- pic — Number of the PIC on which the physical interface is located (0)
- port — Specific port on a PIC (0 or 1)

By default, the interfaces (for example, **xe-6/0/0** or **xe-2/0/0**) on the ports on the uplink module installed on the device are enabled. You can also specify the maximum

Copyright © 2016, Juniper Networks, Inc.

transmission unit (MTU) size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9192. The default MTU for Gigabit Ethernet interfaces is 1514.

Copper and Fiber Operating Modes

On the 2-Port 10-Gigabit Ethernet XPIM, one copper port and one fiber port is grouped together as port 0, and another copper port and fiber port are grouped as port 1. Only two ports can be active at the same time (one port from port 0 and another port from port 1).

The 2-Port 10-Gigabit Ethernet XPIM can be configured to operate in two copper mode, two fiber mode, or mixed mode (one copper and one fiber). In mixed mode, the two ports should be from different port groups (one port from port 1 and the other from port 2).

Link Speeds

The 2-Port 10-Gigabit Ethernet XPIM ports support the following link speeds for copper and fiber:

- Copper—10/100/1000 Mbps or 10Gbps (full duplex). Half-duplex is only for 10/100 Mbps.
- Fiber—1000 Mbps or 10 Gbps (full duplex). Half-duplex mode is not supported.

To set the link speeds, use the following options:

- **10m**—Sets the link speed to 10 Mbps.
- **10g**—Sets the link speed to 10 Gbps.
- **100m**—Sets the link speed to 100 Mbps.
- **1g**—Sets the link speed to 1 Gbps.

Link Settings

The 2-Port 10-Gigabit Ethernet XPIM includes the following link settings:

- **802.3ad**—Specifies an aggregated Ethernet bundle.
- **auto-negotiation**—Enables autonegotiation of flow control, link mode, and speed.
- **loopback**—Enables loopback.
- **no-auto-negotiation**—Disables autonegotiation of flow control, link mode, and speed.
- **no-loopback**—Disables loopback.

By default, flow control is enabled on all ports, a link speed of 10 Gbps in full duplex is supported, autonegotiation is disabled on the fiber ports, and autonegotiation is enabled on copper ports.

NOTE: Autonegotiation is not supported when the 2-Port 10-Gigabit Ethernet XPIM is operating in fiber mode at a link speed of 10 Gbps.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 2598

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- Requirements on page 2598
- Overview on page 2598
- Configuration on page 2598
- Verification on page 2600

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 2546.

Overview

In this example, you configure the xe-6/0/0 interface, set the operating mode to copper mode, set the operating speed to 10 Gbps, and define a logical interface that you can connect to the 2-Port 10-Gigabit Ethernet XPIM. Additionally, you set the MTU value to 1514, set the link option to no loopback, and enable the interface.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-6/0/0 media-type copper speed 10g unit 0 family inet mtu 1514
set interface xe-6/0/0 gigether-options no-loopback
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a 2-Port 10-Gigabit Ethernet XPIM:

1. Configure the interface.

   ```
   [edit]
   user@host# edit interfaces xe-6/0/0
   ```

2. Configure the operating mode.

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces xe-6/0/0]
user@host# set media-type copper
```

3. Set the operating speed for the XPIM.

```
[edit interfaces xe-6/0/0]
user@host# set speed 10g
```

4. Add the logical interface.

```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet
```

5. Assign the physical interface MTU value.

```
[edit interfaces xe-6/0/0]
user@host# set interface xe-6/0/0 mtu 1514
```

6. Assign the logical interface MTU value.

```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet mtu 1500
```

7. Set the link options.

```
[edit interfaces xe-6/0/0]
user@host# set gigether-options no-loopback
```

8. Disable the interface.

```
[edit interfaces xe-6/0/0]
user@host# set disable
```

9. Enable the interface.

```
[edit interfaces xe-6/0/0]
user@host# delete disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces xe-6/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
   user@host# show interfaces xe-6/0/0
   speed 10g;
   media-type copper;
   gigether-options {
   no-loopback;
}
unit 0 {
   family inet {
   mtu 1514;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying That the Correct Hardware Is Installed on page 2600
- Verifying the FPC Status on page 2600
- Verifying the Interface Settings on page 2601

Verifying That the Correct Hardware Is Installed

Purpose Verify that the 2-Port 10-Gigabit Ethernet XPIM is installed on the device.

Action From operational mode, enter the **show chassis hardware** command.

```
Hardware inventory:
Item      Version   Part number  Serial number   Description
Chassis             AJ0309AC0047   SRX650
Midplane    REV 04   710-023875   TV3993
System IO    REV 04   710-023209   TV4035     SRXSME System IO
Routing Engine  REV 01   710-023224   DT5109     RE-SRXSME-SRE6
FPC 0                    FPC
PIC 0                    4x GE Base PIC
FPC 2                    FPC
PIC 0                    2x 10G gPIM
FPC 6                    FPC
PIC 0                    2x 10G gPIM
Power Supply 0  REV 01   740-024283   TA00049WSSSS   PS 645W AC
```

Verify that the output contains the following values:

- **FPC 2 , PIC 0**—2x 10G gPIM
- **FPC 6, PIC 0**—2x 10G gPIM

Verifying the FPC Status

Purpose Verify the FPC status.

Action From operational mode, enter the **show chassis fpc** command.

```
   Temp   CPU Utilization  (%)  Memory   Utilization (%)
Slot State   (C)   Total Interrupt    DRAM (MB)  Heap Buffer
0 Online   -------------------- CPU less FPC --------------------
1 Empty
2 Online   -------------------- CPU less FPC --------------------
3 Empty
4 Empty
5 Empty
6 Online   -------------------- CPU less FPC --------------------
7 Empty
8 Empty
```

The output should display FPC status as online.

Copyright © 2016, Juniper Networks, Inc.

Verifying the Interface Settings

Purpose Verify that the interface is configured as expected.

Action From operational mode, enter the **show interface xe-6/0/0** command.

```
Physical interface: xe-6/0/0, Enabled, Physical link is Up
 Interface index: 144, SNMP ifIndex: 501
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 10Gbps,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 6 Copyright © 2010, Juniper Networks, Inc.
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Current address: 00:1f:12:e0:80:a8, Hardware address: 00:1f:12:e0:80:a8
 Last flapped : 1970-01-01 00:34:22 PST (07:26:29 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None

Logical interface xe-6/0/0.0 (Index 72) (SNMP ifIndex 503)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 25
 Output packets: 25
 Security: Zone: HOST
 Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
 ospf pgm pim rip router-discovery rsvp sap vrrp
 Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.10, Broadcast: 10.10.10.255
```

Verify the following information in the command output:

- Physical interface—xe-6/0/0, Enabled, Physical link is Up
- MTU—1514
- Link mode—Full duplex
- Speed—10 Gbps
- Loopback—Disabled
- Flow control—Enabled

Related Documentation

- Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 2595
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 118

Configuring Ethernet OAM Link Fault Management

- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 2603
- Example: Configuring Ethernet OAM Link Fault Management on page 2605

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

The Ethernet interfaces on SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

NOTE: For SRX550 devices, LFM is supported only on devices that have 16-port or 24-port GPIMs.

The following OAM LFM features are supported:

- Discovery and link monitoring—The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The device performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

  You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote fault detection—Remote fault detection uses flags and events. Flags convey Link Fault (a loss of signal), Dying Gasp (an unrecoverable condition such as a power failure), and Critical Event (an unspecified vendor-specific critical event). You can

Copyright © 2016, Juniper Networks, Inc.

specify the periodic OAM PDU sending interval for fault detection. SRX Series devices use the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote loopback—Remote loopback mode ensures link quality between the device and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

  Junos OS can place a remote data terminal equipment (DTE) into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

Table 273 lists the interfaces modes supported.

Table 273: Supported Interface Modes

Interfaces	Mode
Physical interface (fe/ge)	Family • ccc • ethernet-switching • inet6 • inet • iso • mpls • tcc
	IFD encapsulations • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • ethernet-tcc • extended-vlan-tcc

Copyright © 2016, Juniper Networks, Inc.

Table 273: Supported Interface Modes *(continued)*

Interfaces	Mode
Aggregated Ethernet interface (Static or LACP lag)	Family • ethernet-switching • inet • mpls • iso • inet6
	IFD encapsulations • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • vlan-ccc

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on page 2605

Example: Configuring Ethernet OAM Link Fault Management

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet or Fast Ethernet interface:

- Requirements on page 2605
- Overview on page 2606
- Configuration on page 2606
- Verification on page 2608

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 R2 or later for SRX Series Services Gateways
- Any two models of SRX Series devices connected directly

Before you begin:

- Establish basic connectivity. See the Getting Started Guide for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 2546.

Copyright © 2016, Juniper Networks, Inc.

- Ensure that you configure the interfaces as per the interface modules listed in "Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways" on page 2603

Overview

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two SRX Series devices connected directly. Before you begin configuring Ethernet OAM LFM on these two devices, connect the two devices directly through supported interfaces. See "Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways" on page 2603.

Figure 119 shows the topology used in this example.

Figure 119: Ethernet LFM with SRX Series Devices

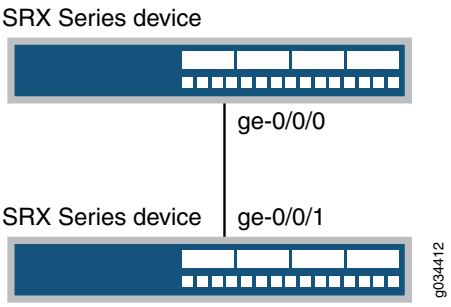

NOTE: For more information about configuring Ethernet OAM Link Fault Management, see Junos® OS Ethernet Interfaces.

Configuration

To configure Ethernet OAM LFM, perform these tasks:

- Configuring Ethernet OAM Link Fault Management on Device 1 on page 2606
- Configuring Ethernet OAM Link Fault Management on Device 2 on page 2607

Configuring Ethernet OAM Link Fault Management on Device 1

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols oam ethernet link-fault-management interface ge-0/0/0
set protocols oam ethernet link-fault-management interface ge-0/0/0 link-discovery
    active
set protocols oam ethernet link-fault-management interface ge-0/0/0 pdu-interval 800
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Ethernet OAM LFM on device 1:

1. Enable IEEE 802.3ah OAM support.

```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0
```

2. Specify that the interface initiates the discovery process.

```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0 link-discovery active
```

3. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.

```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface pdu-interval 800
```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@device1# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/0 {
          pdu-interval 800;
          link-discovery active;
        }
      }
    }
  }
}
```

Configuring Ethernet OAM Link Fault Management on Device 2

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols oam ethernet link-fault-management interface ge-0/0/1
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-interval 800
set protocols oam ethernet link-fault-management interface ge-0/0/1 negotiation-options
  allow-remote-loopback
```

Step-by-Step Procedure

To configure Ethernet OAM LFM on device 2:

1. Enable OAM on the peer interface.

```
[edit protocols oam ethernet link-fault-management]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@device2# set interface ge-0/0/1
```

2. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 pdu-interval 800
```

3. Enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Results From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@device2# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/1 {
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}
```

Verification

Verify the OAM LFM Configuration

Purpose Verify that OAM LFM is configured properly.

Action From operational mode, enter the **show oam ethernet link-fault-management** command.

```
user@device1>show oam ethernet link-fault-management
```

```
Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported
```

Meaning The output displays the MAC address and the discovery state is **Send Any** if OAM LFM has been configured properly.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 2603

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 119

Configuring Power over Ethernet

- Understanding Power over Ethernet on page 2611
- Example: Configuring PoE on All Interfaces on page 2613
- Example: Configuring PoE on an Individual Interface on page 2615
- Example: Disabling a PoE Interface on page 2618

Understanding Power over Ethernet

Power over Ethernet (PoE) is the implementation of the IEEE 802.3 AF and IEEE 802.3 AT standards that allow both data and electrical power to pass over a copper Ethernet LAN cable.

The SRX Series devices support PoE on Ethernet ports. PoE ports transfer electrical power and data to remote devices over standard twisted-pair cable in an Ethernet network. PoE ports allow you to plug in devices that require both network connectivity and electrical power, such as VoIP and IP phones and wireless LAN access points.

You can configure the SRX Series device to act as power sourcing equipment (PSE), supplying power to powered devices that are connected on designated ports.

This topic contains the following sections:

- SRX Series Services Gateway PoE Specifications on page 2611
- PoE Classes and Power Ratings on page 2612
- PoE Options on page 2613

SRX Series Services Gateway PoE Specifications

Table 274 lists the PoE specifications for the SRX320 and SRX340 devices

Table 274: PoE Specifications for the SRX320 and SRX340 Devices

Specifications	For SRX320 Device	For SRX340 Device
Supported standards	• IEEE 802.3 AF • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE+) • Legacy (pre-standards)

Copyright © 2016, Juniper Networks, Inc.

Table 274: PoE Specifications for the SRX320 and SRX340 Devices *(continued)*

Specifications	For SRX320 Device	For SRX340 Device
Supported ports	Supported on two Gigabit Ethernet ports and two Fast Ethernet ports (**ge-0/0/0**, **ge-0/0/1**, **fe-0/0/2**, and **fe-0/0/3**).	Supported on all 16 Gigabit Ethernet ports (**ge-0/0/0** to **ge-0/0/15**).
Total PoE power sourcing capacity	50 W	150 W
Default per port power limit	15.4 W	15.4 W
Maximum per port power limit	30 W	30 W
Power management modes	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.

PoE Classes and Power Ratings

A powered device is classified based on the maximum power that it draws across all input voltages and operational modes. When class-based power management mode is configured on the SRX Series devices, power is allocated taking into account the maximum power ratings defined for the different classes of devices.

Table 275 lists the classes and their power ratings as specified by the IEEE standards.

Table 275: SRX Series Devices PoE Specifications

Class	Usage	Minimum Power Levels Output from PoE Port
0	Default	15.4 W
1	Optional	4.0 W
2	Optional	7.0 W
3	Optional	15.4 W
4	Reserved	Class 4 power devices are eligible to receive power up to 30 W according to IEEE standards.

Copyright © 2016, Juniper Networks, Inc.

PoE Options

When configuring PoE, you must enable the PoE interface in order for the port to provide power to a connected, powered device. In addition, you can configure the following PoE features:

- Port priority—Sets port priority. When it is not possible to maintain power to all connected ports, lower priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower priority port will be powered off automatically if available power is insufficient to power on the higher priority port. (For the ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.)
- Maximum available wattage power available to a port—Sets the maximum amount of power that can be supplied to the port. The default wattage per port is 15.4 watts.
- PoE power consumption logging—Allows logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default. The default telemetry duration is 1 hour. The default telemetry interval is 5 minutes.
- PoE power management mode—Has two modes:
  - Class—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE standards.
  - Static—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port.
- Reserve power—Reserves the specified amount of power for the gateway in case of a spike in PoE consumption. The default is 0.

Related Documentation

- Understanding Ethernet Interfaces on page 2541
- Example: Configuring PoE on All Interfaces on page 2613
- Example: Configuring PoE on an Individual Interface on page 2615
- Example: Disabling a PoE Interface on page 2618

Example: Configuring PoE on All Interfaces

This example shows how to configure PoE on all interfaces.

- Requirements on page 2614
- Overview on page 2614
- Configuration on page 2614
- Verification on page 2615

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, configure Ethernet interfaces. See "Example: Creating an Ethernet Interface" on page 2546.

Overview

This example shows how to configure PoE on all interfaces on a device. In this example, you set the power port priority to low and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetries settings, and you set the PoE management mode to static. Finally, you set the reserved power consumption to 15 watts in case of a spike in PoE consumption.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set poe interface all priority low maximum-power 15.4 telemetries
set poe management static guard-band 15
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PoE on all interfaces:

1. Enable PoE.

   ```
   [edit]
   user@host# edit poe interface all
   ```

2. Set the power port priority.

   ```
   [edit poe interface all]
   user@host# set priority low
   ```

3. Set the maximum PoE wattage available for a port.

   ```
   [edit poe interface all]
   user@host# set maximum-power 15.4
   ```

4. Enable logging of PoE power consumption.

   ```
   [edit poe interface all]
   user@host# set telemetries
   ```

5. Set the PoE management mode.

   ```
   [edit]
   user@host# set poe management static
   ```

6. Reserve power wattage in case of a spike in PoE consumption.

   ```
   [edit]
   user@host# set poe guard-band 15
   ```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show poe interface all** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show poe interface all
  priority low;
  maximum-power 15.4;
telemetries;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the Status of PoE Interfaces

Purpose Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used here is the SRX340 Services Gateway.)

Action From operational mode, enter the **show poe interface all** command.

```
user@host> show poe interface all
```

```
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      Searching   15.4W     Low      0.0W                0
 ge-0/0/1 Enabled      Powered-up  15.4W     High     6.6W                0
 ge-0/0/2 Disabled     Disabled    15.4W     Low      0.0W                0
 ge-0/0/3 Disabled     Disabled    15.4W     Low      0.0W                0
```

The **show poe interface all** command lists PoE interfaces configured on the SRX 240 device, including information on status, priority, power consumption, and class. This output shows that the device has four PoE interfaces of which two are enabled with default values. One port has a device connected that is drawing power within expected limits.

Related Documentation

- Understanding Power over Ethernet on page 2611
- Example: Configuring PoE on an Individual Interface on page 2615
- Example: Disabling a PoE Interface on page 2618

Example: Configuring PoE on an Individual Interface

This example shows how to configure PoE on an individual interface.

- Requirements on page 2616
- Overview on page 2616
- Configuration on page 2616
- Verification on page 2617

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure Ethernet interfaces. See "Example: Creating an Ethernet Interface" on page 2546.
- Configure PoE on all interfaces. See "Example: Configuring PoE on All Interfaces" on page 2613.

Overview

This example shows how to configure PoE on the ge-0/0/0 interface. In this example, you set the power port priority to high and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetries settings, and you set the PoE management mode to static. Finally, you set the reserved power to 15 watts in case of a spike in PoE consumption.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set poe interface ge-0/0/0 priority high maximum-power 15.4 telemetries
set poe management static guard-band 15
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PoE:

1. Enable PoE.

   ```
   [edit]
   user@host# edit poe interface ge-0/0/0
   ```

2. Set the power port priority.

   ```
   [edit poe interface ge-0/0/0]
   user@host# set priority high
   ```

3. Set the maximum PoE wattage available for a port.

   ```
   [edit poe interface ge-0/0/0]
   user@host# set maximum power 15.4
   ```

4. Enable logging of PoE power consumption.

   ```
   [edit poe interface ge-0/0/0]
   user@host# set telemetries
   ```

5. Set the PoE management mode.

   ```
   [edit]
   user@host# set poe management static
   ```


Copyright © 2016, Juniper Networks, Inc.

6. Reserve power wattage in case of a spike in PoE consumption.

```
[edit]
user@host# set poe guard-band 15
```

Results From configuration mode, confirm your configuration by entering the **show poe interface ge-0/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show poe interface ge-0/0/0
  priority high;
  maximum-power 15.4;
telemetries;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying the Status of PoE Interfaces on page 2617
- Verifying the Telemetry Data (History) for the Specified Interface on page 2617
- Verifying PoE Global Parameters on page 2618

Verifying the Status of PoE Interfaces

Purpose Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used in this example is the SRX340 Services Gateway.)

Action From operational mode, enter the **show poe interface ge-0/0/1** command.

```
user@host> show poe interface ge-0/0/1
PoE interface status:
PoE interface                : ge-0/0/1
Administrative status        : Enabled
Operational status           : Powered-up
Power limit on the interface : 15.4 W
Priority                     : High
Power consumed               : 6.6 W
Class of power device        : 0
```

The **show poe interface ge-0/0/1** command lists PoE interfaces configured on the SRX340 device, with their status, priority, power consumption, and class.

Verifying the Telemetry Data (History) for the Specified Interface

Purpose Verify the PoE interface's power consumption over a specified period.

Action From operational mode, enter the **show poe telemetries interface** command.

For all records:

```
user@host> show poe telemetries interface ge-0/0/1 all
```

Copyright © 2016, Juniper Networks, Inc.

```
Sl No Timestamp Power Voltage
1 Fri Jan 04 11:41:15 2009 5.1 W 47.3 V
2 Fri Jan 04 11:40:15 2009 5.1 W 47.3 V
3 Fri Jan 04 11:39:15 2009 5.1 W 47.3 V
4 Fri Jan 04 11:38:15 2009 0.0 W 0.0 V
5 Fri Jan 04 11:37:15 2009 0.0 W 0.0 V
6 Fri Jan 04 11:36:15 2009 6.6 W 47.2 V
7 Fri Jan 04 11:35:15 2009 6.6 W 47.2 V
```

For a specific number of records:

```
user@host> show poe telemetries interface ge-0/0/1 5
Sl No Timestamp Power Voltage
1 Fri Jan 04 11:31:15 2009 6.6 W 47.2 V
2 Fri Jan 04 11:30:15 2009 6.6 W 47.2 V
3 Fri Jan 04 11:29:15 2009 6.6 W 47.2 V
4 Fri Jan 04 11:28:15 2009 6.6 W 47.2 V
5 Fri Jan 04 11:27:15 2009 6.6 W 47.2 V
```

The telemetry status displays the power consumption history for the specified interface, provided telemetry has been configured for that interface.

Verifying PoE Global Parameters

Purpose Verify global parameters such as guard band, power limit, and power consumption.

Action From operational mode, enter the **show poe controller** command.

```
user@host> show poe controller
Controller  Maximum   Power          Guard band  Management
index       power     consumption
   0        150.0 W    0.0 W         0 W         Static
```

The **show poe controller** command lists the global parameters configured on the SRX Series device such as controller index, maximum power, power consumption, guard band, and management mode along with their status.

Related Documentation

- Understanding Power over Ethernet on page 2611
- Example: Configuring PoE on All Interfaces on page 2613
- Example: Disabling a PoE Interface on page 2618

Example: Disabling a PoE Interface

This example shows how to disable PoE on all interfaces or on a specific interface.

- Requirements on page 2619
- Overview on page 2619
- Configuration on page 2619
- Verification on page 2619

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure PoE on all interfaces. See "Example: Configuring PoE on All Interfaces" on page 2613.
- Configure PoE on an individual interface. See "Example: Configuring PoE on an Individual Interface" on page 2615.

Overview

In this example, you disable PoE on all interfaces and on a specific interface, which in this case is ge-0/0/0.

Configuration

Step-by-Step Procedure

To disable PoE on interfaces:

1. Disable PoE on all interfaces.

   ```
   [edit]
   user@host# set poe interface all disable
   ```

2. Disable PoE on a specific interface.

   ```
   [edit]
   user@host# set poe interface ge-0/0/0 disable
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show poe interface** command.

Related Documentation

- Understanding Power over Ethernet on page 2611

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 38

Configuring Interface Encapsulation

- Interface Encapsulation Overview on page 2623
- Configuring Point-to-Point Protocol over Ethernet on page 2631
- Configuring PPPoE-Based Radio-to-Router Protocol on page 2649
- Configuring R2CP Radio-to-Router Protocol on page 2657

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 120

Interface Encapsulation Overview

- Understanding Physical Encapsulation on an Interface on page 2623
- Understanding Frame Relay Encapsulation on an Interface on page 2624
- Understanding Point-to-Point Protocol on page 2625
- Understanding High-Level Data Link Control on page 2628

Understanding Physical Encapsulation on an Interface

Encapsulation is the process by which a lower level protocol accepts a message from a higher level protocol and places it in the data portion of the lower level frame. As a result, datagrams transmitted through a physical network have a sequence of headers: the first header for the physical network (or Data Link Layer) protocol, the second header for the Network Layer protocol (IP, for example), the third header for the Transport Layer protocol, and so on.

The following encapsulation protocols are supported on physical interfaces:

- Frame Relay Encapsulation. See "Understanding Frame Relay Encapsulation on an Interface" on page 2624.
- Point-to-Point Protocol. See "Understanding Point-to-Point Protocol" on page 2625.
- Point-to-Point Protocol over Ethernet. See "Understanding Point-to-Point Protocol over Ethernet" on page 2631.
- High-Level Data Link Control. See "Understanding High-Level Data Link Control" on page 2628.

Related Documentation

- Understanding Interfaces on page 2441
- Understanding Frame Relay Encapsulation on an Interface on page 2624
- Understanding Point-to-Point Protocol on page 2625
- Understanding High-Level Data Link Control on page 2628

Copyright © 2016, Juniper Networks, Inc.

Understanding Frame Relay Encapsulation on an Interface

The Frame Relay packet-switching protocol operates at the Physical Layer and Data Link Layer in a network to optimize packet transmissions by creating virtual circuits between hosts. Figure 120 shows a typical Frame Relay network.

Figure 120: Frame Relay Network

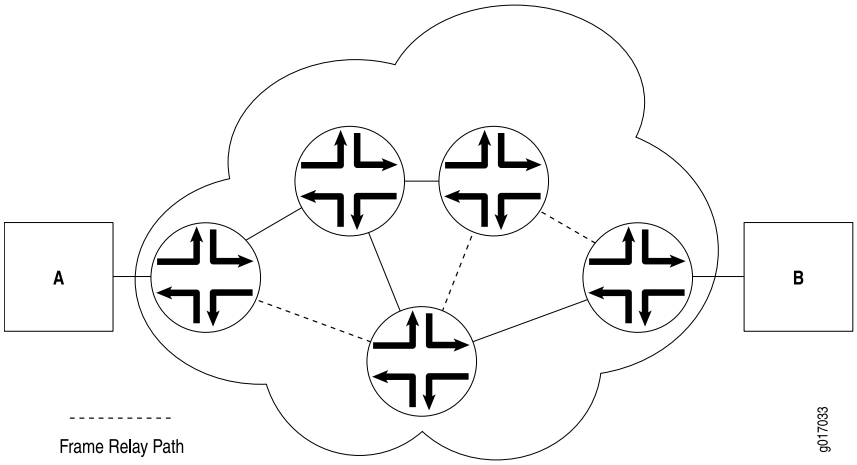

Figure 120 shows multiple paths from Host A to Host B. In a typical routed network, traffic is sent from device to device with each device making routing decisions based on its own routing table. In a packet-switched network, the paths are predefined. Devices switch a packet through the network according to predetermined next-hops established when the virtual circuit is set up.

This topic contains the following sections:

- Virtual Circuits on page 2624
- Switched and Permanent Virtual Circuits on page 2624
- Data-Link Connection Identifiers on page 2625
- Congestion Control and Discard Eligibility on page 2625

Virtual Circuits

A virtual circuit is a bidirectional path between two hosts in a network. Frame Relay virtual circuits are logical connections between two hosts that are established either by a call setup mechanism or by an explicit configuration.

A virtual circuit created through a call setup mechanism is known as a switched virtual circuit (SVC). A virtual circuit created through an explicit configuration is called a permanent virtual circuit (PVC).

Switched and Permanent Virtual Circuits

Before data can be transmitted across an SVC, a signaling protocol like ISDN must set up a call by the exchange of setup messages across the network. When a connection is established, data is transmitted across the SVC. After data transmission, the circuit is

Copyright © 2016, Juniper Networks, Inc.

torn down and the connection is lost. For additional traffic to pass between the same two hosts, a subsequent SVC must be established, maintained, and terminated.

Because PVCs are explicitly configured, they do not require the setup and teardown of SVCs. Data can be switched across the PVC whenever a host is ready to transmit. SVCs are useful in networks where data transmission is sporadic and a permanent circuit is not needed.

Data-Link Connection Identifiers

An established virtual circuit is identified by a data-link connection identifier (DLCI). The DLCI is a value from 16 through 1022. (Values 1 through 15 are reserved.) The DLCI uniquely identifies a virtual circuit locally so that devices can switch packets to the appropriate next-hop address in the circuit. Multiple paths that pass through the same transit devices have different DLCIs and associated next-hop addresses.

Congestion Control and Discard Eligibility

Frame Relay uses the following types of congestion notification to control traffic within a Frame Relay network. Both are controlled by a single bit in the Frame Relay header.

- Forward explicit congestion notification (FECN)
- Backward explicit congestion notification (BECN)

Traffic congestion is typically defined in the buffer queues on a device. When the queues reach a predefined level of saturation, traffic is determined to be congested. When traffic congestion occurs in a virtual circuit, the device experiencing congestion sets the congestion bits in the Frame Relay header to 1. As a result, transmitted traffic has the FECN bit set to 1, and return traffic on the same virtual circuit has the BECN bit set to 1.

When the FECN and BECN bits are set to 1, they provide a congestion notification to the source and destination devices. The devices can respond in either of two ways: to control traffic on the circuit by sending it through other routes, or to reduce the load on the circuit by discarding packets.

If devices discard packets as a means of congestion (flow) control, Frame Relay uses the discard eligibility (DE) bit to give preference to some packets in discard decisions. A DE value of 1 indicates that the frame is of lower importance than other frames and more likely to be dropped during congestion. Critical data (such as signaling protocol messages) without the DE bit set is less likely to be dropped.

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623

Understanding Point-to-Point Protocol

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. PPP is made up of three primary components:

- Link Control Protocol (LCP)—Establishes working connections between two points.
- Authentication protocol—Enables secure connections between two points.

Copyright © 2016, Juniper Networks, Inc.

- Network control protocol (NCP)—Initializes the PPP protocol stack to handle multiple Network Layer protocols, such as IPv4, IPv6, and Connectionless Network Protocol (CLNP).

This topic contains the following sections:

- Link Control Protocol on page 2626
- PPP Authentication on page 2626
- Network Control Protocols on page 2627
- Magic Numbers on page 2627
- CSU/DSU Devices on page 2628

Link Control Protocol

LCP is responsible for establishing, maintaining, and tearing down a connection between two endpoints. LCP also tests the link and determines whether it is active. LCP establishes a point-to-point connection as follows:

1. LCP must first detect a clocking signal on each endpoint. However, because the clocking signal can be generated by a network clock and shared with devices on the network, the presence of a clocking signal is only a preliminary indication that the link might be functioning.
2. When a clocking signal is detected, a PPP host begins transmitting PPP Configure-Request packets.
3. If the remote endpoint on the point-to-point link receives the Configure-Request packet, it transmits a Configure-Acknowledgement packet to the source of the request.
4. After receiving the acknowledgement, the initiating endpoint identifies the link as established. At the same time, the remote endpoint sends its own request packets and processes the acknowledgement packets. In a functioning network, both endpoints treat the connection as established.

During connection establishment, LCP also negotiates connection parameters such as FCS and HDLC framing. By default, PPP uses a 16-bit FCS, but you can configure PPP to use either a 32-bit FCS or a 0-bit FCS (no FCS). Alternatively, you can enable HDLC encapsulation across the PPP connection.

After a connection is established, PPP hosts generate Echo-Request and Echo-Response packets to maintain a PPP link.

PPP Authentication

PPP's authentication layer uses a protocol to help ensure that the endpoint of a PPP link is a valid device. Authentication protocols include the Password Authentication Protocol (PAP), the Extensible Authentication Protocol (EAP), and the Challenge Handshake Authentication Protocol (CHAP). CHAP is the most commonly used.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Support for user id and the password to comply with full ASCII character set is supported through RFC 2486.

The user can enable or disable the RFC 2486 support under the PPP options. The RFC 2486 is disabled by default, and enable the support globally use the command set access ppp-options compliance rfc 2486".

CHAP ensures secure connections across PPP links. After a PPP link is established by LCP, the PPP hosts at either end of the link initiate a three-way CHAP handshake. Two separate CHAP handshakes are required before both sides identify the PPP link as established.

CHAP configuration requires each endpoint on a PPP link to use a shared secret (password) to authenticate challenges. The shared secret is never transmitted over the wire. Instead, the hosts on the PPP connection exchange information that enables both to determine that they share the same secret. Challenges consist of a hash function calculated from the secret, a numeric identifier, and a randomly chosen challenge value that changes with each challenge. If the response value matches the challenge value, authentication is successful. Because the secret is never transmitted and is required to calculate the challenge response, CHAP is considered very secure.

PAP authentication protocol uses a simple two-way handshake to establish identity. PAP is used after the link establishment phase (LCP up), during the authentication phase. Junos OS can support PAP in one direction (egress or ingress), and CHAP in the other.

Network Control Protocols

After authentication is completed, the PPP connection is fully established. At this point, any higher level protocols (for example, IP protocols) can initialize and perform their own negotiations and authentication.

PPP NCPs include support for the following protocols. IPCP and IPv6CP are the most widely used on SRX Series devices.

- IPCP—IP Control Protocol
- IPv6CP—IPv6 Control Protocol
- OSINLCP—OSI Network Layer Control Protocol (includes IS-IS, ES-IS, CLNP, and IDRP)

Magic Numbers

Hosts running PPP can create "magic" numbers for diagnosing the health of a connection. A PPP host generates a random 32-bit number and sends it to the remote endpoint during LCP negotiation and echo exchanges.

In a typical network, each host's magic number is different. A magic number mismatch in an LCP message informs a host that the connection is not in loopback mode and traffic is being exchanged bidirectionally. If the magic number in the LCP message is the same as the configured magic number, the host determines that the connection is in loopback mode, with traffic looped back to the transmitting host.

Copyright © 2016, Juniper Networks, Inc.

Looping traffic back to the originating host is a valuable way to diagnose network health between the host and the loopback location. To enable loopback testing, telecommunications equipment typically supports channel service unit/data service unit (CSU/DSU) devices.

CSU/DSU Devices

A channel service unit (CSU) connects a terminal to a digital line. A data service unit (DSU) performs protective and diagnostic functions for a telecommunications line. Typically, the two devices are packaged as a single unit. A CSU/DSU device is required for both ends of a T1 or T3 connection, and the units at both ends must be set to the same communications standard.

A CSU/DSU device enables frames sent along a link to be looped back to the originating host. Receipt of the transmitted frames indicates that the link is functioning correctly up to the point of loopback. By configuring CSU/DSU devices to loop back at different points in a connection, network operators can diagnose and troubleshoot individual segments in a circuit.

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623

Understanding High-Level Data Link Control

High-Level Data Link Control (HDLC) is a bit-oriented, switched and nonswitched link-layer protocol. HDLC is widely used because it supports half-duplex and full-duplex connections, point-to-point and point-to-multipoint networks, and switched and nonswitched channels.

This topic contains the following sections:

- HDLC Stations on page 2628
- HDLC Operational Modes on page 2629

HDLC Stations

Nodes within a network running HDLC are called stations. HDLC supports three types of stations for data link control:

- Primary stations—Responsible for controlling the secondary and combined other stations on the link. Depending on the HDLC mode, the primary station is responsible for issuing acknowledgement packets to allow data transmission from secondary stations.
- Secondary stations—Controlled by the primary station. Under normal circumstances, secondary stations cannot control data transmission across the link with the primary station, are active only when requested by the primary station, and can respond to the primary station only (not to other secondary stations). All secondary station frames are response frames.
- Combined stations—A combination of primary and secondary stations. On an HDLC link, all combined stations can send and receive commands and responses without

Copyright © 2016, Juniper Networks, Inc.

any permission from any other stations on the link and cannot be controlled by any other station.

HDLC Operational Modes

HDLC runs in three separate modes:

- Normal Response Mode (NRM)—The primary station on the HDLC link initiates all information transfers with secondary stations. A secondary station on the link can transmit a response of one or more information frames only when it receives explicit permission from the primary station. When the last frame is transmitted, the secondary station must wait for explicit permission before it can transmit more frames.

  NRM is used most widely for point-to-multipoint links, in which a single primary station controls many secondary stations.

- Asynchronous Response Mode (ARM)—The secondary station can transmit either data or control traffic at any time, without explicit permission from the primary station. The primary station is responsible for error recovery and link setup, but the secondary station can transmit information at any time.

  ARM is used most commonly with point-to-point links, because it reduces the overhead on the link by eliminating the need for control packets.

- Asynchronous Balance Mode (ABM)—All stations are combined stations. Because no other station can control a combined station, all stations can transmit information without explicit permission from any other station. ABM is not a widely used HDLC mode.

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 121

Configuring Point-to-Point Protocol over Ethernet

- Understanding Point-to-Point Protocol over Ethernet on page 2631
- Understanding PPPoE Interfaces on page 2634
- Example: Configuring PPPoE Interfaces on page 2635
- Understanding PPPoE Ethernet Interfaces on page 2641
- Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 2641
- Understanding CHAP Authentication on a PPPoE Interface on page 2642
- Example: Configuring CHAP Authentication on a PPPoE Interface on page 2643
- Verifying Credit-Flow Control on page 2644
- Verifying PPPoE Interfaces on page 2645
- Verifying R2CP Interfaces on page 2646
- Displaying Statistics for PPPoE on page 2647
- Setting Tracing Options for PPPoE on page 2647

Understanding Point-to-Point Protocol over Ethernet

Point-to-Point Protocol over Ethernet (PPPoE) combines PPP, which typically runs over broadband connections, with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator. PPPoE enables service providers to maintain access control through PPP connections and also manage multiple hosts at a remote site.

PPPoE connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device—a Juniper Networks device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

To use PPPoE, you must initiate a PPPoE session, encapsulate Point-to-Point Protocol (PPP) packets over Ethernet, and configure the device as a PPPoE client. To provide a PPPoE connection, each PPP session must learn the Ethernet address of the remote peer and establish a unique session identifier during the PPPoE discovery and session stages.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Juniper Networks devices with asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) interfaces can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections.

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the *discovery stage*, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the *session stage*, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

This topic contains the following sections:

- PPPoE Discovery Stage on page 2632
- PPPoE Session Stage on page 2633

PPPoE Discovery Stage

To initiate a PPPoE session, a host must first identify the Ethernet MAC address of the remote peer and establish a unique PPPoE session ID for the session. Learning the remote Ethernet MAC address is called *PPPoE discovery*.

During the PPPoE discovery process, the host does not discover a remote endpoint on the Ethernet network. Instead, the host discovers the access concentrator through which all PPPoE sessions are established. Discovery is a client/server relationship, with the host (a device running Junos OS) acting as the client and the access concentrator acting as the server. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.

NOTE: A device cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE Active Discovery Initiation (PADI)—The client initiates a session by broadcasting a PADI packet to the LAN to request a service.
2. PPPoE Active Discovery Offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.

Copyright © 2016, Juniper Networks, Inc.

3. PPPoE Active Discovery Request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.
4. PPPoE Active Discovery Session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session:
   - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.
   - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends a PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A device supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions per device.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. Magic numbers, echo requests, and all other PPP traffic behave exactly as in normal PPP sessions. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE Active Discovery Termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.

NOTE: If PPPoE session is already up and the user restarts the PPPoE daemon, a new PPPoE daemon with a new PID starts while the existing session is not terminated.

If PPPoE session is already down and user restarts the PPPoE daemon, the PPPoE discovery establishes a new session.

The PPPoE session is not terminated for the following configuration changes:

- Changing idle time out value
- Changing auto rec timer value
- Deleting idle time out
- Deleting auto rec timer
- Add new auto rec time

Copyright © 2016, Juniper Networks, Inc.

- Add new idle time out
- Change negotiate address to static address
- Change static ip address to a new static ip address
- Changing default chap secrete

The PPPoE session is terminated for the following configuration changes:

- Add ac name
- Delete chap ppp options
- Add new chap ppp options
- Configure uifd mac

NOTE: **When the MTU for an underlying physical interface is changed, it brings down the PPPoE session. For PPPoE, an MTU greater than 1492 cannot be achieved.**

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623
- Understanding PPPoE Interfaces on page 2634
- Understanding PPPoE Ethernet Interfaces on page 2641
- *Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces*
- Understanding CHAP Authentication on a PPPoE Interface on page 2642
- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650

Understanding PPPoE Interfaces

The device's Point-to-Point Protocol over Ethernet (PPPoE) interface to the access concentrator can be a Fast Ethernet interface, a Gigabit Ethernet interface, a redundant Ethernet interface, an ATM-over-ADSL interface, or an ATM-over-SHDSL interface. The PPPoE configuration is the same for all interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Ethernet, use a PPPoE encapsulation.
- If the interface is ATM-over-ADSL or ATM-over-SHDSL, use a PPPoE over ATM encapsulation.

To configure a PPPoE interface, you create an interface with a logical interface unit 0, then specify a logical Ethernet or ATM interface as the underlying interface for the PPPoE session. You then specify other PPPoE options, including the access concentrator and PPPoE session parameters.

Copyright © 2016, Juniper Networks, Inc.

NOTE: PPPoE over redundant Ethernet (reth) interface is supported on SRX300, SRX320, and SRX340 devices. This feature allows an existing PPPoE session to continue without starting a new PPP0E session in the event of a failover.

Related Documentation

- Understanding Point-to-Point Protocol on page 2625
- Example: Configuring PPPoE Interfaces on page 2635

Example: Configuring PPPoE Interfaces

This example shows how to configure a PPPoE interface.

- Requirements on page 2635
- Overview on page 2635
- Configuration on page 2635
- Disabling the End-of-List Tag on page 2639

Requirements

Before you begin, configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 2546.

Overview

In this example, you create the PPPoE interface pp0.0 and specify the logical Ethernet interface ge-0/0/1.0 as the underlying interface. You also set the access concentrator, set the PPPoE session parameters, and set the MTU of the IPv4 family to **1492**.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pp0 unit 0 pppoe-options underlying-interface ge-0/0/1.0
  access-concentrator ispl.com auto-reconnect 100 idle-timeout 100 client service-name
  video@ispl.com
set interfaces pp0 unit 0 family inet mtu 1492 negotiate-address
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a PPPoE interface:

1. Create a PPPoE interface.

   [edit]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# edit interfaces pp0 unit 0
```

2. Configure PPPoE options.

```
[edit interfaces pp0 unit 0]
user@host# set pppoe-options underlying-interface ge-0/0/1.0 access-concentrator
  ispl.com auto-reconnect 100 idle-timeout 100 client service-name video@ispl.com
```

3. Configure the MTU.

```
[edit interfaces pp0 unit 0]
user@host# set family inet mtu 1492
```

4. Configure the PPPoE interface address.

```
[edit interfaces pp0 unit 0]
user@host# set family inet negotiate-address
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces pp0
unit 0 {
  pppoe-options {
    underlying-interface ge-0/0/1.0;
    idle-timeout 100;
    access-concentrator ispl.com;
    service-name "vide0@ispl.com";
    auto-reconnect 100;
    client;
  }
  family inet {
    mtu 1492;
    negotiate-address;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying PPPoE Interfaces on page 2636
- Verifying PPPoE Sessions on page 2637
- Verifying the PPPoE Version on page 2638
- Verifying PPPoE Statistics on page 2638

Verifying PPPoE Interfaces

Purpose Verify that the PPPoE device interfaces are configured properly.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show interfaces pp0** command.

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 67, SNMP ifIndex: 317
  Type: PPPoE, Link-level type: PPPoE, MTU: 9192
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface pp0.0 (Index 1) (SNMP ifIndex 330)
    Flags: Point-To-Point SNMP-Traps 16384 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 3304,
      Session AC name: isp1.com, AC MAC address: 00:90:1a:40:f6:4c,
      Service name: video@isp1.com, Configured AC name: isp1.com,
      Auto-reconnect timeout: 60 seconds
      Underlying interface: ge-5/0/0.0 (Index 71)
  Input packets : 23
  Output packets: 22
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 16 (00:00:26 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Success
    Protocol inet, MTU: 1492
      Flags: Negotiate-Address
      Addresses, Flags: Kernel Is-Preferred Is-Primary
        Destination: 211.211.211.2, Local: 211.211.211.1
```

The output shows information about the physical and the logical interfaces. Verify the following information:

- The physical interface is enabled and the link is up.
- The PPPoE session is running on the correct logical interface.
- For state, the state is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
  - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-5/0/0.0.
  - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

Verifying PPPoE Sessions

Purpose Verify that a PPPoE session is running properly on the logical interface.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show pppoe interfaces** command.

```
user@host> show pppoe interfaces
pp0.0 Index 67
  State: Session up, Session ID: 31,
  Service name: video@isp1.com, Configured AC name: isp1.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: ge-0/0/1.0 Index 69
```

The output shows information about the PPPoE sessions. Verify the following information:

- The PPPoE session is running on the correct logical interface.
- For state, the session is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
  - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-0/0/1.0.
  - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

NOTE: To clear a PPPoE session on the pp0.0 interface, use the **clear pppoe sessions pp0.0** command. To clear all sessions on the interface, use the **clear pppoe sessions** command.

Verifying the PPPoE Version

Purpose Verify the version information of the PPPoE protocol configured on the device interfaces.

Action From operational mode, enter the **show pppoe version** command.

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
   PPPoE protocol               = Enabled
   Maximum Sessions             = 256
   PADI resend timeout          = 2 seconds
   PADR resend timeout          = 16 seconds
   Max resend timeout           = 64 seconds
   Max Configured AC timeout    = 4 seconds
```

The output shows PPPoE protocol information. Verify the following information:

- The correct version of the PPPoE protocol is configured on the interface.
- For PPPoE protocol, the PPPoE protocol is enabled.

Verifying PPPoE Statistics

Purpose Verify the statistics information about PPPoE interfaces.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show pppoe statistics** command.

```
user@host> show pppoe statistics
Active PPPoE sessions: 4
  PacketType                        Sent         Received
    PADI                             502                0
    PADO                               0              219
    PADR                             219                0
    PADS                               0              219
    PADT                               0              161
    Service name error                 0                0
    AC system error                    0               13
    Generic error                      0                0
    Malformed packets                  0               41
    Unknown packets                    0                0
  Timeout
    PADI                              42
    PADO                               0
    PADR                               0
```

The output shows information about active sessions on PPPoE interfaces. Verify the following information:

- Total number of active PPPoE sessions running on the interfac
- For packet type, the number of packets of each type sent and received during the PPPoE session

Disabling the End-of-List Tag

During the PPPoE discovery stage, any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client. When a client receives a PADO packet, and if it encounters the **End-of-List** tag in the PADO packet, tags after the **End-of-List** tag are ignored and the complete information is not processed correctly. As a result, the PPPoE connection is not established correctly.

You can configuring the **ignore-eol-tag** option to disable the **End-of-List** tag in the PADO packet.

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To disable the **End-of-List** tag:

1. Create a PPPoE interface.

   ```
   [edit]
   user@host# set interfaces pp0 unit 0
   ```

2. Configure PPPoE options.

   ```
   [edit interfaces pp0 unit 0]
   user@host# set pppoe-options ignore-eol-tag
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
pppoe-options {
  ignore-eol-tag;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verifying That the End-of-List Tag Is Disabled

Purpose Verify the status of the **End-of-List** tag in the PPPoE configuration.

Action From operational mode, enter the **show interfaces pp0.0** command.

```
user@host> show pppoe interfaces pp0.0
Logical interface pp0.0 (Index 78) (SNMP ifIndex 541)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 3,
      Session AC name: cell, Remote MAC address: 00:26:88:f7:77:83,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: Never, Idle timeout: Never,
      Underlying interface: ge-0/0/3.0 (Index 77)
     Ignore End-Of-List tag: Enable
```

```
user@host> show pppoe interfaces pp0.0 extensive
pp0.0 Index 74
  State: Session up, Session ID: 1,
  Service name: None,
  Session AC name: cell, Configured AC name: None,
  Remote MAC address: 00:26:88:f7:77:83,
  Session uptime: 00:02:03 ago,
  Auto-reconnect timeout: 10 seconds, Idle timeout: Never,
  Underlying interface: ge-0/0/3.0 Index 73
  Ignore End-of-List tag: Enable
  PacketType                      Sent         Received
    PADI                            23                0
    PADO                             0                5
    PADR                            11                0
    PADS                             0                2
    PADT                             2                0
    Service name error               0                0
    AC system error                  0                0
    Generic error                    0                0
    Malformed packets                0                0
    Unknown packets                  0                0
  Timeout
    PADI                             3
    PADO                             0
    PADR                             3
  Receive Error Counters
    PADI                             0
    PADO                             0
```


Copyright © 2016, Juniper Networks, Inc.

```
PADR                                0
PADS                                0
```

The output shows information about active sessions on PPPoE interfaces. Verify that the **Ignore End-of-List tag: Enable** option is set.

Related Documentation

- Understanding PPPoE Interfaces on page 2634

Understanding PPPoE Ethernet Interfaces

During a Point-to-Point Protocol over Ethernet (PPPoE) session, the device encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 121 shows a typical PPPoE session between a device and an access concentrator on the Ethernet loop.

Figure 121: PPPoE Session on the Ethernet Loop

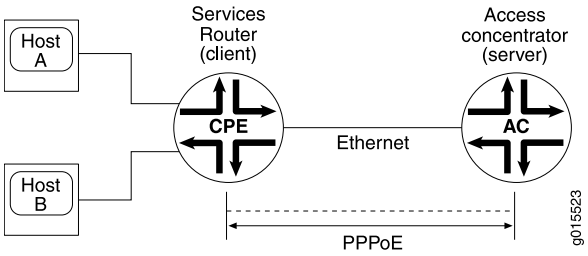

To configure PPPoE on an Ethernet interface, you configure encapsulation on the logical interface.

Related Documentation

- Understanding Point-to-Point Protocol over Ethernet on page 2631
- Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 2641

Example: Configuring PPPoE Encapsulation on an Ethernet Interface

This example shows how to configure PPPoE encapsulation on an Ethernet interface.

- Requirements on page 2641
- Overview on page 2642
- Configuration on page 2642
- Verification on page 2642

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 2546.
- Configure a PPPoE encapsulation interface. See "Example: Configuring PPPoE Interfaces" on page 2635.

Overview

In this example, you configure PPPoE encapsulation on the ge-0/0/1 interface.

Configuration

Step-by-Step Procedure

To configure PPPoE encapsulation:

1. Enable PPPoE encapsulation on the interface.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 encapsulation ppp-over-ether
   ```

2. Commit the configuration if you are done configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces ge-0/0/1** command.

Related Documentation

- Understanding PPPoE Ethernet Interfaces on page 2641

Understanding CHAP Authentication on a PPPoE Interface

For interfaces with Point-to-Point Protocol over Ethernet (PPPoE) encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and be authenticated by its peer.

If you set the **passive** option to handle incoming CHAP packets only, the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not set the **passive** option, the interface always challenges its peer.

You can configure Remote Authentication Dial-In User Service (RADIUS) authentication of PPP sessions using CHAP. CHAP enables you to send RADIUS messages through a routing instance to customer RADIUS servers in a private network.

Related Documentation

- Understanding Point-to-Point Protocol over Ethernet on page 2631
- Example: Configuring CHAP Authentication on a PPPoE Interface on page 2643

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring CHAP Authentication on a PPPoE Interface

This example shows how to configure CHAP authentication on a PPPoE interface.

- Requirements on page 2643
- Overview on page 2643
- Configuration on page 2643
- Verification on page 2644

Requirements

Before you begin:

- Configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 2546.
- Configure a PPPoE interface. See "Example: Configuring PPPoE Interfaces" on page 2635.
- Configure PPPoE encapsulation on an ATM-over-ADSL interface. See *Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface*.

Overview

In this example, you configure a CHAP access profile, and then apply it to the PPPoE interface pp0. You also configure the hostname to be used in CHAP challenge and response packets, and set the passive option for handling incoming CHAP packets.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set access profile A-ppp-client client client1 chap-secret my-secret
set interfaces pp0 unit 0 ppp-options chap access-profile A-ppp-client local-name
  A-ge-0/0/1.0 passive
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure CHAP on a PPPoE interface:

1. Configure a CHAP access profile.

   ```
   [edit]
   user@host# set access profile A-ppp-client client client1 chap-secret my-secret
   ```

2. Enable CHAP options on the interface.

   ```
   [edit]
   user@host# edit interfaces pp0 unit 0 ppp-options chap
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Configure the CHAP access profile on the interface.

   ```
   [edit interfaces pp0 unit 0 ppp-options chap]
   user@host# set access-profile A-ppp-client
   ```

4. Configure a hostname for the CHAP challenge and response packets.

   ```
   [edit interfaces pp0 unit 0 ppp-options chap]
   user@host# set local-name A-ge-0/0/1.0
   ```

5. Set the passive option to handle incoming CHAP packets only.

   ```
   [edit interfaces pp0 unit 0 ppp-options chap]
   user@host# set passive
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces
  pp0 {
  unit 0 {
  ppp-options {
  chap {
  access-profile A-ppp-client;
  local-name A-ge-0/0/1.0;
  passive;
  }
  }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying CHAP Authentication

Purpose Verify that CHAP is enabled on the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Understanding CHAP Authentication on a PPPoE Interface on page 2642

Verifying Credit-Flow Control

Purpose Display PPPoE credit-flow control information about credits on each side of the PPPoE session when credit processing is enabled on the interface.

Copyright © 2016, Juniper Networks, Inc.

Action user@host> **show pppoe interface detail**

```
pp0.51 Index 73
  State: Session up, Session ID: 3,
  Service name: None,
  Configured AC name: None, Session AC name: None,
  Remote MAC address: 00:22:83:84:2e:81,
  Session uptime: 00:05:48 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/4.1 Index 72
  PADG Credits: Local: 12345, Remote: 6789, Scale factor: 128 bytes
  PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
    Quality: 85, Resources 65, Latency 100 msec.
  Dynamic bandwidth: 3 Kbps

pp0.1000 Index 71
  State: Down, Session ID: 1,
  Service name: None,
  Configured AC name: None, Session AC name: None,
  Remote MAC address: 00:00:00:00:00:00,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PADG Credits: enabled
  Dynamic bandwidth: enabled
```

Related Documentation

- Understanding CHAP Authentication on a PPPoE Interface on page 2642
- Verifying Credit-Flow Control on page 2644

Verifying PPPoE Interfaces

Purpose Display PPPoE interfaces information.

Action

- To display PPPoE interface information:

  user@host> **show pppoe interfaces pp0.51 detail**

  ```
  pp0.51 Index 75
    State: Session up, Session ID: 1,
    Service name: None,
    Configured AC name: None, Session AC name: None,
    Remote MAC address: 00:11:22:33:44:55,
    Session uptime: 00:04:18 ago,
    Auto-reconnect timeout: Never, Idle timeout: Never,
    Underlying interface: ge-0/0/1.0 Index 70
    PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
      Quality: 85, Resources 65, Latency 100 msec.
    Dynamic bandwidth: 3 Kbps
  ```

- To display PPPoE terse interface information:

  user@host> **show pppoe interfaces terse pp0.51**

  ```
  Interface    Admin Link Proto    Local         Remote
      pp0.51       up  up   inet   5.1.1.1     --> 5.1.1.2
            inet6    fe80::21f:12ff:fed2:2918/64
                               feee::5:1:1:1/126
  ```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding PPPoE Interfaces on page 2634
- Example: Configuring PPPoE Interfaces on page 2635

Verifying R2CP Interfaces

Purpose Display R2CP interfaces information.

Action

- To display R2CP interface information:

  root@host> **show r2cp interfaces**

```
Interface: ge-0/0/3.51
  Nodes: 0
```

- To display R2CP information:

  root@host> **show r2cp radio extensive**

```
  Node Packet Type          Sent        Received          Errors
  MIM                          -               1               0
  ROM                          1               -               -
  Heartbeats                   0               0               0
  Node Term                    0               0               0
  Node Term Ack                0               0               -

  Heartbeat Timeouts              0
  Node Term Timeouts              0


  Session Packet Type       Sent        Received          Errors
  Init                         -               1               0
  Init ACK                     1               -               -
  Update                       -               0               0
  Terminate                    0               0               0
  Terminate ACK                0               0               0

  Terminate Timeouts              0
```

- To display R2CP session information:

  root@host> **show r2cp sessions extensive**

```
Session: 1
  Destination MAC address 01:02:03:04:05:06
   Status: Established VLANs 201
   Virtual channel:  2
   Session Update:  last received: 3.268 seconds
     Current bandwidth: 22000 Kbps, Maximum 22000 Kbps
     Quality: 100, Resources 100, Latency 100 msec.
   Effective bandwidth: 952 Kbps, last change: 51.484 seconds
     Updates below threshold: 1


   Session Packet Type       Sent        Received          Errors
   Init                         -               1               0
   Init ACK                     1               -               -
   Update                       -               0               0
   Terminate                    0               0               0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Terminate ACK                 0               0               0

    Terminate Timeouts              0
```

Related Documentation

- Understanding PPPoE Interfaces on page 2634
- Example: Configuring PPPoE Interfaces on page 2635

Displaying Statistics for PPPoE

Purpose Display PPPoE statistics.

Action

```
user@host> show interfaces pp0.51 statistics
Logical interface pp0.51 (Index 75) (SNMP ifIndex 137)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: None, Remote MAC address: 00:22:83:84:2f:03,
      Underlying interface: ge-0/0/4.1 (Index 74)
    Input packets : 20865
    Output packets: 284636
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 943 (00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 5.1.1.2, Local: 5.1.1.1
    Protocol inet6, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fed2:2918
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: feee::5:1:1:0/126, Local: feee::5:1:1:1
```

Related Documentation

- Understanding CHAP Authentication on a PPPoE Interface on page 2642
- Verifying Credit-Flow Control on page 2644

Setting Tracing Options for PPPoE

To trace the operations of the router's PPPoE process, include the traceoptions statement at the [edit protocols pppoe] hierarchy level:

```
[edit protocols pppoe]
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
```

Copyright © 2016, Juniper Networks, Inc.

```
    flag flag;
    level severity-level;
    no-remote-trace;
}
```

To specify more than one tracing operation, include multiple **flag** statements.

You can specify the following flags in the **traceoptions** statement:

- **all**—All areas of code
- **config**—Configuration code
- **events**—Event code
- **gres**—Gres code
- **init**—Initialization code
- **interface-db**—Interface database code
- **memory**—Memory management code
- **protocol**—PPPoE protocol processing code
- **rtsock**—Routing socket code
- **session-db**—Session management code
- **signal**—Signal handling code
- **state**—State handling code
- **timer**—Timer code
- **ui**—User interface code

Related Documentation

- Understanding PPPoE Interfaces on page 2634
- Example: Configuring PPPoE Interfaces on page 2635

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 122

Configuring PPPoE-Based Radio-to-Router Protocol

- PPPoE-Based Radio-to-Router Protocols Overview on page 2649
- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650
- Configuring PPPoE-Based Radio-to-Router Protocols on page 2652
- Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 2652
- Credit Flow Control for PPPoE on page 2655
- PPPoE Credit-Based Flow Control Configuration on page 2655

PPPoE-Based Radio-to-Router Protocols Overview

Support for PPPoE-based radio-to-router protocols includes the following extensions to the PPPoE protocol:

- Messages that define how an external device provides the router with timely information about the quality of a link connection
- A flow control mechanism that indicates how much data the router can forward

The router uses the information provided in these PPPoE messages to dynamically adjust the interface speed. When OSPF is notified of this change, it adjusts the cost of the link and updates the routing tables accordingly.

The radio provides ground-to-ground or ground-to-air communications with like devices. When the radio picks up a signal from another device, it initiates a PPPoE session with a directly connected router. The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote routers. The remote radio then forwards traffic over an independent PPPoE session between the remote radio and the router to which it is connected. The two routers exchange LCP and IPCP messages to configure the link and exchange OSPF messages to establish the network topology.

The router and radio are deployed in highly dynamic environments, such as moving vehicles. The quality of the radio link between the routers can vary significantly as a vehicle moves behind an obstruction. Each radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective routers through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the

Copyright © 2016, Juniper Networks, Inc.

PPPoE Discovery Protocol [RFC2516]. The router transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the router detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. An event message notifies OSPF of the change, which then triggers OSPF to announce any resulting routing topology changes to its neighbors.

The PPPoE-based radio-to-router protocol notifies the router about neighbors joining or leaving the network and to create and maintain OSPF adjacencies over the dynamic links established between them. The costs assigned to these links are based on network conditions and flow control information sent by the radios. The calculations and requests to update interface speeds are performed by routines in a common library.

When PPPoE is used for applications, such as mobile radio, the radio links have variable bandwidth. So a mobile radio can function in a PPPoE environment, PPPoE messaging includes PADQ messages, which enable a link cost to be propagated to OSPF through the evaluation of various link quality metrics. The router uses information from these notifications along with user-configured parameters to calculate interface link costs that are used by the routing protocols.

A radio can send an optional PADQ at any time to query or report link quality metrics. When transmitting PPP streams over radio links, the quality of the link directly affects the throughput. The PADQ packet is used by the radio modem to report link metrics.

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can also grant each other forwarding credits. The grantee can forward traffic to the peer only when it has a sufficient number of credits to do so. Credit-based forwarding allows both sides of the session to agree to use a non-default credit scaling factor during the PADR and PADS message exchange. Although this is used on both sides of the session, this feature provides the radio client with a flow control mechanism that throttles traffic by limiting the number of credits it grants to the router.

Related Documentation

- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650

Understanding the PPPoE-Based Radio-to-Router Protocol

Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

For example, a high-band networking waveform (HNW) radio provides ground-to-ground or ground-to-air communications with like devices. When the HNW picks up a signal from another device, it initiates a PPPoE session with a directly connected device (router). The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote devices. The remote radio then forwards traffic to a remote device using an independent PPPoE session. The two devices exchange Link Control Protocol (LCP) and Internet Protocol Control Protocol (IPCP) messages to configure the link and exchange OSPF messages to establish the network topology.

Copyright © 2016, Juniper Networks, Inc.

Each HNW radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective devices through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the PPPoE Discovery Protocol (RFC 2516). The device transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the device detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. OSPF is notified of the change and announces any resulting routing topology changes to its neighbors.

The CLI statement, **radio-router**, indicates that metrics announcements received on the interface will be processed by the device. When a PPPoE logical interface refers to this as an underlying interface, the device then processes incoming PADQ messages and uses information from the host's messages to control the flow of traffic and manage the speed of the link, resulting in a corresponding adjustment of the OSPF cost. If this option is not specified, then PADQ messages received over the underlying interface are ignored.

The following options are available within the **radio-router** configuration statement:

- **bandwidth**, **resource**, **latency**, and **quality** —These statements provide control over the weights used when transforming PADQ link metrics into an interface speed for the virtual link:
  - **bandwidth**—Weight of current (vs. maximum) data rate
  - **resource**—Resource weight
  - **latency**—Latency weight
  - **quality**—Relative link quality weight

  All four weights accept values from 0 through 100. The default value for all four weights is 100.
- **credit**—This statement supports the credit-based flow control extensions described in RFC 4938, The statement enables PPPoE peers to grant each other forwarding credits. The grantee is then allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. The subsequent credit interval statement controls how frequently the device generates credit announcement messages. The **interval** sub-statement, which controls the grant rate interval, accepts values from 1 through 60 seconds.
- **threshold**—This statement specifies how much of a difference is required between the calculated and the current interface speeds. The **threshold** value, expressed as a percentage, defaults to 10.

The following hierarchy provides another view of the **radio-router** configuration statements.

```
interfaces{
  interface-name {
    radio-router {
      bandwidth;
      credit {
        interval;
      }
```

```
            latency;
            quality;
            resource;
            threshold;
        }
    }
}
```

Related Documentation

- Understanding Point-to-Point Protocol over Ethernet on page 2631
- *Example: Configuring the PPPoE-Based Radio-to-Router Protocol*

Configuring PPPoE-Based Radio-to-Router Protocols

To configure the PPPoE-based radio-to-router protocol:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Configure radio-router on the logical Ethernet interface.
3. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
4. Configure the operational mode as server.
5. (Optional) Identify the access concentrator by a unique name.
6. Specify how many seconds to wait before attempting to reconnect.
7. Provide a name for the type of service provided by the access concentrator.
8. Configure the maximum transmission unit (MTU) of the interface.
9. Configure the MTU size for the protocol family.
10. Disable the sending of keepalive messages on the logical interface.

Related Documentation

- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650
- Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 2652

Example: Configuring the PPPoE-Based Radio-to-Router Protocol

This example shows how to configure the PPPoE-based radio-to-router protocol.

- Requirements on page 2652
- Overview on page 2653
- Configuration on page 2653
- Verification on page 2654

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

1. Configure network interfaces. See "Example: Creating an Ethernet Interface" on page 2546.
2. Configure PPPoE interfaces. See "Example: Configuring PPPoE Interfaces" on page 2635.
3. Configure PPPoE encapsulation on an Ethernet interface. See "Example: Configuring PPPoE Encapsulation on an Ethernet Interface" on page 2641.
4. Configure PPPoE encapsulation on an ATM-over-ADSL interface. See *Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface*.
5. Configure CHAP authentication on a PPPoE interface. See "Example: Configuring CHAP Authentication on a PPPoE Interface" on page 2643.

Overview

In this example, you configure the ge-3/0/3 interface and set the bandwidth, resource, latency, and quality to **100**. You also set the threshold value to **10**, and then configure options on the logical interface.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces ge-3/0/3 unit 1 radio-router bandwidth 100 resource 100 latency 100 quality
  100 threshold 10
set interfaces pp0 unit 1 pppoe-options underlying-interface ge-3/0/3 server
set interfaces pp0 unit 1 family inet unnumbered-address lo0.0 destination 192.168.1.2
set interfaces pp0 unit 1 family inet6 address lo0.0 destination fec0:1:1:1::2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the PPPoE-based radio-to-router protocol:

1. Enable the PPPoE-based radio-to-router protocol.

   ```
   [edit]
   user@host# edit interfaces ge-3/0/3 unit 1 radio-router
   ```

2. Set the interface speed for the virtual link.

   ```
   [edit interfaces ge-3/0/3 unit 1 radio-router]
   user@host# set bandwidth 100 resource 100 latency 100 quality 100
   ```

3. Set the calculated and current interface speeds, as a percentage.

   ```
   [edit interfaces ge-3/0/3 unit 1 radio-router]
   user@host# set threshold 10
   ```

4. Configure options on the logical interface.

   ```
   [edit interfaces pp0 unit 1]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set pppoe-options underlying-interface ge-3/0/3
user@host# set pppoe-options server
user@host# set family inet unnumbered-address lo0.0 destination 192.168.1.2
user@host# set family inet6 address lo0.0 destination fec0:1:1:1::2
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show interfaces** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  user@host# show interfaces ge-3/0/3 {
  unit 1
  radio-router {
  bandwidth 100;
  resource 100;
  latency 100;
  quality 100;
  threshold 10;
  }
  }
}
  ...
  pp0 {
    unit 1 {
    pppoe-options {
    underlying-interface ge-3/0/3;
    server;
  }
  family inet {
    unnumbered-address lo0.0 destination 192.168.1.2;
    }
    family inet6;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the PPPoE-based Radio-to-Router Protocol

Purpose Verify the PPPoE-Based radio-to-router protocol.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650

Copyright © 2016, Juniper Networks, Inc.

Credit Flow Control for PPPoE

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can grant each other forwarding credits. The grantee is allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. When credit-based forwarding is used on both sides of the session, the radio client can throttle traffic by limiting the number of credits it grants to the router.

The **interfaces** statement includes the **radio-router** attribute, which contains the parameters used for rate-based scheduling and OSPF link cost calculations. It also includes the **credit** attribute to indicate that credit-based packet scheduling is supported on the PPPoE interfaces that reference this underlying interface. Interfaces that set the **encapsulation** attribute support the PPPoE Active Discovery Grant (PADG) and PPPoE Active Discovery Credit (PADC) messages in the same way that the **radio-router** attribute provides active support for the PPPoE Active Discovery Quality (PADQ) message.

The **credit interval** parameter controls how frequently the router generates credit announcement messages. For PPPoE this corresponds to the interval between PADG credit announcements for each session.

Related Documentation

- PPPoE-Based Radio-to-Router Protocols Overview on page 2649
- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650
- Configuring PPPoE-Based Radio-to-Router Protocols on page 2652

PPPoE Credit-Based Flow Control Configuration

This example shows a PPPoE credit-based flow control configuration.

```
[edit interfaces ge-0/0/1]
unit 0 {
   encapsulation ppp-over-ether;
   radio-router {
      credit {
         interval 10;
      }
      bandwidth 80;
      threshold 5;
   }
}
```

Related Documentation

- Understanding the PPPoE-Based Radio-to-Router Protocol on page 2650
- Configuring PPPoE-Based Radio-to-Router Protocols on page 2652

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 123

Configuring R2CP Radio-to-Router Protocol

- R2CP Radio-to-Router Protocol Overview on page 2657
- Configuring the R2CP Radio-to-Router Protocol on page 2658

R2CP Radio-to-Router Protocol Overview

The Network Centric Waveform (NCW) radio-specific radio-to-router control protocol (R2CP) is similar to the PPPoE radio-to-router protocol. Both of these protocols exchange dynamic metric changes in the network that the routers use to update the OSPF topologies.

In radio-router topologies, the router connects to the radio over a Gigabit Ethernet link and the radio transmits packets over the radio frequency (RF) link. The radio periodically sends metrics to the router, which uses RF link characteristics and other data to inform the router on the shaping and OSPF link capacity. The router uses this information to shape the data traffic and provide the OSPF link cost for its SPF calculations. The radio functions like a Layer 2 switch and can only identify remote radio-router pairs using the Layer 2 MAC addresses. With R2CP the router receives metrics for each neighboring router, identified by the MAC address of the remote router. The R2CP daemon translates the MAC addresses to link the local IPv6 address and sends the metrics for each neighbor to OSPF. Processing these metrics is similar to the handling of PPPoE PADQ metrics. Unlike PPPoE, which is a point-to-point link, these R2CP neighbors are treated as nodes in a broadcast LAN.

You must configure each neighbor node with a per unit scheduler for CoS. The scheduler context defines the attributes of Junos class-of-service. To define CoS for each radio, you can configure virtual channels to limit traffic. You need to configure virtual channels for as many remote radio-router pairs as there are in the network. You configure virtual channels on a logical interface. Each virtual channel can be configured to have a set of eight queues with a scheduler and an optional shaper. When the radio initiates the session with a peer radio-router pair, a new session is created with the remote MAC address of the router and the VLAN over which the traffic flows. Junos OS chooses from the list of free virtual channels and assigns the remote MAC and the eight CoS queues and the scheduler to this remote MAC address. All traffic destined to this remote MAC address is subjected to the CoS that is defined in the virtual channel.

Copyright © 2016, Juniper Networks, Inc.

A virtual channel group is a collection of virtual channels. Each radio can have only one virtual channel group assigned uniquely. If you have more than one radio connected to the router, you must have one virtual channel group for each local radio-to-router pair. Although a virtual channel group is assigned to a logical interface, a virtual channel is not the same as a logical interface. The only features supported on a virtual channel are queuing, packet scheduling, and accounting. Rewrite rules and routing protocols apply to the entire logical interface.

All nodes in the R2CP network are in a broadcast LAN. The point-to-multipoint over LAN protocol supports advertising different bandwidth information for neighbors on a broadcast link. The network link is a point-to-multipoint link in the OSPFv3 link state database, which uses existing OSPF neighbor discovery to provide automatic discovery without configuration. It enables each node to advertise a different metric to every other node in the network to accurately represent the cost of communication. The **p2mp-over-lan** interface type under the OSPFv3 interface configuration enables you to configure the interface. OSPFv3 then uses LAN procedures for neighbor discovery and flooding, but represents the interface as point-to-multipoint in the link state database.

The interface type and router LSA are available under the following hierarchies:

[protocols ospf3 area *area-id* interface *interface-name*]

[routing-instances *routing-instances-name* protocols ospf3 area *area-id* interface *interface-name*]

For example:

```
protocols {
   ospf3 {
      area 0.0.0.0 {
         interface ge-0/0/2.0 {
            interface-type p2mp-over-lan;
         }
      }
   }
}
```

Related Documentation

- Configuring the R2CP Radio-to-Router Protocol on page 2658

Configuring the R2CP Radio-to-Router Protocol

To configure the R2CP protocol:

1. Configure the interfaces.

   The following example creates four logical interfaces on ge-0/0/2, using unit 52 for R2CP control messages and units 101-193 for data traffic. The **per-unit-scheduler** statement is required for R2CP.

   ```
   interfaces {
      ge-0/0/2 {
         per-unit-scheduler;
   ```

 Copyright © 2016, Juniper Networks, Inc.

```
    vlan-tagging;
    unit 52 {
      vlan-id 52;
      family inet {
        address 52.1.1.1/24;
      }
    }
    unit 101 {
      vlan-id 101;
      family inet {
        address 101.1.1.1/24;
      }
    }
    unit 102 {
      vlan-id 102;
      family inet {
        address 102.1.1.1/24;
      }
    }
    unit 103 {
      vlan-id 103;
      family inet {
        address 103.1.1.1/24;
      }
    }
  }
}
```

2. Configure the R2CP protocol.

   The following example configures g2-0/0/2.52 as the interface for R2CP control messages, vg1 as the virtual-channel group, and ge-0/0/2.101-103 as data interfaces using the radio-interface statement.

```
protocols {
  r2cp {
    radio myRadio {
      interface ge-0/0/2.52;
      virtual-channel-group vg1;
      radio-interface ge-0/0/2.101;
      radio-interface ge-0/0/2.102;
      radio-interface ge-0/0/2.103;
    }
  }
}
```

3. Configure class of service.

   The following example defines virtual-channels, their initial shaping-rates, and the virtual-channel-group to which they belong. It also makes the association between radio-interface interfaces and virtual-channel-group. In the class of service configuration, the **vc-shared-scheduler** configuration statement is required for each interface configured as a radio interface in the R2CP protocol configuration.

```
class-of-service {
  virtual-channels {
    vc1;
```

```
    vc2;
    vc3;
    vc4;
  }
  virtual-channel-groups {
    vg1 {
      vc1 {
        scheduler-map sm;
        shaping-rate 15m;
        default;
      }
      vc2 {
        scheduler-map sm;
        shaping-rate 20m;
      }
      vc3 {
        scheduler-map sm;
        shaping-rate 20m;
      }
      vc4 {
        scheduler-map sm;
        shaping-rate 20m;
      }
    }
  }
  forwarding-classes {
    queue 0 DATA-queue;
  }
  interfaces {
    ge-0/0/2 {
      unit 101 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
      unit 102 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
      unit 103 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
    }
  }
  scheduler-maps {
    sm {
      forwarding-class DATA-queue scheduler sm-scheduler;
    }
  }
  schedulers {
    sm-scheduler {
      transmit-rate percent 20;
      buffer-size percent 20;
      priority low;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

}

Related Documentation

- R2CP Radio-to-Router Protocol Overview on page 2657

Copyright © 2016, Juniper Networks, Inc.

PART 39

Configuration Statements and Operational Commands

- Configuration Statements on page 2665
- Operational Commands on page 2759

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 124

Configuration Statements

- Chassis Configuration Statement Hierarchy on page 2668
- Class-of-Service Configuration Statement Hierarchy on page 2671
- Interfaces Configuration Statement Hierarchy on page 2675
- PoE Configuration Statement Hierarchy on page 2690
- accept-source-mac on page 2691
- access-point-name on page 2692
- apply-groups on page 2692
- arp-resp on page 2693
- authentication-method (Interfaces) on page 2693
- bandwidth (Interfaces) on page 2694
- bundle (Interfaces) on page 2694
- cbr rate on page 2695
- cellular-options on page 2695
- classifiers (CoS) on page 2696
- client-identifier (Interfaces) on page 2697
- code-points (CoS) on page 2697
- compression-device (Interfaces) on page 2698
- credit (Interfaces) on page 2698
- data-rate on page 2699
- disable (PoE) on page 2699
- dhcp (Interfaces) on page 2700
- duration (PoE) on page 2700
- encapsulation (Interfaces) on page 2701
- family inet (Interfaces) on page 2702
- family inet6 on page 2705
- flag (Interfaces) on page 2707
- flexible-vlan-tagging (Interfaces) on page 2708
- flow-control (Interfaces) on page 2708

Copyright © 2016, Juniper Networks, Inc.

- flow-monitoring (Services) on page 2709
- forwarding-classes (CoS) on page 2710
- fpc (Interfaces) on page 2711
- gratuitous-arp-reply on page 2711
- gsm-options on page 2712
- guard-band (PoE) on page 2712
- hub-assist on page 2713
- inline-jflow (Forwarding Options) on page 2713
- interface (PIC Bundle) on page 2714
- interface (PoE) on page 2715
- interfaces (CoS) on page 2716
- interval (Interfaces) on page 2717
- interval (PoE) on page 2717
- ipv4-template (Services) on page 2718
- ipv6-template (Services) on page 2718
- lacp (Interfaces) on page 2719
- latency (Interfaces) on page 2719
- lease-time on page 2720
- line-rate (Interfaces) on page 2720
- link-speed (Interfaces) on page 2721
- loopback (Interfaces) on page 2721
- loss-priority (CoS Loss Priority) on page 2722
- loss-priority (CoS Rewrite Rules) on page 2723
- loss-priority-maps (CoS Interfaces) on page 2724
- loss-priority-maps (CoS) on page 2724
- management (PoE) on page 2725
- maximum-power (PoE) on page 2725
- media-type (Interfaces) on page 2726
- minimum-links (Interfaces) on page 2726
- native-vlan-id (Interfaces) on page 2727
- next-hop-tunnel on page 2727
- no-dns-propagation on page 2728
- option-refresh-rate (Services) on page 2728
- pic-mode (Chassis T1 Mode) on page 2729
- periodic (Interfaces) on page 2730
- ppp-over-ether on page 2730
- pppoe on page 2731

Copyright © 2016, Juniper Networks, Inc.

- pppoe-options on page 2732
- priority (PoE) on page 2733
- profile (Access) on page 2734
- profiles on page 2737
- promiscuous-mode (Interfaces) on page 2738
- quality (Interfaces) on page 2738
- r2cp on page 2739
- radio-router (Interfaces) on page 2740
- redundancy-group (Interfaces) on page 2741
- redundant-ether-options on page 2741
- redundant-parent (Interfaces Fast Ethernet) on page 2742
- redundant-parent (Interfaces Gigabit Ethernet) on page 2742
- resource (Interfaces) on page 2743
- retransmission-attempt on page 2743
- retransmission-interval (Interfaces) on page 2744
- roaming-mode on page 2744
- scheduler-map (CoS Virtual Channels) on page 2745
- select-profile on page 2745
- server-address on page 2746
- shaping-rate (CoS Interfaces) on page 2747
- simple-filter (Interfaces) on page 2748
- sip-password on page 2748
- sip-user-id on page 2749
- source-address-filter (Interfaces) on page 2750
- source-filtering (Interfaces) on page 2751
- speed (Interfaces) on page 2751
- telemetries (PoE) on page 2752
- template-refresh-rate (Services) on page 2752
- threshold (Interfaces) on page 2753
- traceoptions (Interfaces) on page 2753
- update-server on page 2754
- vbr rate on page 2754
- vdsl-profile on page 2755
- vendor-id (Interfaces) on page 2755
- vlan-tagging (Interfaces) on page 2756
- web-authentication (Interfaces) on page 2757

Copyright © 2016, Juniper Networks, Inc.

Chassis Configuration Statement Hierarchy

Use the statements in the **chassis** configuration hierarchy to configure alarms, aggregated devices, clusters, the Routing Engine, and other chassis properties.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority number;
      }
    }
    sonet {
      device-count number;
    }
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
    ethernet {
      link-down (ignore | red | yellow);
    }
    integrated-services {
      failure (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
    serial {
      cts-absent (ignore | red | yellow);
      dcd-absent (ignore | red | yellow);
      dsr-absent (ignore | red | yellow);
      loss-of-rx-clock (ignore | red | yellow);
      loss-of-tx-clock (ignore | red | yellow);
    }
    services {
      hw-down (ignore | red | yellow);
      linkdown (ignore | red | yellow);
      pic-hold-reset (ignore | red | yellow);
      pic-reset (ignore | red | yellow);
      rx-errors (ignore | red | yellow);
      sw-down (ignore | red | yellow);
      tx-errors (ignore | red | yellow);
    }
    t3 {
      ais (ignore | red | yellow);
      exz (ignore | red | yellow);
      ferf (ignore | red | yellow);
```


Copyright © 2016, Juniper Networks, Inc.

```
      idle (ignore | red | yellow);
      lcv (ignore | red | yellow);
      lof (ignore | red | yellow);
      los (ignore | red | yellow);
      pll (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
  }
  cluster {
    configuration-synchronize {
      no-secondary-bootup-auto;
    }
    control-link-recovery;
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    network-management {
      cluster-master;
    }
    redundancy-group  group-number {
      gratuitous-arp-count number;
      hold-down-interval number;
      interface-monitor interface-name {
        weight number;
      }
      ip-monitoring {
        family {
          inet {
            ipv4-address {
              interface {
                logical-interface-name;
                secondary-ip-address ip-address;
              }
              weight number;
            }
          }
        }
        global-threshold number;
        global-weight number;
        retry-count number;
        retry-interval seconds;
      }
      node (0 | 1 ) {
        priority number;
      }
      preempt;
    }
    reth-count number;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (world-readable | no-world-readable);
        size maximum-file-size;
      }
      flag flag;
```

Copyright © 2016, Juniper Networks, Inc.

```
      level {
        (alert | all | critical | debug | emergency | error | info | notice | warning);
      }
      no-remote-trace;
    }
  }
  config-button {
    no-clear;
    no-rescue;
  }
  craft-lockout;
  fpc slot-number {
    offline;
    pic slot-number {
      aggregate-ports;
      framing {
        (e1 | e3 | sdh | sonet | t1 | t3);
      }
      ingress-policer-overhead bytes
      max-queues-per-interface (4 | 8);
      mlfr-uni-nni-bundles number;
      no-multi-rate;
      np-cache;
      port slot-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
        speed (oc12-stm4 | oc3-stm1 | oc48-stm16);
      }
      q-pic-large-buffer (large-scale | small-scale);
      services-offload {
        low-latency;
        per-session-statistics;
      }
      shdsl {
        pic-mode (1-port-atm | 2-port-atm | 4-port-atm | efm);
      }
      sparse-dlcis;
      traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode (egress-only | ingress-and-egress);
      }
      tunnel-queuing;
    }
    services-offload;
  }
  ioc-npc-connectivity {
    ioc slot-number {
      npc (npc-slot-number | none);
    }
  }
  maximum-ecmp (16 | 32 | 64);
  network-services (ethernet | IP);
  routing-engine {
    bios {
      no-auto-upgrade;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
    usb-wwan {
      port 1;
    }
  }
  usb {
    storage {
      disable;
    }
  }
}
```

Related Documentation

- cluster (Chassis) on page 3563
- *ip-monitoring*

Class-of-Service Configuration Statement Hierarchy

Use the statements in the **class-of-service** configuration hierarchy to configure class-of-services (CoS) features.

```
class-of-service {
  adaptive-shapers adaptive-shaper-name {
    trigger becn {
      shaping-rate (absolute-rate | percent percent);
    }
  }
  application-traffic-control {
    rate-limiters rate-limiter-name {
      bandwidth-limit kbps;
      burst-size-limit bytes;
    }
    rule-sets rule-set-name {
      rule rule-name {
          match {
            application [application-name];
            application-any;
            application-group [application-group-name];
            application-known;
            application-unknown;
          }
          then {
            dscp-code-point dscp-value;
            forwarding-class class-name;
            log;
            loss-priority (high | low |medium-high | medium-low);
            rate-limit {
              loss-priority-high;
              client-to-server rate-limiter;
              server-to-client rate-limiter;
            }
          }
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier name
    {
    forwarding-class class-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-points [alias-or-bit-string ];
      }
    }
    import (classifier-name | default);
  }
}
code-point-aliases {
  (dscp | dscp-ipv6 |exp |ieee-802.1 |ieee-802.1ad |inet-precedence) alias-name{
    dscp-bits;
    }
}
drop-profiles profile-name {
  fill-level percent {
    drop-probability number;
  }
  interpolate {
    drop-probability [number];
    fill-level [percent];
  }
}
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    next-hop-map next-hop-map-name {
      forwarding-class class-name {
        discard;
        lsp-next-hop [lsp-regular-expression];
        next-hop [next-hop-identifier];
        non-lsp-next-hop;
      }
    }
  }
  fragmentation-maps fragmentation-map-name {
    forwarding-class forwarding-class-name {
      drop-timeout milliseconds;
        (fragment-threshold bytes |no-fragmentation) ;
        multilink-class number;
      }
    }
  host-outbound-traffic {
    dscp-code-point static-dscp-code-point;
    forwarding-class class-name;
    tcp {
      raise-internet-control-priority;
    }
  }
  interfaces interface-name {
    input-traffic-control-profile profile-name;
      output-traffic-control-profile profile-name;
      output-traffic-control-profile-remaining profile-name;
      scheduler-map scheduler-map;
      shaping-rate bps;
      unit  logical-unit-number  {
        adaptive-shaper adaptive-shaper-name;
        classifiers {
          (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
        }
        forwarding-class class-name;
        input-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        loss-priority-maps {
          frame-relay-de {
            (lpmap-name | default);
          }
        }
        output-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        rewrite-rules {
          (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad
            |inet-precedence)
        }
        scheduler-map scheduler-map-name;
          shaping-rate {
          rate;
```

```
        }
        vc-shared-scheduler;
        virtual-channel-group group-name;
      }
    }
  }
  loss-priority-maps {
    frame-relay-de loss-priority-map-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-points [bit-string];
      }
    }
  }
  rewrite-rules {
    (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      rewrite-rule-name {
      forwarding-class forwarding-class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-point alias-or-bit-string;
        }
        import (default | rewrite-rule-name);
      }
    }
  }
  scheduler-maps scheduler-map-name {
    forwarding-class class-name {
      scheduler scheduler-name;
      }
    }
  schedulers scheduler-name {
    buffer-size {
      exact;
      (percent percent | remainder percent  | temporal microseconds) ;
    }
    drop-profile-map {
      loss-priority (any | high | low | medium-high | medium-low);
      protocol any ;
      drop-profile profile;
    }
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (absolute-rate | percent percent);
    transmit-rate <exact> (percent percent | rate bits | remainder percent);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
  traffic-control-profiles profile-name {
    delay-buffer-rate ( absolute-rate | cps cells-per-second | percent percent);
```


Copyright © 2016, Juniper Networks, Inc.

```
        guaranteed-rate (absolute-rate | percent percent);
        overhead-accounting (bytes bytes | cell-mode | frame-mode);
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    tri-color;
    virtual-channel-groups virtual-channel-group-name {
      virtual-channel-name {
        default;
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    }
    virtual-channels virtual-channel-name;
  }
```

Related Documentation

- SSL Proxy Overview on page 531
- Understanding Interfaces on page 2441

Interfaces Configuration Statement Hierarchy

Use the statements in the **interfaces** configuration hierarchy to configure interfaces on the device.

NOTE: For a gigabit ethernet interface, the **gigether-options** and **ether-options** are identical, but only the **gigether-options** are documented.

```
interfaces {
  interface-name {
    accounting-profile name;
    clocking (external | internal);
    dce;
    description text;
    disable;
    e1-options {
      bert-algorithm algorithm;
      bert-error-rate rate;
      bert-period seconds;
      fcs (16 | 32);
      framing (g704 | g704-no-crc4 | unframed);
      idle-cycle-flag (flags | ones);
      invert-data data;
      loopback (local | remote);
      start-end-flag (shared | filler);
      timeslots time-slot-range;
    }
    e3-options {
      bert-algorithm algorithm;
      bert-error-rate rate;
      bert-period seconds;
      compatibility-mode {
```

Copyright © 2016, Juniper Networks, Inc.

```
      digital-link {
        subrate value;
      }
      kentrox {
        subrate value;
      }
      larscom;
    }
    fcs (16 | 32);
    framing (g.751 | g.832);
    idle-cycle-flag value;
    invert-data;
    loopback (local | remote);
    (no-payload-scrambler | payload-scrambler);
    (no-unframed | -unframed);
    start-end-flag (filler | shared);
  }
  encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc |
    ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc |
    extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls
    | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
  fastether-options {
    802.3ad interface-name {
      (backup | primary);
      lacp {
        port-priority port-number;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
  }
  flexible-vlan-tagging;
  gigether-options {
    802.3ad interface-name {
      (backup | primary);
      lacp {
        port-priority port-number;
      }
    }
    (auto-negotiation <remote-fault> (local-interface-offline | local-interface-online)
      | no-auto-negotiation);
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth [number];
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
}
gratuitous-arp-reply;
hierarchical-scheduler {
    maximum-hierarchy-levels 2;
}
hold-time {
    down milliseconds;
    up milliseconds;
}
keepalives {
    down-count number;
    interval number;
    up-count number;
}
link-mode (full-duplex | half-duplex);
lmi {
    lmi-type (ansi | c-lmi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte number;
    t392dce number;
}
logical-tunnel-options {
    per-unit-mac-disable;
}
mac mac-address;
mtu bytes;
native-vlan-idvlan-id;
no-gratuitous-arp-request;
no-keepalives;
optics-options {
    alarm {
        low-light-alarm (link-down | syslog);
    }
    warning {
        low-light-warning (link-down | syslog);
    }
    wavelength wavelength-options;
}
otn-options {
    bytes {
        transmit-payload-type number];
    }
    fec (efec | gfec | none);
    (laser-enable | no-laser-enable);
    (line-loopback | no-line-loopback);
    rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
    trigger {
        oc-lof {
```

Copyright © 2016, Juniper Networks, Inc.

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  oc-lom {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  oc-los {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  oc-wavelength-lock {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-ais {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-lck {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-oci {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-sd {
```


Copyright © 2016, Juniper Networks, Inc.

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-bbe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-es {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-ses {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-uas {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-ttim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  opu-ptim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-ais {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
```

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-exe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-iae {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-sd {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-bbe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-es {
```


Copyright © 2016, Juniper Networks, Inc.

```
            hold-time {
              down milliseconds;
              up milliseconds;
              }
            ignore;
          }
          otu-tca-ses {
            hold-time {
              down milliseconds;
              up milliseconds;
              }
            ignore;
          }
          otu-tca-uas {
            hold-time {
              down milliseconds;
              up milliseconds;
              }
            ignore;
          }
          otu-ttim {
            hold-time {
              down milliseconds;
              up milliseconds;
              }
            ignore;
          }
        }
        tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi |
          otu-dapi |otu-expected-receive-dapi | otu-expected-receive-sapi |otu-sapi);
      }
      passive-monitor-mode;
      (per-unit-scheduler | no-per-unit-schedule);
      port-mirror-instance;
      ppp-options {
        chap {
          access-profile name;;
          default-chap-secret secret;
          local-name name;
          no-rfc2486;
          passive;
        }
        compression {
          acfc;
          pfc;
        }
        dynamic-profile (dynamic-profile | junos-default-profile);
        lcp-max-conf-req number;
        lcp-restart-timer milliseconds;
        loopback-clear-timer seconds;
        ncp-max-conf-req number;
        ncp-restart-timer milliseconds;
        no-termination-request;
        pap {
          access-profile name;
          default-password password;
```

```
        local-name name;
        local-password password;
        no-rfc2486;
        passive;
    }
}
promiscuous-mode;
receive-bucket {
    overflow {
        discard;
        tag;
    }
    rate number;
    threshold number;
}
redundant-pseudo-interface-options {
    redundancy-group number;
}
satop-options {
    excessive-packet-loss-rate {
        sample-period milliseconds;
        threshold percentage;
    }
    idle-pattern number;
    (jitter-buffer-auto-adjust | jitter-buffer-latency milliseconds | jitter-buffer-packets
        number;
    payload-size number;
}
speed (100m |10m | 1g);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        cascade-port;
        link-mode (full-duplex | half-duplex);
        speed (100m |10m | 1g);
        vlan-id number;
    }
}
t1-options {
    alarm-compliance {
        accunet-t1-5-service;
    }
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flags (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
```


Copyright © 2016, Juniper Networks, Inc.

```
    timeslots time-slot-range;
}
t3-options {
    bert-algorithm algorithm ;
    bert-error-rate rate ;
    bert-period seconds ;
    (cbit-parity | no-cbit-parity);
    compatibility-mode {
        adtran {
            subrate value;
        }
        digital-link {
            subrate value;
        }
        kentrox {
            subrate value;
        }
        larscom;
            subrate value;
        }
        verilink;
            subrate value;
        }
    }
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag (flags | ones);
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (no-payload-scrambler | payload-scrambler);
    (no-unframed | unframed);
    start-end-flag value (filler | shared);
}
traceoptions {
    flag (all | event | ipc | media);
}
transmit-bucket {
    overflow {
        discard;
    }
    rate number;
    threshold number;
}
(traps | no-traps);
unit unit-number {
    accept-source-mac {
        mac-address mac-address;
    }
    accounting-profile name;
    arp-resp (restricted | unrestricted);
    backup-options {
        interface interface-name;
    }
    bandwidth bandwidth;
    description text;
```

Copyright © 2016, Juniper Networks, Inc.

```
disable;
encapsulation (dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge |
  vlan-ccc | vlan-vpls |vlan-tcc);
family {
  ccc {
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
    }
    policer {
      input input-policer-name;
      output output-policer-name;
    }
  }
  ethernet-switching {
    bridge-domain-type (svlan| bvlan);
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
    }
    interface-mode (access | trunk);
    native-vlan-id native-vlan-id;
    policer {
      input input-policer-name;
      output outputpolicer-name;
    }
    port-mode (access | tagged-access | trunk);
    reflective-relay;
    vlan-id vlan-id;
    vlan members [vlan-id];
    vlan-rewrite {
      translate {
        from-vlan-id;
        to-vlan-id ;
      }
    }
  }
  inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
    address (source–address/prefix) {
      arp destination-address {
        (mac mac-address | multicast-mac multicast-mac-address);
        publish publish-address;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key-value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds
      (preempt <hold-timeseconds> | no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
          priority-cost value;
        }
      }
      virtual-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  dhcp {
    client-identifier {
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
```

Copyright © 2016, Juniper Networks, Inc.

```
            user-id (ascii string| hexadecimal string);
          }
          lease-time (length | infinite);
          retransmission-attempt value;
          retransmission-interval seconds;
          server-address server-address;
          update-server;
          vendor-id vendor-id ;
        }
        filter {
          group number;
          input filter-name;
          input-list [filter-name];
          output filter-name;
          output-list [filter-name];
        }
        mtu value;
        no-neighbor-learn;
        no-redirects;
        policer {
          arp arp-name;
          input input-name;
          output output-name;
        }
        primary;
        rpf-check {
          fail-filter filter-name;
          mode {
            loose;
          }
        }
        sampling {
          input;
          output;
          simple-filter;
        }
        targeted-broadcast {
          (forward-and-send-to-re |forward-only);
        }
        unnumbered-address {
          interface-name;
          preferred-source-address preferred-source-address;
        }
      }
      inet6 {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            ouput;
          }
        }
        address source–address/prefix {
          eui-64;
          ndp address {
            (mac mac-address | multicast-mac multicast-mac-address);
```


Copyright © 2016, Juniper Networks, Inc.

```
        publish;
      }
      preferred;
      primary;
      vrrp-inet6-group group_id {
        (accept-data | no-accept-data);
        advertisements-threshold number;
        authentication-key value;
        authentication-type (md5 | simple);
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        (preempt <hold-time seconds>| no-preempt );
        priority value;
        track {
          interface interface-name {
            bandwidth-threshold value;
            priority-cost value;
          }
          priority-hold-time seconds;
          route route-address{
            routing-instance routing-instance;
          }
        }
        virtual-inet6-address [address];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group value;
          active-interface interface-name;
        }
      }
      web-authentication {
        http;
        https;
        redirect-to-https;
      }
    }
    (dad-disable | no-dad-disable);
    dhcpv6-client {
      client-ia-type (ia-na | ia-pd);
      client-identifier duid-type (duid-ll | duid-llt | vendor);
      client-type (autoconfig | stateful);
      rapid-commit;
      req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server |
        sip-domain | sip-server |time-zone | vendor-spec);
      retransmission-attempt number;
      update-router-advertisement {
        interface interface-name;
      }
      update-server;
    }
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    mtu value;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
      input input-name;
      output output-name;
    }
    rpf-check {
      fail-filter filter-name;
      mode {
        loose;
      }
    }
    sampling {
      input;
      output;
    }
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
  }
  iso {
    address source-address;
    mtu value;
  }
  mlfr-end-to-end {
    bundle bundle-name;
  }
  mlfr-uni-nni {
    bundle bundle-name;
  }
  mlppp {
    bundle bundle-name;
  }
  mpls {
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
    }
    mtu mtu-value;
    policer {
      input input-name;
      output output-name;
    }
  }
  tcc {
    policer {
      input input-name;
      output output-name;
    }
    proxy {
```


Copyright © 2016, Juniper Networks, Inc.

```
            inet-address inet-address;
        }
        remote {
            inet-address inet-address;
            mac-address mac-address;
        }
    }
    vpls {
        filter {
            group number;
            input filter-name;
            input-list [filter-name];
            output filter-name;
            output-list [filter-name];
        }
        policer {
            input input-name;
            output output-name;
        }
    }
}
input-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number ;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interface-shared-with {
    psd-name;
}
native-inner-vlan-id value;
(no-traps | traps);
output-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        no-rfc2486;
        passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
```

Copyright © 2016, Juniper Networks, Inc.

```
            access-profile name;
            default-password password;
            local-name name;
            local-password password;
            no-rfc2486;
            passive;
          }
        }
        proxy-arp (restricted | unrestricted);
        radio-router {
          bandwidth number;
          credit {
            interval number;
          }
          data-rate number;
          latency number;
          quality number;
          resource number;
          threshold number;
        }
        swap-by-poppush;
        traps;
        vlan-id vlan-id;
        vlan-id-range vlan-id-range;
        vlan members [vlan-id];
        vlan-id-range vlan-id1-vlan-id2;
        vlan-tags {
          (inner vlan-id | inner-range vlan-id1-vlan-id2);
          inner-list [vlan-id];
          outer vlan-id;
        }
      }
      vlan-tagging;
    }
  }
```

Related Documentation

- Understanding Interfaces on page 2441

PoE Configuration Statement Hierarchy

To configure Power over Ethernet options, use the configuration statements in the **poe** configuration hierarchy. Statement descriptions that are exclusive to the SRX Series devices running Junos OS are described in this chapter.

```
poe {
  fpc slot-number {
    maximum-power watts;
    priority (high | low);
  }
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
```


Copyright © 2016, Juniper Networks, Inc.

```
        telemetries {
          disable;
          duration hours;
          interval minutes;
        }
      }
      management (class | static);
    }
```

Related Documentation

- Example: Configuring PoE on All Interfaces on page 2613

accept-source-mac

Syntax

```
accept-source-mac {
  mac-address mac-address;
}
```

Hierarchy Level [edit interfaces *interface-name* unit logical-unit-number]

Release Information Statement introduced in Junos OS Release 11.4.

Description For Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interfaces, specify the MAC addresses from which the interface can receive packets. Ensure that you update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. If you do not update the MAC address, the interface cannot receive packets from the new card.

NOTE:

- Software-based MAC limiting is supported on SRX300, SRX320, and SRX340 devices. A maximum of 32 MAC addresses is supported per device.

Options ***mac-address*** —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: ***nn:nn:nn:nn:nn:nn*** (for example, **00:11:22:33:44:55**) or ***nnnn:nnnn:nnnn*** (for example, **0011.2233.4455**). You can configure up to 32 source addresses. To specify more than one address, include multiple ***mac-addresses*** in the **source-address-filter** statement.

Required Privilege Level interface—To view this statement in the configuration..
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

access-point-name

Syntax **access-point-name** *apn*;

Hierarchy Level **[edit interfaces** *interface-name* **cellular-options gsm-options profiles** *profile-name***]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the access point name (APN) provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options ***apn***—Access point name.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *3G Wireless Modem Overview*

apply-groups

Syntax apply-groups;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio-router]

Release Information Statement introduced in Junos OS Release 9.6.
Statement modified in Junos OS Release 15.1.

Description Apply the groups from which to inherit configuration data. If **radio-router** is set without any other attributes specified, the first four values become 100 and threshold stays at 10, and capacity, margin, and delay are deprecated. If **radio-router** is set, do not change the OSPF reference-bandwidth value because this generates an incorrect link cost.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring PPPoE-Based Radio-to-Router Protocols on page 2652
- *Configuring Tunnel Interfaces on MX Series Routers*

 Copyright © 2016, Juniper Networks, Inc.

arp-resp

Syntax arp-resp (restricted|unrestricted);

Hierarchy Level [edit interfaces *interfaces-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure Address Resolution Protocol (ARP) response on the interface.

Options
- **restricted**—Enable restricted proxy ARP response on the interface. This is the default.
- **unrestricted**—Enable unrestricted ARP response on the interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Static ARP Entries on Ethernet Interfaces on page 2545

authentication-method (Interfaces)

Syntax **authentication-method (pap | chap | none);**

Hierarchy Level **[edit interfaces *interface-name* cellular-options gsm-options profiles *profile-name*]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the authentication method for connection to a Global System for Mobile Communications (GSM) cellular network.

Options
- **pap**—Password Authentication Protocol.
- **chap**—Challenge Handshake Authentication Protocol.
- **none**—No authentication method is used.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *3G Wireless Modem Overview*

Copyright © 2016, Juniper Networks, Inc.

bandwidth (Interfaces)

Syntax bandwidth *bandwidth*;

Hierarchy Level [edit interfaces *interface-name* radio-router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This option controls the weight of the current (vs. maximum) data rate (value 0–100).

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

bundle (Interfaces)

Syntax bundle *bundle-name*;

Hierarchy Level **[edit interfaces *interface-name* unit *logical-unit-number* family mlppp]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify the logical interface name the link joins.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

cbr rate

Syntax cbr rate;

Hierarchy Level [edit interfaces interface-name atm-options vpi vpi-identifier shaping]

Release Information Command introduced in Junos OS Release 9.5.

Description For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.

Options
- CBR Value—Constant bandwidth utilization (range: 33,000 through 1,199,920)
- CDVT—Cell delay variation tolerance in microseconds (range: 1 through 9999)

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

cellular-options

Syntax
```
cellular-options {
  roaming-mode (home only | automatic)
  gsm-options {
    select-profile profile-name;
    profiles {
      profile-name {
        sip-user-id simple-ip-user-id;
        sip-password simple-ip-password;
        access-point-name apn;
        authentication-method (pap | chap | none);
      }
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure options for connecting a 3G wireless modem interface to a cellular network.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *3G Wireless Modem Configuration Overview*

Copyright © 2016, Juniper Networks, Inc.

classifiers (CoS)

Syntax

```
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
    forwarding-class forwarding-class-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-point alias-or-bit-string ;
    }
    import (default | user-defined;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.2

Description Configure a user-defined behavior aggregate (BA) classifier.

Options

- *classifier-name*—User-defined name for the classifier.
- import (default | *user-defined*)—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type **dscp** and you specify **import default**, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify **import mymap**, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named **mymap**.
- forwarding-class *class-name*—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
- loss-priority *level*—Specify a loss priority for this forwarding class: **high**, **low**, **medium-high**, **medium-low**.
- code-points (*alias* | *bits*)—Specify a code-point alias or the code points that map to this forwarding class.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

client-identifier (Interfaces)

Syntax

```
client-identifier {
    (ascii string | hexadecimal string);
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* dhcp]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify an ASCII or hexadecimal identifier for the Dynamic Host Configuration Protocol (DHCP) client. The DHCP server identifies a client by a client-identifier value.

Options

- **ascii** ***ascii*** —Identifier consisting of ASCII characters.
- **hexadecimal** ***hexadecimal*** —Identifier consisting of hexadecimal characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

code-points (CoS)

Syntax code-points [*aliases*] [*bit-patterns*];

Hierarchy Level [edit class-of-service classifiers (dscp | ieee-802.1) *classifier-name* forwarding-class *class-name* loss-priority *level*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure one or more code-point aliases or bit sets to apply to a forwarding class.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP code points or code point aliases.

Options ***aliases***—Name of the alias or aliases.

bit-patterns—Value of the code-point bits, in decimal form.

Required Privilege Level interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441
- Example: Configuring BA Classifiers on Transparent Mode Devices on page 2941

Copyright © 2016, Juniper Networks, Inc.

compression-device (Interfaces)

Syntax compression-device *name*;

Hierarchy Level [edit interfaces *interface-name* unit (Interfaces) *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the compression interface for voice services traffic.

Options *name*—Name of the AC.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

credit (Interfaces)

Syntax
```
credit {
  interval number;
}
```

Hierarchy Level [edit interfaces *interface-name* radio—router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This parameter controls credit-based scheduling parameters and includes an interval option to set the grant rate interval to a value between 1–60 seconds.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

data-rate

Syntax data-rate *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio-router]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the weight of the resource factor when calculating an effective data rate.

Options ***weight***—Factor used to calculate data rate.
Range: 0 through 100
Default: 100

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

disable (PoE)

Syntax disable;

Hierarchy Level [edit poe interface (all | *interface-name*)]
[edit poe interface (all | *interface-name*) telemetries]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disables the PoE capabilities of the port. If PoE capabilities are disabled for a port, the port operates as a standard network access port. If the disable statement is specified after the telemetries statement, logging of PoE power consumption for the port is disabled. To disable monitoring and retain the stored interval and duration values for possible future use, you can specify the disable sub statement in the sub stanza for telemetries. Similarly for retaining the port configuration but disabling the PoE feature on the port, disable can be used in sub stanza for interface.

Default The PoE capabilities are automatically enabled when a PoE interface is set. Specifying the telemetries statement enables monitoring of PoE per-port power consumption.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Disabling a PoE Interface on page 2618

Copyright © 2016, Juniper Networks, Inc.

dhcp (Interfaces)

Syntax

```
dhcp {
  client-identifier {
    (ascii string | hexadecimal string);
  }
  lease-time (length | infinite);
  retransmission-attempt value;
  retransmission-interval seconds;
  server-address server-address;
  update-server;
  vendor-id vendor-id ;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the Dynamic Host Configuration Protocol (DHCP) client.

Options The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

duration (PoE)

Syntax duration *hours*;

Hierarchy Level [edit poe interface (all | *interface-name*) telemetries]

Release Information Statement introduced in Junos OS Release 9.5.

Description Modifies the duration for which telemetry records are stored. If telemetry logging continues beyond the specified duration, the older records are discarded one by one as new records are collected.

Options hours— Hours for which telemetry data should be retained.

Range: 1 through 24 hours

Default: 1 hour

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE on All Interfaces on page 2613

Copyright © 2016, Juniper Networks, Inc.

encapsulation (Interfaces)

Syntax encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | frame-relay-port-ccc | vlan-ccc | vlan-vpls);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify logical link layer encapsulation.

Options

- **cisco-hdlc**—For normal mode (when the device is using only one B-channel). Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points
- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623

Copyright © 2016, Juniper Networks, Inc.

family inet (Interfaces)

Syntax

```
inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address (source–address/prefix) {
    arp destination-address {
      (mac mac-address | multicast-mac multicast-mac-address);
      publish publish-address;
    }
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key-value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds
      (preempt <hold-time seconds> | no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
          priority-cost value;
        }
      }
      virtual-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  dhcp {
    client-identifier {
```


Copyright © 2016, Juniper Networks, Inc.

```
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      user-id (ascii string| hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp arp-name;
    input input-name;
    output output-name;
  }
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
    simple-filter;
  }
  targeted-broadcast {
    (forward-and-send-to-re |forward-only);
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
  }
```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement introduced in a prior release of Junos OS.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.

NOTE: You use family inet to assign an IPv4 address. You use family inet6 to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

family inet6

Syntax

```
inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      ouput;
    }
  }
  address source–address/prefix {
    eui-64;
    ndp address {
      (mac mac-address | multicast-mac multicast-mac-address);
      publish;
    }
    preferred;
    primary;
    vrrp-inet6-group group_id {
      (accept-data | no-accept-data);
      advertisements-threshold number;
      authentication-key value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      (preempt <hold-time seconds>| no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold value;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
        }
      }
      virtual-inet6-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  (dad-disable | no-dad-disable);
  dhcpv6-client {
    client-ia-type (ia-na | ia-pd);
    client-identifier duid-type (duid-ll | duid-llt | vendor);
```

Copyright © 2016, Juniper Networks, Inc.

```
    client-type (autoconfig | stateful);
    rapid-commit;
    req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-domain
      | sip-server |time-zone | vendor-spec);
    retransmission-attempt number;
    update-router-advertisement {
      interface interface-name;
    }
    update-server;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  nd6-stale-time seconds;
  no-neighbor-learn;
  policer {
    input input-name;
    output output-name;
  }
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
  }
  unnumbered-address {
    interface-name;
    preferred-source-address preferred-source-address;
  }
}
```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement supported in Junos 10.2 for SRX Series devices.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.

NOTE: You use family inet6 to assign an IPv6 address. You use family inet to assign an IPv4 address. An interface can be configured with both an IPv4 and IPv6 address.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

flag (Interfaces)

Syntax flag

Hierarchy Level [edit interfaces interface-name traceoptions]

Release Information Statement introduced in Junos OS Release 10.1.

Description Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options
- **all**—Enable all interface trace flags.
- **event** —Trace interface events.
- **cache**—Enable interface flags for Web filtering cache maintained on the routing table.
- **enhanced**—Enable interface flags for processing through Enhanced Web Filtering.
- **ipc**—Trace interface IPC messages.
- **media**—Trace interface media changes.
- **critical**—Trace critical events.
- **major**—Trace major events.

NOTE:
- MTU is limited to 1518 on this interface.
- Cache and enhanced options are applicable only to Enhanced Web Filtering.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

flexible-vlan-tagging (Interfaces)

Syntax flexible-vlan-tagging;

Hierarchy Level [edit interfaces *interface*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Simultaneously supports transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

Options **native-vlan-id**—Configures a VLAN identifier for single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring VLAN Tagging on page 2485

flow-control (Interfaces)

Syntax **(flow-control | no-flow-control);**

Hierarchy Level **[edit interfaces *interface-name* fastether-options],**
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* redundant-ether-options]

Release Information Statement modified in Junos OS Release 9.2.

Description For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the device to the remote side of the connection. Enabling flow control is useful when the device is a Gigabit Ethernet switch.

Default Flow control is the default behavior.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

flow-monitoring (Services)

Syntax

```
flow-monitoring {
  version9 {
    template template-name {
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      ipv4-template;
      ipv6-template;
      option-refresh-rate {
        packets packets;
        seconds seconds;
      }
      template-refresh-rate {
        packets packets;
        seconds seconds;
      }
    }
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure flow monitoring.

Options **version9**—Version 9 configuration.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

forwarding-classes (CoS)

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- *class-name*—Display the forwarding class name assigned to the internal queue number.

  NOTE: This option is supported only on high-end SRX Series devices, including the SRX1500, SRX5400, SRX5600, and SRX5800.

  NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **policing-priority**—Layer 2 policing. One forwarding class can be configured as **premium** and others are configured as **normal**.
- **priority**—Fabric priority value:
  - **high**—Forwarding class's fabric queuing has high priority.
  - **low**—Forwarding class's fabric queuing has low priority.
- *queue-number*—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, either **high** or **low**.

  NOTE: The spu-priority option is only supported on SRX1500, and SRX5000 line devices.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppQoS on page 589

fpc (Interfaces)

Syntax fpc;

Hierarchy Level [edit interfaces pic-set pic-set-name]

Release Information Command introduced in Junos OS Release 9.6.

Description Sets the PIC bundle and the FPC slot.

Options
- ***apply-groups***—Inherit configuration data from these groups.
- ***apply-groups-except***—Do not inherit configuration data from these groups.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]
[edit interfaces **interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Interfaces Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

gsm-options

Syntax

```
gsm-options {
  select-profile profile-name;
  profiles {
    profile-name {
      sip-user-id simple-ip-user-id;
      sip-password simple-ip-password;
      access-point-name apn;
      authentication-method (pap | chap | none);
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* cellular-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the 3G wireless modem interface to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *3G Wireless Modem Overview*

guard-band (PoE)

Syntax guard-band *watts*;

Hierarchy Level [edit poe]

Release Information Statement introduced in Junos OS Release 9.5.

Description Reserves the specified amount of power for the SRX Series device in case of a spike in PoE consumption.

Options ***watts***—Amount of power to be reserved for the SRX Series device in case of a spike in PoE consumption.

Range: 0 through 19 W

Default: 0 W

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Power over Ethernet on page 2611

Copyright © 2016, Juniper Networks, Inc.

hub-assist

Syntax hub-assist *weight*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio-router]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the weight of the resource factor when calculating an effective interface bandwidth.

Options ***weight***—Factor used to calculate interface bandwidth.

Range: 0 through 100

Default: 100

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE-Based Radio-to-Router Protocols on page 2652

inline-jflow (Forwarding Options)

Syntax

```
inline-jflow {
  flow-export-rate number;
  source-address ip-address;
}
```

Hierarchy Level [edit forwarding-options sampling instance *instance- name* family inet output]
[edit forwarding-options sampling instance *instance- name* family inet6 output]

Release Information Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.

Description Specify Inline processing of sampled packets.

Options

- **flow-export-rate *value***—Flow export rate of monitored packets in kpps. The range is from 1 through 400.
- **source-address *address***—Address to use for generating monitored packets.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

interface (PIC Bundle)

Syntax interface *interface-name*;

Hierarchy Level [edit interfaces pic-set pic-set-name]

Release Information Command introduced in Junos OS Release 9.6.

Description Sets the PIC bundle and the interface.

Options
- ***apply-groups***— Groups from which to inherit configuration data.
- ***apply-groups-except***— Do not inherit configuration data from these groups.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

interface (PoE)

Syntax

```
interface (all | interface-name) {
  disable;
  maximum-power watts;
  priority (high | low);
  telemetries {
    disable;
    duration hours;
    interval minutes;
  }
}
```

Hierarchy Level [edit poe]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable a PoE interface for a PoE port. The PoE interface must be enabled in order for the port to provide power to a connected powered device.

Default The PoE interface is enabled by default

Options

- **all**— Apply the configuration to all interfaces on the SRX Series device that have not been explicitly configured otherwise.
- **interface-name**— Explicitly configure a specific interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Power over Ethernet on page 2611

Copyright © 2016, Juniper Networks, Inc.

interfaces (CoS)

Syntax

```
interfaces interface-name {
  input-traffic-control-profile profile-name;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
    scheduler-map scheduler-map;
    shaping-rate bps;
    unit  logical-unit-number  {
      adaptive-shaper adaptive-shaper-name;
      classifiers {
        (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
      }
      forwarding-class class-name;
      input-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      loss-priority-maps {
        frame-relay-de {
          (lpmap-name | default);
        }
      }
      output-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      rewrite-rules {
        (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      }
      scheduler-map scheduler-map-name;
        shaping-rate {
        rate;
      }
      vc-shared-scheduler;
      virtual-channel-group group-name;
    }
  }
}
```

Hierarchy Level [edit class-of-service interface *interface-name* unit *number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Associate the class-of-service configuration elements with an interface.

Options interface *interface-name* unit *number*—The user-specified interface name and unit number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

interval (Interfaces)

Syntax interval *seconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio-router credit]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the frequency that the router generates credit announcement messages.

Options ***seconds***—Interval between PADG credit announcements for each session.
Range: 0 through 60
Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

interval (PoE)

Syntax interval *minutes*;

Hierarchy Level [edit poe interface (all | *interface-name*) telemetries]

Release Information Statement introduced in Junos OS Release 9.5.

Description Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.

Options ***minutes***—Interval at which data is logged.
Range: 1 through 30 minutes
Default: 5 minutes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

ipv4-template (Services)

Syntax ipv4-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify that the flow monitoring version 9 template is used only for IPv4 records.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

ipv6-template (Services)

Syntax ipv6-template;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify that the flow monitoring version 9 template is used only for IPv6 records.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

lacp (Interfaces)

Syntax

```
lacp {
  port-priority port-number;
}
```

Hierarchy Level [edit interfaces *interface-name* redundant-ether-options]

Release Information Statement introduced in Junos OS Release 10.2.

Description For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).

Options

- **active**—Initiate transmission of LACP packets.
- **passive**—Respond to LACP packets.

Default: If you do not specify **lacp** as either **active** or **passive**, LACP remains off (the default).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding LACP on Standalone Devices on page 2571

latency (Interfaces)

Syntax latency *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio—router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This option controls the latency weight (value 0–100).

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

lease-time

Syntax lease-time (*length* | infinite);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet dhcp]

Release Information Statement introduced in Junos OS Release 9.2.

Description Request a specific lease time for the IP address.

Default If no lease time is requested by client, then the server sends the lease time. The default lease time on a Junos OS DHCP server is one day.

Options ***seconds*** —Request a lease time of a specific duration.

Range: 60 through 2147483647 seconds

infinite—Request that the lease never expire.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

line-rate (Interfaces)

Syntax line-rate

Hierarchy Level [edit interfaces *interfaces name* shdsl-options]

Release Information Command introduced in Junos OS Release 10.0.

Description Specify a line rate for an G.SHDSL interface.

Options
- **auto**— Automatically selects a line rate.
- **value** — Select the values between 192 kbps and 22784 kbps for the speed of transmission of data on the G.SHDSL connection.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring the G.SHDSL Interface on SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

link-speed (Interfaces)

Syntax **link-speed** ***speed*;**

Hierarchy Level **[edit interfaces** ***interface-name*** **redundant-ether-options]**

Release Information Statement modified in Junos OS Release 9.2.

Description For redundant Ethernet interfaces in a chassis cluster only, set the required link speed.

Options ***speed*** —For redundant Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

loopback (Interfaces)

Syntax **(loopback | no-loopback);**

Hierarchy Level **[edit interfaces** ***interface-name*** **redundant-ether-options]**

Release Information Statement modified in Junos OS Release 9.2.

Description For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces, enable or disable loopback mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

loss-priority (CoS Loss Priority)

Syntax loss-priority *level* code-points[*values*];

Hierarchy Level [edit class-of-service loss-priority-maps frame-relay-de *map-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Map CoS values to a loss priority.

Options ***level*** can be one of the following:

- **high**—Packet has high loss priority.
- **medium-high**—Packet has medium-high loss priority.
- **medium-low**—Packet has medium-low loss priority.
- **low**—Packet has low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

loss-priority (CoS Rewrite Rules)

Syntax loss-priority *level*;

Hierarchy Level [edit class-of-service rewrite-rules *type rewrite-name* forwarding-class *class-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and packet loss priority (PLP). The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern.

Options ***level*** can be one of the following:

- **high**—The rewrite rule applies to packets with high loss priority.
- **low**—The rewrite rule applies to packets with low loss priority.
- **medium-high**—The rewrite rule applies to packets with medium-high loss priority.
- **medium-low**—The rewrite rule applies to packets with medium-low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

loss-priority-maps (CoS Interfaces)

Syntax

```
loss-priority-maps {
  frame-relay-de (map-name | default);
}
```

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Assign the loss priority map to a logical interface.

Options

- **default**—Apply default loss priority map. The default map contains the following:

  ```
  loss-priority low code-point 0;
  loss-priority high code-point 1;
  ```

- ***map-name***—Name of loss priority map to be applied.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

loss-priority-maps (CoS)

Syntax

```
loss-priority-maps {
  frame-relay-de loss-priority-map-name {
    loss-priority (high | low | medium-high | medium-low) {
      code-points [bit-string];
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.2.

Description Map the loss priority of incoming packets based on CoS values.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

management (PoE)

Syntax management (class | static);

Hierarchy Level [edit poe]

Release Information Statement introduced in Junos OS Release 9.5.

Description Designates how the SRX Series device allocates power to the PoE ports.

Default static

Options
- **static**—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port.
- **class**—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE 802.3 AF standard.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PoE on All Interfaces on page 2613

maximum-power (PoE)

Syntax maximum-power *watts*;

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.5.

Description Maximum amount of power that can be supplied to the port.

Default 15.4 W

Options **Watts**—The maximum number of watts that can be supplied to the port.

Range —0 through 15.4

Default—15.4 W

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PoE on All Interfaces on page 2613

Copyright © 2016, Juniper Networks, Inc.

media-type (Interfaces)

Syntax media-type

Hierarchy Level [edit interfaces *interface-name* media-type]

Release Information Command introduced in Junos OS Release 10.2.

Description Configure the operating modes for the 2-Port 10 Gigabit Ethernet XPIM.

Options
- copper
- fiber

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

minimum-links (Interfaces)

Syntax **minimum-links** *number*;

Hierarchy Level **[edit interfaces** *interface-name* **redundant-ether-options]**

Release Information Statement added in Junos OS Release 10.1.

Description For redundant Ethernet interfaces configured as 802.3ad redundant Ethernet interface link aggregation groups (LAGs) in a chassis cluster only, set the required minimum number of physical child links on the primary node that must be working to prevent the interface from being down. Interfaces configured as redundant Ethernet interface LAGs typically have between 4 and 16 physical interfaces, but only half, those on the primary node, are relevant to the minimum-links setting.

If the number of operating interfaces on the primary node falls below the configured value, it will cause the interface to be down even if some of the interfaces are still working.

Options *number*—For redundant Ethernet interface link aggregation group links, specify the number of physical child links on the primary node in the redundant Ethernet interface that must be working. The default **minimum-links** value is 1. The maximum value is half of the total number of physical child interfaces bound to the redundant Ethernet interface being configured or 8, whichever is smaller.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

native-vlan-id (Interfaces)

Syntax native-vlan-id*vlan-id*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.

Options ***vlan-id***—Configure a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration

Related Documentation
- Understanding Interfaces on page 2441

next-hop-tunnel

Syntax **next-hop-tunnel** ***gateway-address*** **ipsec-vpn** ***vpn-name***;

Hierarchy Level **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*** **family** ***family-name*****]**

Release Information Statement introduced in Junos OS Release 9.5.

Description For the secure tunnel (**st**) interface, create entries in the Next-Hop Tunnel Binding (NHTB) table, which is used to map the next-hop gateway IP address to a particular IP Security (IPsec) Virtual Private Network (VPN) tunnel. NHTB allows the binding of multiple IPsec VPN tunnels to a single IPsec tunnel interface.

Options
- ***gateway-address***—Next-hop gateway IP address.
- **ipsec-vpn** ***vpn-name*** —VPN to which the next-hop gateway IP address is mapped.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

no-dns-propagation

Syntax no-dns-propagation;

Hierarchy Level [edit interface *interface-name* unit *unit-number* family *inet* | *inet6 dhcp-client*]

Release Information Statement introduced in Junos OS Release 12.1X47-D35.

Description Disable the propagation of DNS information to the kernel.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

option-refresh-rate (Services)

Syntax option-refresh-rate

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the option refresh rate.

Options
- **packets**—Specify the number of packets. The range is from 1 through 480,000.
- **seconds**—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461

Copyright © 2016, Juniper Networks, Inc.

pic-mode (Chassis T1 Mode)

Syntax pic-mode (clear-channel);

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* ethernet]

Release Information Statement added in Junos OS Release 10.2.

Description Configure normal T1 mode or channelized T1 mode.

Options

- **clear-channel**—(default) Normal T1 mode.
- **ct1**—Channelized T1 mode.

NOTE: When chassis clustering is enabled, it is necessary to indicate in the command which node is being configured. In such circumstances, the edit chassis fpc command becomes edit chassis node *node-id* fpc.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

periodic (Interfaces)

Syntax **periodic (fast | slow);**

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options lacp]**

Release Information Statement introduced in Junos OS Release 10.2.

Description For redundant Ethernet interfaces in a chassis cluster only, configure the interval at which the interfaces on the remote side of the link transmit link aggregation control protocol data units (PDUs) by configuring the **periodic** statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval.

Options

- **fast**—Transmit link aggregation control PDUs every second.
- **slow**—Transmit link aggregation control PDUs every 30 seconds.

Default: fast

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

ppp-over-ether

Syntax ppp-over-ether;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number encapsulation*]

Release Information Statement introduced before Junos OS Release 11.2.
This encapsulation is supported for Redundant Ethernet interface in Junos OS Release 11.2.

Description This encapsulation is used for underlying interfaces of pp0 interfaces. This encapsulation is supported on Fast Ethernet interface, Gigabit Ethernet interface, and Redundant Ethernet interface. When Redundant Ethernet interface is used as underlying interface, an existing pppoe session can be continued in case of failover.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

pppoe

Syntax

```
pppoe {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable users to connect to a network of hosts over a bridge or access concentrator.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the Point-to-Point Protocol over Ethernet process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

pppoe-options

Syntax

```
pppoe-options {
  access-concentrator name ;
  auto-reconnect seconds;
  (client | server);
  ignore-eol-tag;
  service-name name;
  underlying-interface interface-name;
}
```

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

Release Information Statement modified in Junos OS Release 12.3X48 to include **ignore-eol-tag** statement.

Description Configure PPP over Ethernet-specific interface properties.

Options

access-concentrator *name*—(SRX Series devices with Point-to-Point Protocol over Ethernet (PPPoE) interfaces) Configure the name of the access concentrator. If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.

auto-reconnect *seconds*—Configure the amount of time to wait before reconnecting after a session has terminated.

client —Configure the device to operate in the PPPoE client mode.

idle-timeout *seconds*—Configure the maximum time that a session can be idle.

ignore-eol-tag—Disable the **End-of-List** tag to process the tags after the **End-of-List** tag in a PPPoE Active Discovery Offer (PADO) packet.

service-name *name*—Configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.

server—Configure the device to operate in the PPPoE server mode.

underlying-interface *interface-name*—Configure the interface on which PPP over Ethernet is running.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PPPoE Interfaces on page 2635

 Copyright © 2016, Juniper Networks, Inc.

priority (PoE)

Syntax priority (high | low);

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.5.

Description Sets the priority of individual ports. When it is not possible to maintain power to all connected ports, lower-priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower-priority port will be powered off automatically if available power is insufficient to power on the higher-priority port. Note that for ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.

Default low

Options value—high or low:

- **high**—Specify that this port is to be treated as high priority in terms of power allocation
- **low**—Specify that this port is to be treated as low priority in terms of power allocation.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE on All Interfaces on page 2613

Copyright © 2016, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order [accounting-method];
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order [accounting-method];
  address-assignment pool pool-name;
  authentication-order [ldap | none | password | radius | securid];
  authorization-order [jsrc];
  client client-name {
    chap-secret chap-secret;
    client-group [ group-names ];
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
```


Copyright © 2016, Juniper Networks, Inc.

```
radius {
  accounting-server [server];
  attributes {
    exclude {
      acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
      acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
      acc-loop-cir-id [access-request | accounting-start | accounting-stop];
      accounting-authentic [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-delay-time [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-session-id [access-request];
      accounting-terminate-cause [accounting-off];
      act-data-rate-dn [access-request | accounting-start | accounting-stop];
      act-data-rate-up [access-request | accounting-start | accounting-stop];
      act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      act-interlv-delay-up [access-request | accounting-start | accounting-stop];
      att-data-rate-dn [access-request | accounting-start | accounting-stop];
      att-data-rate-up [access-request | accounting-start | accounting-stop];
      called-station-id [access-request | accounting-start | accounting-stop];
      calling-station-id [access-request | accounting-start | accounting-stop];
      class [access-request | accounting-start | accounting-stop];
      delegated-ipv6-prefix [accounting-start | accounting-stop];
      dhcp-gi-address [access-request | accounting-start | accounting-stop];
      dhcp-mac-address [access-request | accounting-start | accounting-stop];
      dhcp-options [access-request | accounting-start | accounting-stop];
      downstream-calculated-qos-rate [access-request | accounting-start |
        accounting-stop];
      dsl-forum-attributes [access-request | accounting-start | accounting-stop];
      dsl-line-state [access-request | accounting-start | accounting-stop];
      dsl-type [access-request | accounting-start | accounting-stop];
      dynamic-iflset-name [accounting-start | accounting-stop];
      event-time-stamp [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      framed-interface-id [access-request | accounting-start | accounting-stop];
      framed-ip-address [access-request | accounting-start | accounting-stop];
      framed-ip-netmask [access-request | accounting-start | accounting-stop];
      framed-ip-route [access-request | accounting-start | accounting-stop];
      framed-ipv6-pool [accounting-start | accounting-stop];
      framed-ipv6-prefix [accounting-start | accounting-stop];
      framed-ipv6-route [accounting-start | accounting-stop];
      framed-pool [accounting-start | accounting-stop];
      input-filter [accounting-start | accounting-stop];
      input-gigapackets [accounting-stop];
      input-gigawords [accounting-stop];
      input-ipv6-gigawords [accounting-stop];
      input-ipv6-octets [accounting-stop];
      input-ipv6-packets [accounting-stop];
      interface-description [access-request | accounting-start | accounting-stop];
      l2c-downstream-data [access-request | accounting-start | accounting-stop];
      l2c-upstream-data [access-request | accounting-start | accounting-stop];
      max-data-rate-dn [access-request | accounting-start | accounting-stop];
      max-data-rate-up [access-request | accounting-start | accounting-stop];
      max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      max-interlv-delay-up [access-request | accounting-start | accounting-stop];
      min-data-rate-dn [access-request | accounting-start | accounting-stop];
```

Copyright © 2016, Juniper Networks, Inc.

```
            min-data-rate-up [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
            nas-identifier [access-request | accounting-start | accounting-stop];
            nas-port [access-request | accounting-off | accounting-on | accounting-start |
              accounting-stop];
            nas-port-id [access-request | accounting-start | accounting-stop];
            nas-port-type [access-request | accounting-start | accounting-stop];
            output-filter [accounting-start | accounting-stop];
            output-gigapackets [accounting-stop];
            output-gigawords [accounting-stop];
            output-ipv6-gigawords [accounting-stop];
            output-ipv6-octets [accounting-stop];
            output-ipv6-packets [accounting-stop];
            upstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];
          }
          ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
          }
        }
      authentication-server [server];
      radius-options {
        request-rate number;
        revert-interval seconds;
      }
      radius-server server-address {
        accounting-port port-number
        max-outstanding-requests number-of--outstanding-requests;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
      }
      service {
        accounting-order {
          activation-protocol;
          radius;
        }
      }
      session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
      }
    }
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Copyright © 2016, Juniper Networks, Inc.

Description Create a profile containing a set of attributes that define device management access.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441
- Understanding User Authentication for Security Devices on page 5173
- Layer 2 Transparent and Switching Overview on page 2893

profiles

Syntax

```
profiles {
  profile-name {
    sip-user-id simple-ip-user-id;
    sip-password simple-ip-password;
    access-point-name apn;
    authentication-method (pap | chap | none);
  }
}
```

Hierarchy Level [edit interfaces *interface-name* cellular-options gsm-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network. You can configure up to 16 profiles.

Options ***profile-name***—Name of the profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

promiscuous-mode (Interfaces)

Syntax promiscuous-mode;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Enable promiscuous mode on Layer 3 Ethernet interfaces. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or Services Processing Unit regardless of the destination MAC address of the packet.

You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and on aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 2551

quality (Interfaces)

Syntax quality *<value>*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* radio—router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This option controls relative link quality weight (value 0–100).

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

r2cp

Syntax

```
r2cp {
    command binary-file-path;
    disable;
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the Radio-to-Router Control Protocol (R2CP) used to exchange dynamic metric changes in the network that routers use to update the OSPF topologies.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the Radio-to-Router Control Protocol process.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

radio-router (Interfaces)

Syntax

```
radio-router {
  bandwidth number;
  credit {
    interval number;
  }
  data-rate number;
  latency number;
  quality number;
  resource number;
  threshold number;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 10.1.

Description Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

redundancy-group (Interfaces)

Syntax redundancy-group *number*;

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the redundancy group that a redundant Ethernet interface belongs to.

Options ***number*** —Number of the redundancy group that the redundant interface belongs to. Failover properties of the interface are inherited from the redundancy group.

Range: 1 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

redundant-ether-options

Syntax
```
redundant-ether-options {
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    periodic (fast | slow);
  }
  link-speed speed;
  (loopback | no-loopback);
  minimum-links number;
  redundancy-group number;
  source-address-filter mac-address;
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level **[edit interfaces *interface-name*]**

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Ethernet redundancy options for a chassis cluster.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Example: Enabling Eight Queue Class of Service on Redundant Ethernet Interfaces*
- *Example: Configuring Chassis Cluster Redundant Ethernet Interfaces for IPv4 and IPv6 Addresses*

Copyright © 2016, Juniper Networks, Inc.

redundant-parent (Interfaces Fast Ethernet)

Syntax redundant-parent *interface-name* ;

Hierarchy Level [edit interfaces *interface-name* fastether-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Fast Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options ***interface*** —Parent redundant interface of the Fast Ethernet interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

redundant-parent (Interfaces Gigabit Ethernet)

Syntax redundant-parent *interface-name*;

Hierarchy Level [edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Gigabit Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options ***interface*** —Parent redundant interface of the Gigabit Ethernet interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

resource (Interfaces)

Syntax resource *number*;

Hierarchy Level [edit interfaces *interface-name* radio—router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This option controls the resource weight (value 1–100).

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

retransmission-attempt

Syntax retransmission-attempt *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet dhcp]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the number of times the device retransmits a Dynamic Host Control Protocol (DHCP) packet if a DHCP server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.

Options *number*—Number of retransmit attempts..
Range: 0 through 6
Default: 4

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

retransmission-interval (Interfaces)

Syntax retransmission-interval *seconds*;

Hierarchy Level [edit **interfaces** *interface-name* unit *logical-unit-number* family *family-name* dhcp]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the time between successive retransmission attempts.

Options ***seconds*** —Number of seconds between successive retransmission.
Range: 4 through 64 seconds
Default: 4 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

roaming-mode

Syntax **roaming-mode (home-only | automatic)**

Hierarchy Level **[edit interfaces** *interface-name* **cellular-options]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify whether the 3G wireless modem interface can access networks other than the home network.

Options
- **home-only**—No roaming is allowed.
- **automatic**—Allows access to networks other than the home network. This is the default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

scheduler-map (CoS Virtual Channels)

Syntax scheduler-map *map-name*;

Hierarchy Level [edit class-of-service virtual-channel-groups *group-name virtual-channel-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Apply a scheduler map to this virtual channel.

Options ***map-name***—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- default (CoS) on page 1632
- shaping-rate (CoS Virtual Channels) on page 1649
- virtual-channel-group (CoS Interfaces) on page 1651
- virtual-channel-groups on page 1652
- virtual-channels on page 1650

select-profile

Syntax **select-profile** *profile-name*

Hierarchy Level **[edit interfaces** *interface-name* **cellular-options gsm-options]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Select the active profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options ***profile-name***—Name of a configured profile that is to be used to establish a data call.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

server-address

Syntax server-address *ip-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet dhcp]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.

Default The client accepts the first offer it receives from any DHCP server.

Options ***ip-address***—DHCP server address.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

shaping-rate (CoS Interfaces)

Syntax shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

Logical and physical interface traffic shaping is mutually exclusive. This means you can include the **shaping-rate** statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level or the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, but not both.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

Default If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options *rate*—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: For logical interfaces, 1000 through 32,000,000,000 bps.

For physical interfaces, 1000 through 160,000,000,000 bps.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

simple-filter (Interfaces)

Syntax simple-filter;

Hierarchy Level [edit interfaces *interfaces-name* unit *logical-unit-number* family *family-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Apply a simple filter to an interface. You can apply simple filters on ingress interfaces only.

Options input ***filter-name***: Name of one filter to evaluate when packets are received on the interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

sip-password

Syntax sip-password *simple-ip-password*;

Hierarchy Level [edit interfaces *interface-name* cellular-options gsm-options profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the password provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options ***simple-ip-password***—Password.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541
- *3G Wireless Modem Overview*

Copyright © 2016, Juniper Networks, Inc.

sip-user-id

Syntax **sip-user-id** *simple-ip-user-id*;

Hierarchy Level **[edit interfaces** *interface-name* **cellular-options gsm-options profiles** *profile-name***]**

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the username provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options ***simple-ip-user-id***—Username.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *3G Wireless Modem Overview*

Copyright © 2016, Juniper Networks, Inc.

source-address-filter (Interfaces)

Syntax **source-address-filter** *mac-address* ;

Hierarchy Level **[edit interfaces** *interface-name* **redundant-ether-options]**

Release Information Statement modified in Junos OS Release 9.2.

Description For redundant Ethernet interfaces, specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering.

Be sure to update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. Otherwise, the interface cannot receive packets from the new card.

NOTE:

- Software based MAC limiting is supported on SRX300, SRX320, and SRX340 devices.

  A maximum of 32 devices are supported per device.

Options *mac-address* —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: ***nn:nn:nn:nn:nn:nn:nn*** (for example, **00:11:22:33:44:55**) or ***nnnn:nnnn:nnnn*** (for example, **0011.2233.4455**). You can configure up to 64 source addresses. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Ethernet Interfaces on page 2541

 Copyright © 2016, Juniper Networks, Inc.

source-filtering (Interfaces)

Syntax **(source-filtering | no-source-filtering);**

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement modified in Junos OS Release 9.2.

Description For redundant Ethernet interfaces, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

By default, source address filtering is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541

speed (Interfaces)

Syntax speed (100m |10m | 1g);

Hierarchy Level [edit interfaces *interface-name* speed]

Release Information Command introduced in Junos OS Release 10.2.

Description Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.

Options
- 100m — Link speed of 100 Mbps
- 10g — Link speed of 10 Gbps
- 10m — Link speed of 10 Mbps
- 1g — Link speed of 1 Gbps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Ethernet Interfaces on page 2541
- Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 2598

Copyright © 2016, Juniper Networks, Inc.

telemetries (PoE)

Syntax

```
telemetries {
  disable;
  duration hours;
  interval minutes;
}
```

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.5.

Description Allow logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default.

Default If the telemetries statement is specified, logging is enabled with the default values for interval and duration.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE on All Interfaces on page 2613

template-refresh-rate (Services)

Syntax template-refresh-rate;

Hierarchy Level [edit services flow-monitoring version9 template *template-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the template refresh rate.

Options

- **packets**—Specify the number of packets. The range is from 1 through 480,000.
- **seconds**—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

threshold (Interfaces)

Syntax threshold <*value*>;

Hierarchy Level [edit interfaces *interface-name* radio–router]

Release Information Statement introduced in Junos OS Release 10.1.

Description This option controls the percentage of bandwidth change required for routing updates.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

traceoptions (Interfaces)

Syntax traceoptions

Hierarchy Level [edit interfaces interface-name traceoptions]

Release Information Command introduced in Junos OS Release 10.1.

Description Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options **flag** - Tracing parameters

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- PPPoE-Based Radio-to-Router Protocols Overview on page 2649

Copyright © 2016, Juniper Networks, Inc.

update-server

Syntax update-server;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet dhcp]

Release Information Statement introduced in Junos OS Release 9.2 for SRX Series devices
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

vbr rate

Syntax vbr rate;

Hierarchy Level [edit interfaces interface-name atm-options vpi vpi-identifier shaping]

Release Information Command introduced in Junos OS Release 9.5.

Description For ATM encapsulation only, define a variable bit rate bandwidth utilization in the traffic-shaping profile.

Options
- Burst Size—The maximum burst size that can be sent at the peak rate.
- Peak Rate—The maximum instantaneous rate at which the user will transmit.
- Sustained Rate—The average rate as measured over a long interval.
- CDVT—Cell Delay Variation Tolerance in microseconds (range: 1 – 9999).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

vdsl-profile

Syntax vdsl-profile

Hierarchy Level [edit interfaces interface-name vdsl-options]

Release Information Command introduced in Junos OS Release 10.1.

Description Configure the type of VDSL2 profiles. A profile is a table that contains a list of preconfigured VDSL2 settings.

Options
- Auto (default)
- 8a
- 8b
- 8c
- 8d
- 12a
- 12b
- 17a

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- VDSL2 Interface Support on SRX Series Devices on page 2503

vendor-id (Interfaces)

Syntax vendor-id *vendor-id* ;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* dhcp]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.

Options *vendor-id* —vendor class ID.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

vlan-tagging (Interfaces)

Syntax vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level [edit interfaces *interface*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.

Options **native-vlan-id**—Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

NOTE: The native-vlan-id can be configured only when either flexible-vlan-tagging mode or interface-mode trunk is configured.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Tagging on page 2485

 Copyright © 2016, Juniper Networks, Inc.

web-authentication (Interfaces)

Syntax

```
web-authentication {
  http;
  https;
  redirect-to-https;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* address *address*]

Release Information Statement introduced in Junos OS Release 9.2.
Support for **https** and **redirect-to-https** introduced for high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Enable the Web authentication process for firewall user authentication.

Options

http—Enable HTTP service.

https—Enable authentication through HTTPS.

redirect-to-https—Redirect Web authentication to HTTPS.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 125

Operational Commands

- clear dhcpv6 server binding (Local Server)
- clear ethernet-switching statistics mac-learning
- clear interfaces statistics swfabx
- clear ipv6 neighbors
- clear lacp statistics interfaces
- restart (Reset)
- show chassis fpc (View)
- show chassis hardware (View)
- show ethernet-switching mac-learning-log (View)
- show ethernet-switching table (View)
- show igmp-snooping route (View)
- show interfaces (SRX Series)
- show interfaces diagnostics optics
- show interfaces flow-statistics
- show interfaces queue
- show interfaces statistics (View)
- show interfaces terse zone
- show ipv6 neighbors
- show lacp interfaces (View)
- show lacp statistics interfaces (View)
- show oam ethernet link-fault-management
- show poe controller (View)
- show pppoe interfaces
- show pppoe statistics
- show poe telemetries
- show services accounting
- show services accounting aggregation (View)

Copyright © 2016, Juniper Networks, Inc.

- show services accounting aggregation template (View)
- show services accounting flow-detail (View)

Copyright © 2016, Juniper Networks, Inc.

clear dhcpv6 server binding (Local Server)

Syntax

clear dhcpv6 server binding
<all | *client-id* | *ip-address* | *session-id*>
<interface *interface-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 10.4.

Description

Clear the binding state of a DHCPv6 client from the client table on the DHCPv6 local server.

Options

- all—(Optional) Clear the binding state for all DHCPv6 clients.
- *client-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified client ID (option 1).
- *ip-address*—(Optional) Clear the binding state for the DHCPv6 client with the specified address.
- *session-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified session ID.
- interface *interface-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
- routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

clear

Related Documentation

- *show dhcpv6 server binding (View)*

Copyright © 2016, Juniper Networks, Inc.

clear ethernet-switching statistics mac-learning

Syntax clear ethernet-switching statistics mac-learning

Release Information Command introduced in Junos OS Release 10.1.

Description Clear the media access control (MAC) learning statistics.

Options
- **none**—Clear MAC learning statistics on all interfaces.
- **interface *interface-name***—(Optional) Clear MAC learning statistics on the specified interface.

Required Privilege Level view

Related Documentation
- show ethernet-switching table (View) on page 2808show ethernet-switching table

List of Sample Output
clear ethernet-switching statistics mac-learning on page 2762
clear ethernet-switching statistics mac-learning interface interface-name on page 2762

Sample Output

clear ethernet-switching statistics mac-learning

```
user@host> clear ethernet-switching statistics mac-learning
```

clear ethernet-switching statistics mac-learning interface interface-name

```
user@host> clear ethernet-switching statistics mac-learning interface interface-name
```


Copyright © 2016, Juniper Networks, Inc.

clear interfaces statistics swfabx

Syntax clear interfaces statistics <swfab0 | swfab1>

Release Information Command introduced in Junos OS Release 11.1.

Description Clears interface statistics for the specified swfab interface.

Required Privilege Level clear

Related Documentation
- show interfaces swfabx on page 2028

List of Sample Output clear interfaces statistics <swfab0 | swfab1> on page 2763

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces statistics <swfab0 | swfab1>

```
user@host> clear interfaces statistics <swfab0 | swfab1>
```

Copyright © 2016, Juniper Networks, Inc.

clear ipv6 neighbors

Syntax clear ipv6 neighbors
<all | host *hostname*>

Release Information Command introduced in Junos OS Release 12.1X45-D10.

Description Clear IPv6 neighbor cache information.

Options **none**—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level clear

Related Documentation
- show ipv6 neighbors on page 2861

List of Sample Output clear ipv6 neighbors on page 2764

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
11:11::2             00:19:e2:4b:61:83  deleted
   12:12::2             00:19:e2:4b:61:83  deleted
   10:1::2              00:00:0a:00:00:00  deleted
```


Copyright © 2016, Juniper Networks, Inc.

clear lacp statistics interfaces

Syntax **clear lacp statistics interfaces** <*interface-name*>

Release Information Command modified in Junos OS Release 10.2.

Description Clear the LACP statistics. If you do not specify an interface name, LACP statistics for all interfaces are cleared.

Options ***interface-name***—(Optional) Name of an interface.

Required Privilege Level clear

Related Documentation

- show lacp statistics interfaces (View) on page 2867
- Verifying LACP on Redundant Ethernet Interfaces on page 2579

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

restart (Reset)

Syntax

restart
<application-identification |application-security |audit-process |commitd-service |chassis-control | class-of-service |database-replication |datapath-trace-service |ddns |dhcp |dhcp-service |dynamic-flow-capture |disk-monitoring |event-processing | ethernet-connectivity-fault-management |ethernet-link-fault-management |extensible-subscriber-services |fipsd |firewall |firewall-authentication-service |general-authentication-service |gracefully |gprs-process |idp-policy |immediately |interface-control | ipmi |ipsec-key-management |jflow-service |jnu-management |jnx-wmicd-service |jsrp-service |kernel-replication |l2-learning |l2cpd-service |lacp |license-service |logical-system-service |mib-process |mountd-service |named-service |network-security |network-security-trace |nfsd-service |ntpd-service |pgm |pic-services-logging |profilerd |pki-service |remote-operations |rest-api |routing |sampling |sampling-route-record |scc-chassisd |secure-neighbor-discovery |security-intelligence |security-log |services |service-deployment |simple-mail-client-service |soft |snmp |static-routed |statistics-service |subscriber-management |subscriber-management-helper |system-log-vital |tunnel-oamd |uac-service |user-ad-authentication |vrrp |web-management >

Release Information

Command introduced before Junos OS Release 9.2

Description

Restart a Junos OS process.

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router to drop calls and interrupt transmission, resulting in possible loss of data.

Options

- application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
- application-security—(Optional) Restart the application security process.
- audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, for analyzing and tracking usage patterns, and for billing a user based upon the amount of time used or the type of services accessed.
- chassis-control—(Optional) Restart the chassis management process.
- class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
- commitd-service—(Optional) Restart the committed services.
- database-replication—(Optional) Restart the database replication process.
- datapath-trace-service—(Optional) Restart the Restart the packet path tracing process.
- ddns—(Optional) Restart the dynamic domain name system, which dynamically updates IP addresses for registered domain names.

Copyright © 2016, Juniper Networks, Inc.

- dhcp—(Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.
- dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.
- disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.
- dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on PIC3 monitoring services cards.
- ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- ethernet-link-fault-management—(Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- event-processing—(Optional) Restart the event process (eventd).
- extensible-subscriber-services—(Optional) Restart the extensible subscriber services process.
- fipsd—(Optional) Restart the fipsd services.
- firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and accepts or rejects packets that are transiting an interface on a router or switch.
- firewall-authentication-service—(Optional) Restart the firewall authentication service process.
- general-authentication-service—(Optional) Restart the general authentication process.
- gprs-process—(Optional) Restart the General Packet Radio Service (GPRS) process.
- gracefully—(Optional) Restart the software process.
- idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.
- immediately—(Optional) Immediately restart the software process.
- interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.
- ipmi—(Optional) Restart the intelligent platform management interface process.
- ipsec-key-management—(Optional) Restart the IPsec key management process.
- jflow-service—(Optional) Restart jflow service process.
- jnu-management—(Optional) Restart jnu management process.
- jnx-wmicd-service—(Optional) Restart jnx wmicd service process.
- jsrp-service—(Optional) Restart the Juniper Services Redundancy Protocol (jsrdp) process, which controls chassis clustering.

Copyright © 2016, Juniper Networks, Inc.

- kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
- lacp—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link. The LACP process allows link aggregation control instances to reach agreement on the identity of the LAG to which a link belongs, moves the link to that LAG, and enables the transmission and reception processes for the link to function in an orderly manner.
- l2cpd-service—(High-end SRX Series only) (Optional) Restart the Layer 2 Control Protocol (L2CP) process, which enables features such as L2 protocol tunneling and nonstop bridging.
- l2-learning—(Optional) Restart the Layer 2 (L2) address flooding and learning process.
- license-service—(Optional) Restart the feature license management process.
- logical-system-service—(Optional) Restart the logical system service process.
- mib-process—(Optional) Restart the MIB version II process, which provides the router's MIB II agent.
- mountd-service—(Optional) Restart the service for Network File System (NFS) mount requests.
- named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
- network-security—(Optional) Restart the network security process.
- network-security-trace—(Optional) Restart the network security trace process.
- nfsd-service—(Optional) Restart the remote NFS server process, which provides remote file access for applications that need NFS-based transport.
- ntpd-service—(Optional) Restart the Network Time Protocol (NTP) process.
- pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.
- pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.
- pki-service—(Optional) Restart the public key infrastructure (PKI) service process.
- profilerd—(Optional) Restart the profiler process.
- remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.
- rest-api—(Optional) Restart the rest api process.
- routing—(Optional) Restart the routing protocol process (rpd).
- sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

Copyright © 2016, Juniper Networks, Inc.

- sampling-route-record—(Optional) Restart the sampling route record process.
- scc-chassisd—(Optional) Restart the scc chassisd process.
- secure-neighbor-discovery—(Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.
- security-intelligence—(Optional) Restart security intelligence process.
- security-log—(Optional) Restart the security log process.
- service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.
- services—(Optional) Restart a service.
- simple-mail-client-service—(Optional) Restart the simple mail client service process.
- snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.
- static-routed—(Optional) Restart the static routed process.
- soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.
- statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.
- subscriber-management—(Optional) Restart the subscriber management process.
- subscriber-management-helper—(Optional) Restart the subscriber management helper process.
- system-log-vital—(Optional) Restart system log vital process.
- tunnel-oamd—(Optional) Restart the tunnel OAM process for L2 tunneled networks.
- uac-service—(Optional) Restart the Unified Access Control (UAC) process.
- user-ad-authentication—(Optional) Restart User ad Authentication process
- vrrp—(Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.
- web-management—(Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation
- *Restart Commands Overview*

List of Sample Output restart interfaces on page 2770

Output Fields When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```


Copyright © 2016, Juniper Networks, Inc.

show chassis fpc (View)

Syntax **show chassis fpc**
<**detail** < *fpc-slot* >| <node (*node-id* | local | primary)>> |
<node (*node-id* | local | primary)> |
<**pic-status** < *fpc-slot* >| <node (*node-id* | local | primary)>>

Release Information Command modified in Junos OS Release 9.2.
Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.

NOTE: On SRX5K-MPC3-40G10G (IOC3), all four PICs cannot be powered on. A maximum of two PICs can be powered on at the same time. By default, PIC0 and PIC1 are online.

Use the **set chassis fpc <slot> pic <pic> power off** command to choose the PICs you want to power on.
When you use the **set chassis fpc <slot> pic <pic> power off** command to power off PIC0 and PIC1, PIC2 and PIC3 are automatically turned on.
When you switch from one set of PICs to another set of PICs using the **set chassis fpc <slot> pic <pic> power off** command again, ensure that there is 60 seconds duration between the two actions, otherwise core files are seen during the configuration.
The Table 185 summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 276: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary

CLI Configuration	PIC Selection
Default (i.e. no CLI configuration)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1, PIC-2 and PIC-3 powered OFF	Online: PIC-0 Offline: PIC-1, PIC-2, PIC-3
PIC-0, PIC-2 and PIC-3 powered OFF	Online: PIC-1 Offline: PIC-0, PIC-2, PIC-3
PIC-0, PIC-1 and PIC-3 powered OFF	Online: PIC-2 Offline: PIC-0, PIC-1, PIC-3
PIC-0, PIC-1 and PIC-2 powered OFF	Online: PIC-3 Offline: PIC-0, PIC-1, PIC-2

Copyright © 2016, Juniper Networks, Inc.

Table 276: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary *(continued)*

CLI Configuration	PIC Selection
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1 and PIC-2 powered OFF	Online: PIC-0, PIC-3 Offline: PIC-1, PIC-2
PIC-0 and PIC-3 powered OFF	Online: PIC-2, PIC-1 Offline: PIC-0, PIC-3
PIC-0 and PIC-1 powered OFF	Online: PIC-2, PIC-3 Offline: PIC-0, PIC-1
All other combinations of PICs being powered OFF (Invalid)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3 Default PICs will be selected for the invalid combinations. Also, a system log message will be displayed to indicate the invalid combination PIC selection.

Description Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options

- none—Display status information for all FPCs.
- **detail**—(Optional) Display detailed FPC status information.
- ***fpc-slot*** —(Optional) Display information about the FPC in this slot.
- **node**—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Copyright © 2016, Juniper Networks, Inc.

- **pic-status**—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see ***fpc-slot***).

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

List of Sample Output
show chassis fpc on page 2774
show chassis fpc (SRX5600 and SRX5800 devices) on page 2774
show chassis fpc (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3) on page 2774
show chassis fpc detail 2 on page 2775
show chassis fpc pic-status (SRX5600 and SRX5800 devices) on page 2775
show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2) on page 2775
show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC) on page 2776
show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured) on page 2776
show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP) on page 2777
show chassis fpc pic-status(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured) on page 2777
show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices) on page 2777
show chassis fpc pic-status for HA(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3) on page 2778

Output Fields Table 186 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 277: show chassis fpc Output Fields

Field Name	Field Description
Slot or Slot State	Slot number and state. The state can be one of the following conditions: • **Dead**—Held in reset because of errors. • **Diag**—Slot is being ignored while the device is running diagnostics. • **Dormant**—Held in reset. • **Empty**—No FPC is present. • **Online**—FPC is online and running. • **Present**—FPC is detected by the device, but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either **Hardware Not Supported** or **Hardware Not In Right Slot**. FPC is coming up but not yet online. • **Probed**—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • **Probe-wait**—Waiting to be probed.
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.

Copyright © 2016, Juniper Networks, Inc.

Table 277: show chassis fpc Output Fields *(continued)*

Field Name	Field Description
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
Start Time	Time when the Routing Engine detected that the FPC was running.
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.
PIC type	(**pic-status** output only) Type of FPC.

Sample Output

show chassis fpc

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)     Memory    Utilization (%)
     Slot State       (C)    Total  Interrupt        DRAM (MB) Heap  Buffer
      0  Online          -------------------- CPU less FPC --------------------
      1  Online          ---------------------  Not Usable --------------------
      2  Online          -------------------- CPU less FPC --------------------
```

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)       Memory    Utilization (%)
Slot State           (C)   Total  Interrupt         DRAM (MB) Heap     Buffer
  0  Empty
  1  Empty
  2  Empty
  3  Online           37       3          0             1024        7     42
  4  Empty
  5  Empty
  6  Online           30       8          0             1024       23     30
  7  Empty
  8  Empty
  9  Empty
 10  Empty
 11  Empty
```

show chassis fpc

Copyright © 2016, Juniper Networks, Inc.

(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   CPU Utilization (%)  Memory

                                    Utilization (%)
Slot State           (C)  Total  Interrupt      1min   5min   15min  DRAM (MB)

                                  Heap     Buffer
  0  Online           36     20          0       20     19     19    1024

                                   4         26
  1  Online           35      8          0        8      8      8    2048

                                  12         14
  2  Online           40     21          0       20     20     20    3584

                                   5         13
```

Sample Output

show chassis fpc detail 2

```
user@host> show chassis fpc detail 2
Slot 2 information:
  State                                 Online
  Temperature                            37
  Total CPU DRAM                       1024 MB
  Total RLDRAM                          0 MB
  Total DDR DRAM                        0 MB
  Start time:                             2012-07-18 07:18:50 PDT
  Uptime:                                 4 days, 21 hours, 51 minutes, 59 seconds

  Max Power Consumption                 0 Watts
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
Slot 3   Online       SRX5k SPC
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
Slot 6   Online       SRX5k DPC 4x 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
```

Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 3   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 6   Offline      SRX5k SPC II
Slot 9   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Online       SRX5k IOC II
  PIC 0  Online       1x 100GE CFP
  PIC 2  Online       2x 40GE QSFP+
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status

Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis fpc pic-status

node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Offline      SRX5k SPC II
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 9   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP
  PIC 1  Online       10x 1GE(LAN) SFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Offline      SRX5k IOC II
```

show chassis fpc pic-status (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
```

Copyright © 2016, Juniper Networks, Inc.

```
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow

node1:
--------------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status for HA (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline

node1:
--------------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline
```


Copyright © 2016, Juniper Networks, Inc.

show chassis hardware (View)

Syntax **show chassis hardware**
<clei-models | detail | extensive | models | node (***node-id*** **| all | local | primary)>**

Release Information Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include **node** option.

Description Display chassis hardware information.

Options

- **clei-models**—(Optional) Display Common Language Equipment Identifier Code (CLEI) barcode and model number for orderable field-replaceable units (FRUs).
- **detail | extensive**—(Optional) Display the specified level of output.
- **models**—(Optional) Display model numbers and part numbers for orderable FRUs.
- **node**—(Optional) For chassis cluster configurations, display chassis hardware information on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Interface Naming Conventions on page 2446

Output Fields Table 188 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 278: show chassis hardware Output Fields

Field Name	Field Description
Item	Chassis component—Information about the backplane; power supplies; fan trays; Routing Engine; each Physical Interface Module (PIM)—reported as FPC and PIC—and each fan, blower, and impeller.
Version	Revision level of the chassis component.
Part Number	Part number for the chassis component.
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the device chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the device chassis.
Assb ID or Assembly ID	Identification number that describes the FRU hardware.

Copyright © 2016, Juniper Networks, Inc.

Table 278: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
FRU model number	Model number of FRU hardware component.
CLEI code	Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.
EEPROM Version	ID EEPROM version used by hardware component: 0x01 (version 1) or 0x02 (version 2).

Copyright © 2016, Juniper Networks, Inc.

Table 278: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
Description	Brief description of the hardware item: • Type of power supply. • Switch Control Board (SCB) Starting with Junos OS Release 12.1X47-D15, the SRX5K-SCBE (SCB2) is introduced. ◦ There are three SCB slots in SRX5800 devices. The third slot can be used for an SCB or an FPC. When an SRX5K-SCB was used , the third SCB slot was used as an FPC. SCB redundancy is provided in chassis cluster mode. ◦ With an SCB2, a third SCB is supported. If a third SCB is plugged in, it provides intra-chassis fabric redundancy. ◦ The Ethernet switch in the SCB2 provides the Ethernet connectivity among all the FPCs and the Routing Engine. The Routing Engine uses this connectivity to distribute forwarding and routing tables to the FPCs. The FPCs use this connectivity to send exception packets to the Routing Engine. ◦ Fabric connects all FPCs in the data plane. The Fabric Manager executes on the Routing Engine and controls the fabric system in the chassis. Packet Forwarding Engines on the FPC and fabric planes on the SCB are connected through HSL2 channels. ◦ SCB2 supports HSL2 with both 3.11 Gbps and 6.22 Gbps (SerDes) link speed and various HSL2 modes. When an FPC is brought online, the link speed and HSL2 mode are determined by the type of FPC. Starting with Junos OS Release 15.1X49-D10, the SRX5K-SCB3 (SCB3) with enhanced midplanes is introduced. ◦ All existing SCB software that is supported by SCB2 is supported on SCB3. ◦ SRX5K-RE-1800X4 (RE2). Mixed Routing Engine use is not supported. ◦ SCB3 works with the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), SRX5K-MPC3-40G10G (IOC3), and SRX5K-SPC-4-15-320 (SPC2) with current midplanes and the new enhanced midplanes. ◦ Mixed SCB use is not supported. If an SCB2 and an SCB3 are used, the system will only power on the master Routing Engine's SCB and will power off the other SCBs. Only the SCB in slot 0 will be powered on and a system log is generated. ◦ SCB3 supports up to 400 Gbps per slot with old midplanes and up to 500 Gbps per slot with new midplanes. ◦ SCB3 supports fabric intra-chassis redundancy. ◦ SCB3 supports the same chassis cluster function as the SRX5K-SCB (SCB1) and the SRX5K-SCBE (SCB2), except for in-service software upgrade (ISSU) and in-service hardware upgrade (ISHU). ◦ SCB3 has a second external Ethernet port. ◦ Fabric bandwidth increasing mode is not supported.

Copyright © 2016, Juniper Networks, Inc.

Table 278: show chassis hardware Output Fields *(continued)*

Field Name	Field Description
	• Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • IOCs Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced. ◦ IOC3 has two types of IOC3 MPCs, which have different built-in MICs: the 24x10GE + 6x40GE MPC and the 2x100GE + 4x10GE MPC. ◦ IOC3 supports SCB3 and SRX5000 line backplane and enhanced backplane. ◦ IOC3 can only work with SRX5000 line SCB2 and SCB3. If an SRX5000 line SCB is detected, IOC3 will be offline, an FPC misconfiguration alarm will be raised, and a system log message is generated. ◦ IOC3 interoperates with SCB2 and SCB3. ◦ IOC3 interoperates with the SRX5K-SPC-4-15-320 (SPC2) and the SRX5K-MPC (IOC2). ◦ The maximum power consumption for one IOC3 is 645W. An enhanced power module must be used. ◦ The IOC3 does not support the following command to set a PIC to go offline or online: **request chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot> <offline | online>** . ◦ IOC3 supports 240 Gbps of throughput with the enhanced SRX5000 line backplane. ◦ Chassis cluster functions the same as for the SRX5000 line IOC2. ◦ IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. ◦ IOC3 supports ISSU and ISHU in chassis cluster mode. ◦ IOC3 supports intra-FPC and and Inter-FPC Express Path (previously known as *services offloading*) with IPv4. ◦ NAT of IPv4 and IPv6 in normal mode and IPv4 for Express Path mode. ◦ All four PICs on the 24x10GE + 6x40GE cannot be powered on. A maximum of two PICs can be powered on at the same time. Use the **set chassis fpc <slot> pic <pic> power off** command to choose the PICs you want to power on. NOTE: Fabric bandwidth increasing mode is not supported on IOC3. • SRX Clustering Module (SCM) • Fan tray • For hosts, the Routing Engine type. ◦ Starting with Junos OS Release 12.1X47-D15, the SRX5K-RE-1800X4 (RE2) Routing Engine is introduced. ◦ The RE2 has an Intel Quad core Xeon processor, 16 GB of DRAM, and a 128-GB solid-state drive (SSD). The number 1800 refers to the speed of the processor (1.8 GHz). The maximum required power for this Routing Engine is 90W. NOTE: The RE2 provides significantly better performance than the previously used Routing Engine, even with a single core.

Copyright © 2016, Juniper Networks, Inc.

show chassis hardware

show chassis hardware

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                CM0715AK0021      SRX1500
Midplane         REV 08   750-058562   ACMA4255          SRX1500
CB 0             REV 08   711-053838   ACMA7529          CPU Board SRX700E
Routing Engine 0          BUILTIN      BUILTIN           SRX Routing Engine
FPC 0            REV 07   711-053832   ACMA3311          FEB
  PIC 0                   BUILTIN      BUILTIN           12x1G-T-4x1G-SFP-4x10G
    Xcvr 12      REV 01   740-014132   61521013          SFP-T
    Xcvr 13      REV 02   740-013111   A281604           SFP-T
    Xcvr 14      REV 02   740-011613   NRN30NV           SFP-SX
    Xcvr 15      REV 02   740-011613   NRN2PWV           SFP-SX
    Xcvr 16      REV 01   740-021308   AJA17B5           SFP+-10G-SR
    Xcvr 17      REV 01   740-021308   MSP056B           SFP+-10G-SR
    Xcvr 18      REV 01   740-031980   AS920WJ           SFP+-10G-SR
    Xcvr 19      REV 01   740-031980   AS92W5N           SFP+-10G-SR
Power Supply 0   REV 01   740-055217   1EDP42500JZ       PS 400W 90-264V AC in
Fan Tray 0                                              SRX1500 0, Front to Back
 Airflow - AFO
Fan Tray 1                                              SRX1500 1, Front to Back
 Airflow - AFO
Fan Tray 2                                              SRX1500 2, Front to Back
 Airflow - AFO
Fan Tray 3                                              SRX1500 3, Front to Back
 Airflow - AFO
```

show chassis hardware (SRX5600 and SRX5800 devices for SRX5K-MPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN12170EAAGA      SRX 5800
Midplane         REV 01   710-041799   ACAX3849          SRX 5800 Backplane
FPM Board        REV 01   710-024632   CAAX7297          Front Panel Display
PDM              Rev 03   740-013110   QCS170250DU       Power Distribution Modu
le
PEM 0            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC i
n
PEM 1            Rev 03   740-034724   QCS17020203C      PS 4.1kW; 200-240V AC i
n
PEM 2            Rev 04   740-034724   QCS17100200A      PS 4.1kW; 200-240V AC i
n
PEM 3            Rev 03   740-034724   QCS17080200M      PS 4.1kW; 200-240V AC i
n
Routing Engine 0 REV 11   740-023530   9012047437        SRX5k RE-13-20
CB 0             REV 09   710-024802   CAAX7202          SRX5k SCB
CB 1             REV 09   710-024802   CAAX7157          SRX5k SCB
FPC 0            REV 07   750-044175   CAAD0791          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 07   750-044175   CAAD0751          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
```

Copyright © 2016, Juniper Networks, Inc.

```
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 2           REV 28  750-020751   CAAW1817         SRX5k DPC 4X 10GE
  CPU           REV 04  710-024633   CAAZ5269         SRX5k DPC PMB
  PIC 0                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
    Xcvr 0      REV 02  740-014289   T10A00404        XFP-10G-SR
  PIC 1                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 2                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 3                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
FPC 6           REV 02  750-044175   ZY2552           SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
FPC 9           REV 10  750-044175   CAAP5932         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 10          REV 22  750-043157   ZH8192           SRX5k IOC II   CPU
     REV 08   711-043360   YX3879           SRX5k MPC PMB
  MIC 0         REV 01  750-049488   YZ2084           10x 10GE SFP+
    PIC 0               BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0    REV 01  740-031980   AMB0HG3          SFP+-10G-SR
      Xcvr 1    REV 01  740-031980   AM20B6F          SFP+-10G-SR
  MIC 1         REV 19  750-049486   CAAH3504         1x 100GE CFP
    PIC 2               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   X000D375         CFP-100G-SR10
FPC 11          REV 07.04.07 750-043157 CAAJ8771      SRX5k IOC II   CPU
     REV 08   711-043360   CAAJ3881           SRX5k MPC PMB
  MIC 0         REV 19  750-049486   CAAH0979         1x 100GE CFP
    PIC 0               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   UP1020Z          CFP-100G-SR10
  MIC 1         REV 08  750-049487   CAAM1160         2x 40GE QSFP+
    PIC 2               BUILTIN      BUILTIN          2x 40GE QSFP+
      Xcvr 0    REV 01  740-032986   QB151094         QSFP+-40G-SR4
      Xcvr 1    REV 01  740-032986   QB160509         QSFP+-40G-SR4
Fan Tray 0      REV 04  740-035409   ACAE0875         Enhanced Fan Tray
Fan Tray 1      REV 04  740-035409   ACAE0876         Enhanced Fan Tray
```

show chassis hardware (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN108DA5AAGA      SRX 5800
Midplane         REV 02   710-013698   TR0037            SRX 5600 Midplane
FPM Board        REV 02   710-014974   JY4635            Front Panel Display
PDM              Rev 02   740-013110   QCS10465005       Power Distribution Module
PEM 0            Rev 03   740-023514   QCS11154040       PS 1.7kW; 200-240VAC in
PEM 2            Rev 02   740-023514   QCS10504014       PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 05   740-015113   1000681023        RE-S-1300
CB 0             REV 05   710-013385   JY4775            SRX5k SCB
FPC 1            REV 17   750-020751   WZ6349            SRX5k DPC 4X 10GE
  CPU            REV 02   710-024633   WZ0718            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0                NON-JNPR     C724XM088         XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 02   740-011571   C831XJ08S         XFP-10G-SR
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
FPC 3            REV 22   750-043157   ZH8189            SRX5k IOC II
```


Copyright © 2016, Juniper Networks, Inc.

```
  CPU            REV 06   711-043360   YX3912           SRX5k MPC PMB
  MIC 0          REV 01   750-055732   CACF9115         20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 2     REV 02   740-013111   B358549          SFP-T
      Xcvr 9     REV 02   740-011613   PNB1FQS          SFP-SX
    PIC 1                 BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 9     REV 02   740-011613   PNB1FFF          SFP-SX
FPC 5            REV 01   750-027945   JW9665           SRX5k FIOC
  CPU
FPC 8            REV 08   750-023996   XA7234           SRX5k SPC
  CPU            REV 02   710-024633   XA1599           SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN          SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
Fan Tray 0       REV 03   740-014971   TP0902           Fan Tray
Fan Tray 1       REV 01   740-014971   TP0121           Fan Tray
```

show chassis hardware (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4 (RE2))

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                               JN122A040AGA      SRX5800
Midplane         REV 01   710-041799   ACRA7817          SRX5800 Backplane
FPM Board        REV 01   760-058099   CACA2100          Front Panel Display
PDM              Rev 03   740-013110   QCS1739517Z       Power Distribution Modu

                                       le
PEM 0            Rev 05   740-034724   QCS17460203K      PS 4.1kW; 200-240V AC i

                                       n
PEM 1            Rev 04   740-034724   QCS172302017      PS 4.1kW; 200-240V AC i

                                       n
Routing Engine 0 REV 01   740-056658   9013040855        SRX5k RE-1800X4
Routing Engine 1
CB 0             REV 01   750-056587   CACG1424          SRX5k SCB II
CB 1             REV 01   750-056587   CACC9307          SRX5k SCB II
CB 2             REV 01   750-056587   CAAZ1128          SRX5k SCB II
FPC 0            REV 10   750-056758   CACS2667          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4092          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACV0038          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 10   750-043157   CACB6877          SRX5k IOC II
  CPU            REV 04   711-043360   CACH6074          SRX5k MPC PMB
```

Copyright © 2016, Juniper Networks, Inc.

```
  MIC 0          REV 19   750-049486   CAAH3504          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020Z           CFP-100G-SR10
  MIC 1          REV 04   750-049488   CACB6429          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AP21RJ5           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   AP21RLJ           SFP+-10G-SR
      Xcvr 2     REV 01   740-030658   AD1148A0AYC       SFP+-10G-USR
      Xcvr 3     REV 01   740-031980   B11E02718         SFP+-10G-SR
FPC 4            REV 10   750-056758   CACW0706          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 7            REV 10   750-056758   CACS2725          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 8            REV 11   750-043157   CABN4955          SRX5k IOC II
  CPU            REV 04   711-043360   CACT9926          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0979          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP2077V           CFP-100G-SR10
FPC 9            REV 10   750-056758   CACW0755          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 10           REV 07   750-044175   CAAD0747          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
Fan Tray 0       REV 04   740-035409   ACAE2294          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE2099          Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1235BC7AGA      SRX5800
Midplane         REV 01   710-024803   ACRC3244          SRX5800 Backplane
FPM Board        REV 01   710-024632   CACA2108          Front Panel Display
PDM              Rev 03   740-013110   QCS1739519B       Power Distribution Module
PEM 0            Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01   740-056658   9009153221        SRX5k RE-1800X4
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
FPC 0            REV 18   750-054877   CACH4004          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4082          SRX5k SPC II
```


Copyright © 2016, Juniper Networks, Inc.

```
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 2             REV 10    750-056758   CACW0713         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 3             REV 11    750-043157   CACA8792         SRX5k IOC II
  CPU               REV 04    711-043360   CACA8809         SRX5k MPC PMB
  MIC 0             REV 19    750-049486   CAAH3485         1x 100GE CFP
    PIC 0                     BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0        REV 01    740-035329   UNM0G3C          CFP-100G-SR10
  MIC 1             REV 04    750-049488   CABX0782         10x 10GE SFP+
    PIC 2                     BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0        REV 01    740-031980   AMB0HX3          SFP+-10G-SR
      Xcvr 1        REV 01    740-031980   ANT0E6V          SFP+-10G-SR
      Xcvr 2        REV 01    740-031980   ANR0ZVY          SFP+-10G-SR
      Xcvr 3        REV 01    740-031980   AP308ZU          SFP+-10G-SR
FPC 4             REV 10    750-044175   CAAS8024         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 7             REV 10    750-056758   CACS5126         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 8             REV 11    750-043157   CACA8798         SRX5k IOC II
  CPU               REV 04    711-043360   CACA8826         SRX5k MPC PMB
  MIC 0             REV 19    750-049486   CAAH0996         1x 100GE CFP
    PIC 0                     BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0        REV 01    740-035329   UP30A6N          CFP-100G-SR10
FPC 9             REV 07    750-044175   CAAD0745         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                     BUILTIN      BUILTIN          SPU Flow
  PIC 1                     BUILTIN      BUILTIN          SPU Flow
  PIC 2                     BUILTIN      BUILTIN          SPU Flow
  PIC 3                     BUILTIN      BUILTIN          SPU Flow
FPC 10            REV 18    750-054877   CACD2570         SRX5k SPC II
  CPU                       BUILTIN      BUILTIN          SRX5k DPC PPC
Fan Tray 0        REV 04    740-035409   ACAE2122         Enhanced Fan Tray
Fan Tray 1        REV 04    740-035409   ACAE2254         Enhanced Fan Tray
```

show chassis hardware (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item            Version  Part number  Serial number     Description
Chassis                               JN1250870AGB      SRX5600
Midplane        REV 01   760-063936   ACRE2578          Enhanced SRX5600 Midplane
```

```
FPM Board        REV 02   710-017254   KD9027            Front Panel Display
PEM 0            Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V A

                                       C in
PEM 1            Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V A

                                       C in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0            REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU            REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4            REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5            REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                 Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board        test     710-017254   test                Front Panel Display
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry
Module
PEM 1            Rev 01   740-038514   QCS120311 00J     DC 2.6kW Power Entry
Module
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 11   750-043157   CACY1592          SRX5k IOC II
```


Copyright © 2016, Juniper Networks, Inc.

```
  CPU            REV 04   711-043360   CACZ8831          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACN0239          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-031980   ARN23HW           SFP+-10G-SR
      Xcvr 8     REV 01   740-031980   ARN2FVW           SFP+-10G-SR
      Xcvr 9     REV 01   740-031980   ARN2YVM           SFP+-10G-SR
FPC 5            REV 10   750-056758   CADA8736          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                Enhanced Fan Tray
```

show chassis hardware detail

show chassis hardware detail (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and (SRX5K-RE-1800X4 (RE2)

```
user@host> show chassis hardware detail
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN122A040AGA      SRX5800
Midplane         REV 01   710-041799   ACRA7817          SRX5800 Backplane
FPM Board        REV 01   760-058099   CACA2100          Front Panel Display
PDM             Rev 03   740-013110   QCS1739517Z       Power Distribution Module
PEM 0            Rev 05   740-034724   QCS17460203K      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 04   740-034724   QCS172302017      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01   740-056658   9013040855        SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200269450529 741 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1
CB 0             REV 01   750-056587   CACG1424          SRX5k SCB II
CB 1             REV 01   750-056587   CACC9307          SRX5k SCB II
CB 2             REV 01   750-056587   CAAZ1128          SRX5k SCB II
FPC 0            REV 10   750-056758   CACS2667          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4092          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACV0038          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
```

Copyright © 2016, Juniper Networks, Inc.

```
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 3           REV 10  750-043157   CACB6877         SRX5k IOC II
  CPU           REV 04  711-043360   CACH6074         SRX5k MPC PMB
  MIC 0         REV 19  750-049486   CAAH3504         1x 100GE CFP
    PIC 0               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   UP1020Z          CFP-100G-SR10
  MIC 1         REV 04  750-049488   CACB6429         10x 10GE SFP+
    PIC 2               BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0    REV 01  740-031980   AP21RJ5          SFP+-10G-SR
      Xcvr 1    REV 01  740-031980   AP21RLJ          SFP+-10G-SR
      Xcvr 2    REV 01  740-030658   AD1148A0AYC      SFP+-10G-USR
      Xcvr 3    REV 01  740-031980   B11E02718        SFP+-10G-SR
FPC 4           REV 10  750-056758   CACW0706         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 7           REV 10  750-056758   CACS2725         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 8           REV 11  750-043157   CABN4955         SRX5k IOC II
  CPU           REV 04  711-043360   CACT9926         SRX5k MPC PMB
  MIC 0         REV 19  750-049486   CAAH0979         1x 100GE CFP
    PIC 0               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01  740-035329   UP2077V          CFP-100G-SR10
FPC 9           REV 10  750-056758   CACW0755         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 10          REV 07  750-044175   CAAD0747         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
Fan Tray 0      REV 04  740-035409   ACAE2294         Enhanced Fan Tray
Fan Tray 1      REV 04  740-035409   ACAE2099         Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item            Version  Part number  Serial number     Description
Chassis                               JN1235BC7AGA      SRX5800
Midplane        REV 01   710-024803   ACRC3244          SRX5800 Backplane
FPM Board       REV 01   710-024632   CACA2108          Front Panel Display
PDM             Rev 03   740-013110   QCS1739519B       Power Distribution Module
PEM 0           Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
PEM 1           Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 01  740-056658   9009153221        SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200298450703 72 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
```


Copyright © 2016, Juniper Networks, Inc.

```
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
FPC 0            REV 18   750-054877   CACH4004          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 18   750-054877   CACH4082          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 10   750-056758   CACW0713          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 11   750-043157   CACA8792          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8809          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH3485          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UNM0G3C           CFP-100G-SR10
  MIC 1          REV 04   750-049488   CABX0782          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-031980   AMB0HX3           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   ANT0E6V           SFP+-10G-SR
      Xcvr 2     REV 01   740-031980   ANR0ZVY           SFP+-10G-SR
      Xcvr 3     REV 01   740-031980   AP308ZU           SFP+-10G-SR
FPC 4            REV 10   750-044175   CAAS8024          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 7            REV 10   750-056758   CACS5126          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 8            REV 11   750-043157   CACA8798          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8826          SRX5k MPC PMB
  MIC 0          REV 19   750-049486   CAAH0996          1x 100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1x 100GE CFP
      Xcvr 0     REV 01   740-035329   UP30A6N           CFP-100G-SR10
FPC 9            REV 07   750-044175   CAAD0745          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 10           REV 18   750-054877   CACD2570          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
Fan Tray 0       REV 04   740-035409   ACAE2122          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE2254          Enhanced Fan Tray
```

Copyright © 2016, Juniper Networks, Inc.

show chassis hardware detail (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 SCB3 with enhanced midplanes and (SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware detail
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1250870AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2578          Enhanced SRX5600 Midplane
FPM Board         REV 02   710-017254   KD9027              Front Panel Display
PEM 0             Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V
AC in
PEM 1             Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V
AC in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
  ad0    3831 MB  UGB30SFA4000T2       SFA4000T2 000027A0 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-043         Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0          SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1          SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2          SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3          SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0              REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0             REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                      BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                    BUILTIN      BUILTIN           SPU Cp
  PIC 1                    BUILTIN      BUILTIN           SPU Flow
  PIC 2                    BUILTIN      BUILTIN           SPU Flow
  PIC 3                    BUILTIN      BUILTIN           SPU Flow
FPC 2             REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU             REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                    BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                    BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                    BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0        REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2        REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                    BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ        REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4             REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU             REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1           REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                  BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7      REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8      REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9      REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5             REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                      BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                    BUILTIN      BUILTIN           SPU Flow
  PIC 1                    BUILTIN      BUILTIN           SPU Flow
  PIC 2                    BUILTIN      BUILTIN           SPU Flow
  PIC 3                    BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                 Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
```


Copyright © 2016, Juniper Networks, Inc.

```
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board         test     710-017254   test                Front Panel Display
PEM 0             Rev 01   740-038514   QCS114111003        DC 2.6kW Power Entry
Module
PEM 1             Rev 01   740-038514   QCS120311O0J        DC 2.6kW Power Entry
Module
Routing Engine 0 REV 01   740-056658   9009186342          SRX5k RE-1800X4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200313161216 109 Compact Flash
  ad1   28843 MB  UGB94BPH32H0S2-KCI   11000160387         Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000       uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087       uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0             REV 01   750-062257   CAEB8178            SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN             SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN             SPU Cp
  PIC 1                   BUILTIN      BUILTIN             SPU Flow
  PIC 2                   BUILTIN      BUILTIN             SPU Flow
  PIC 3                   BUILTIN      BUILTIN             SPU Flow
FPC 4            REV 11   750-043157   CACY1592            SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8831            SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACN0239            10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN             10x 10GE SFP+
      Xcvr 7     REV 01   740-031980   ARN23HW             SFP+-10G-SR
      Xcvr 8     REV 01   740-031980   ARN2FVW             SFP+-10G-SR
      Xcvr 9     REV 01   740-031980   ARN2YVM             SFP+-10G-SR
FPC 5            REV 10   750-056758   CADA8736            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN             SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN             SPU Flow
  PIC 1                   BUILTIN      BUILTIN             SPU Flow
  PIC 2                   BUILTIN      BUILTIN             SPU Flow
  PIC 3                   BUILTIN      BUILTIN             SPU Flow
Fan Tray                                                  Enhanced Fan Tray
```

show chassis hardware extensive node 1

show chassis hardware extensive node 1 (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4)

```
user@host> show chassis hardware extensive node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1235BC7AGA      SRX5800
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               JN1235BC7AGA
 Assembly ID:  0x051a            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SRX5800
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 1a 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 33 35 42 43 37 41 47 41 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   710-024803   ACRC3244          SRX5800 Backplane
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          710-024803        S/N:               S/N ACRC3244
 Assembly ID:  0x091a            Assembly Version:  01.01
 Date:         02-26-2014        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5800 Backplane           FRU Model Number:  SRX5800-BP-A
 Board Information Record:
  Address 0x00: ad 01 08 00 4c 96 14 d3 28 00 00 ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 38 30 33 00 00
  Address 0x20: 53 2f 4e 20 41 43 52 43 33 32 34 34 00 1a 02 07
  Address 0x30: de ff ff ff ad 01 08 00 4c 96 14 d3 28 00 00 ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 42 50 2d 41 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 01   710-024632   CACA2108          Front Panel Display
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          710-024632        S/N:               S/N CACA2108
 Assembly ID:  0x096f            Assembly Version:  01.01
 Date:         02-05-2014        Assembly Flags:    0x00
 Version:      REV 01
 ID: Front Panel Display         FRU Model Number:  SRX5800-CRAFT-A
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 6f 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 36 33 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 41 32 31 30 38 00 05 02 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 43 52 41 46 54 2d 41 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PDM              Rev 03   740-013110   QCS1739519B       Power Distribution Module

 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-013110        S/N:               QCS1739519B
 Assembly ID:  0x0416            Assembly Version:  01.03
 Date:         10-26-2013        Assembly Flags:    0x00
 Version:      Rev 03
 ID: Power Distribution Module
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 16 01 03 52 65 76 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 31 33 31 31 30 00 00
  Address 0x20: 51 43 53 31 37 33 39 35 31 39 42 00 00 1a 0a 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 0            Rev 04   740-034724   QCS17230201Z      PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
```


Copyright © 2016, Juniper Networks, Inc.

```
 P/N:          740-034724        S/N:               QCS17230201Z
 Assembly ID:  0x044b            Assembly Version:  01.04
 Date:         06-04-2013        Assembly Flags:    0x00
 Version:      Rev 04
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  SRX5800-PWR-4100-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4b 01 04 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 34 37 32 34 00 00
  Address 0x20: 51 43 53 31 37 32 33 30 32 30 31 5a 00 04 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 50 57 52 2d 34 31 30 30 2d
  Address 0x60: 41 43 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 05   740-034724   QCS174502014      PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-034724        S/N:               QCS174502014
 Assembly ID:  0x044b            Assembly Version:  01.05
 Date:         11-06-2013        Assembly Flags:    0x00
 Version:      Rev 05
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  SRX5800-PWR-4100-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4b 01 05 52 65 76 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 34 37 32 34 00 00
  Address 0x20: 51 43 53 31 37 34 35 30 32 30 31 34 00 06 0b 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 38 30 30 2d 50 57 52 2d 34 31 30 30 2d
  Address 0x60: 41 43 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 01   740-056658   9009153221        SRX5k RE-1800X4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-056658        S/N:               9009153221
 Assembly ID:  0x0c1a            Assembly Version:  01.01
 Date:         07-22-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k RE-1800X4             FRU Model Number:  SRX5K-RE-1800X4
 Board Information Record:
  Address 0x00: 54 32 30 32 37 45 43 2d 34 34 47 42 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 36 36 35 38 00 00
  Address 0x20: 39 30 30 39 31 35 33 32 32 31 00 00 00 16 07 07
  Address 0x30: dd ff ff ff 54 32 30 32 37 45 43 2d 34 34 47 42
  Address 0x40: 00 00 00 00 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 52 45 2d 31 38 30 30 58 34 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 64 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200298450703 72 Compact Flash
  ad1   114304 MB  VSFA18PI128G-KC     32779-073         Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
```

Copyright © 2016, Juniper Networks, Inc.

```
Routing Engine 1
CB 0             REV 01   750-056587   CACC9541          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACC9541
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-07-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 43 39 35 34 31 00 07 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 01   750-056587   CACG1447          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACG1447
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-07-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 47 31 34 34 37 00 07 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
CB 2             REV 01   750-056587   CACH9058          SRX5k SCB II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-056587        S/N:               S/N CACH9058
 Assembly ID:  0x0c19            Assembly Version:  01.01
 Date:         03-06-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: SRX5k SCB II                FRU Model Number:  SRX5K-SCBE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 19 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 38 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 48 39 30 35 38 00 06 03 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 08 ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware extensive node 1 (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware extensive node 1
```


Copyright © 2016, Juniper Networks, Inc.

```
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               JN124FEC0AGB
 Assembly ID:  0x051b            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x08
 ID: SRX5600
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 1b 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 34 46 45 43 30 41 47 42 08 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midpla

                                    ne
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          760-063936        S/N:               ACRE2946
 Assembly ID:  0x0914            Assembly Version:  01.01
 Date:         03-19-2015        Assembly Flags:    0x08
 Version:      REV 01            CLEI Code:         CLEI-CODE
 ID: SRX5600 Midplane            FRU Model Number:  SRX5600X-CHAS
 Board Information Record:
  Address 0x00: ad 01 08 00 88 a2 5e 12 68 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 14 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 36 33 39 33 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 52 45 32 39 34 36 08 13 03 07
  Address 0x30: df ff ff ff ad 01 08 00 88 a2 5e 12 68 00 ff ff
  Address 0x40: ff ff ff ff 01 43 4c 45 49 2d 43 4f 44 45 20 53
  Address 0x50: 52 58 35 36 30 30 58 2d 43 48 41 53 20 20 20 20
  Address 0x60: 20 20 20 20 20 20 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ba ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        test     710-017254   test              Front Panel Display
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          710-017254        S/N:               test
 Assembly ID:  0x01ff            Assembly Version:  01.00
 Date:         06-18-2007        Assembly Flags:    0x00
 Version:      test
 ID: Front Panel Display
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 01 ff 01 00 74 65 73 74 00 00 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 32 35 34 00 00
  Address 0x20: 74 65 73 74 00 00 00 00 00 00 00 00 00 12 06 07
  Address 0x30: d7 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 ff ff ff ff ff ff ff ff ff ff ff ff
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry
Module
```

```
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-038514        S/N:               QCS114111003
 Assembly ID:  0x044c            Assembly Version:  01.01
 Date:         10-14-2011        Assembly Flags:    0x00
 Version:      Rev 01
 ID: DC 2.6kW Power Entry Module FRU Model Number:  SRX5600-PWR-2400-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4c 01 01 52 65 76 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 38 35 31 34 00 00
  Address 0x20: 51 43 53 31 31 34 31 31 31 30 30 33 00 0e 0a 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 36 30 30 2d 50 57 52 2d 32 34 30 30 2d
  Address 0x60: 44 43 2d 53 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 01   740-038514   QCS120311100J      DC 2.6kW Power Entry
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-038514        S/N:               QCS120311100J
 Assembly ID:  0x044c            Assembly Version:  01.01
 Date:         01-17-2012        Assembly Flags:    0x00
 Version:      Rev 01
 ID: DC 2.6kW Power Entry Module FRU Model Number:  SRX5600-PWR-2400-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 4c 01 01 52 65 76 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 38 35 31 34 00 00
  Address 0x20: 51 43 53 31 32 30 33 31 31 30 30 4a 00 11 01 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 52 58 35 36 30 30 2d 50 57 52 2d 32 34 30 30 2d
  Address 0x60: 44 43 2d 53 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-056658        S/N:               9009186342
 Assembly ID:  0x0c1a            Assembly Version:  01.01
 Date:         02-05-2014        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         COUCATTBAA
 ID: SRX5k RE-1800X4             FRU Model Number:  SRX5K-RE-1800X4
 Board Information Record:
  Address 0x00: 54 32 30 32 37 45 43 2d 34 34 47 42 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 1a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 36 36 35 38 00 00
  Address 0x20: 39 30 30 39 31 38 36 33 34 32 00 00 00 05 02 07
  Address 0x30: de ff ff ff 54 32 30 32 37 45 43 2d 34 34 47 42
  Address 0x40: 00 00 00 00 01 43 4f 55 43 41 54 54 42 41 41 53
  Address 0x50: 52 58 35 4b 2d 52 45 2d 31 38 30 30 58 34 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 64 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200313161216 109 Compact Flash
  ad1   28843 MB  UGB94BPH32H0S2-KCI   11000160387       Disk 1
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
```


Copyright © 2016, Juniper Networks, Inc.

```
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-062257        S/N:               CAEB8178
 Assembly ID:  0x0c59            Assembly Version:  01.01
 Date:         03-19-2015        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         CLEI-CODE
 ID: SRX5k SCB3                  FRU Model Number:  SRX5K-SCB3
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 59 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 32 35 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 42 38 31 37 38 00 13 03 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4c 45 49 2d 43 4f 44 45 20 53
  Address 0x50: 52 58 35 4b 2d 53 43 42 33 20 20 20 20 20 20 20
  Address 0x60: 20 20 20 20 20 20 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 63 ff ff ff ff ff ff ff ff ff ff ff ff
FPC 2            REV 01   750-062243   CAED0386          SRX5k IOC3 24XGE+6XLG
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-062243        S/N:               CAED0386
 Assembly ID:  0x0c57            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5k IOC3 24XGE+6XLG
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 fe 0c 57 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 32 34 33 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 44 30 33 38 36 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
  Address 0x40: f2 06 00 ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  CPU             REV 01   711-062244   CADX8554          SRX5k IOC3 PMB
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-062244        S/N:               CADX8554
 Assembly ID:  0x0c5a            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: SRX5k IOC3 PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0c 5a 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 36 32 32 34 34 00 00
  Address 0x20: 53 2f 4e 20 43 41 44 58 38 35 35 34 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 2a 47 7e 50 10 05 76 5c
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0ab5            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: 12x 10GE SFP+
```

Copyright © 2016, Juniper Networks, Inc.

```
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b5 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 a9 3c 00 00 00 00 0a b6 00 00
 PIC 1                  BUILTIN      BUILTIN           12x 10GE SFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab5            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 12x 10GE SFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b5 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 e9 b4 00 00 00 00 0a b5 00 00
 PIC 2                  BUILTIN      BUILTIN           3x 40GE QSFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab6            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 3x 40GE QSFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b6 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 03 e8 c4 33 24 3a 38 00 00 00 02
 PIC 3                  BUILTIN      BUILTIN           3x 40GE QSFP+
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0ab6            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 3x 40GE QSFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a b6 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```


Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x70: 00 00 00 00 c0 02 ab 1c 00 00 00 00 0a b5 00 00
  WAN MEZZ       REV 01   750-062682   CAEA4788          24x 10GE SFP+ Mezz
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-062682        S/N:               CAEA4788
 Assembly ID:  0x0c76            Assembly Version:  01.01
 Date:         04-28-2015        Assembly Flags:    0x00
 Version:      REV 01
 ID: 24x 10GE SFP+ Mezz
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0c 76 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 32 36 38 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 45 41 34 37 38 38 00 1c 04 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
FPC 4            REV 11   750-043157   CACY1592          SRX5k IOC II
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-043157        S/N:               CACY1592
 Assembly ID:  0x0bd1            Assembly Version:  04.11
 Date:         07-30-2014        Assembly Flags:    0x00
 Version:      REV 11            CLEI Code:         COUIBCWBAA
 ID: SRX5k IOC II                FRU Model Number:  SRX5K-MPC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b d1 04 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 31 35 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 59 31 35 39 32 00 1e 07 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
  Address 0x40: f2 06 00 ff 01 43 4f 55 49 42 43 57 42 41 41 53
  Address 0x50: 52 58 35 4b 2d 4d 50 43 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 92 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 04   711-043360   CACZ8831          SRX5k MPC PMB
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-043360        S/N:               CACZ8831
 Assembly ID:  0x0bd2            Assembly Version:  01.04
 Date:         07-28-2014        Assembly Flags:    0x00
 Version:      REV 04
 ID: SRX5k MPC PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b d2 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 33 33 36 30 00 00
  Address 0x20: 53 2f 4e 20 43 41 43 5a 38 38 33 31 00 1c 07 07
  Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 49 fa 60 10 40 05 76 5c
  MIC 1          REV 04   750-049488   CACN0239          10x 10GE SFP+
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-049488        S/N:               CACN0239
 Assembly ID:  0x0a88            Assembly Version:  02.04
 Date:         02-26-2014        Assembly Flags:    0x00
 Version:      REV 04            CLEI Code:         COUIBCXBAA
```

```
  ID: 10x 10GE SFP+               FRU Model Number:  SRX-MIC-10XG-SFPP
  Board Information Record:
   Address 0x00: 34 01 03 03 ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0a 88 02 04 52 45 56 20 30 34 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 34 39 34 38 38 00 00
   Address 0x20: 53 2f 4e 20 43 41 43 4e 30 32 33 39 00 1a 02 07
   Address 0x30: de ff ff ff 34 01 03 03 ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4f 55 49 42 43 58 42 41 41 53
   Address 0x50: 52 58 2d 4d 49 43 2d 31 30 58 47 2d 53 46 50 50
   Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 9f c0 03 e4 14 55 8a 95 a8 00 00 00 02
     PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
       Xcvr 7     REV 01   740-031980   ARN23HW           SFP+-10G-SR
       Xcvr 8     REV 01   740-031980   ARN2FVW           SFP+-10G-SR
       Xcvr 9     REV 01   740-031980   ARN2YVM           SFP+-10G-SR
FPC 5              REV 10   750-056758   CADA8736          SRX5k SPC II
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-056758        S/N:               CADA8736
  Assembly ID:  0x0b4f            Assembly Version:  01.10
  Date:         09-01-2014        Assembly Flags:    0x00
  Version:      REV 10            CLEI Code:         COUCATLBAB
  ID: SRX5k SPC II                FRU Model Number:  SRX5K-SPC-4-15-320
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ae 01 f2 06 00 ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 4f 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 37 35 38 00 00
   Address 0x20: 53 2f 4e 20 43 41 44 41 38 37 33 36 00 01 09 07
   Address 0x30: de ff ff ff ff ff ff ff ff ff ff ff ff ff ae 01
   Address 0x40: f2 06 00 ff 01 43 4f 55 43 41 54 4c 42 41 42 53
   Address 0x50: 52 58 35 4b 2d 53 50 43 2d 34 2d 31 35 2d 33 32
   Address 0x60: 30 00 00 00 00 00 43 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 50 ff ff ff ff ff ff ff ff ff ff ff ff
   CPU                   BUILTIN      BUILTIN           SRX5k DPC PPC
   PIC 0                 BUILTIN      BUILTIN           SPU Cp
  Jedec Code:   0x0000            EEPROM Version:    0x00
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0x0a20            Assembly Version:  00.00
  Date:         00-00-0000        Assembly Flags:    0x00
  ID: SPU Cp
  Board Information Record:
   Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  I2C Hex Data:
   Address 0x00: 00 00 00 00 0a 20 00 00 00 00 00 00 00 00 00 00
   Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
   Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
   Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x70: 00 00 00 00 de ad be ef 46 10 f0 d8 40 43 43 c0
   PIC 1                 BUILTIN      BUILTIN           SPU Flow
  Jedec Code:   0x0000            EEPROM Version:    0x00
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0x0a21            Assembly Version:  00.00
  Date:         00-00-0000        Assembly Flags:    0x00
  ID: SPU Flow
  Board Information Record:
   Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  I2C Hex Data:
```


Copyright © 2016, Juniper Networks, Inc.

```
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 13 5b d0 40 43 43 c0
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0a21            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SPU Flow
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 0c 66 40 40 43 43 c0
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0a21            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: SPU Flow
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 21 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 41 73 73 65
  Address 0x20: 42 55 49 4c 54 49 4e 00 41 73 73 65 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 46 0e db 00 40 43 43 c0
Fan Tray                                                Enhanced Fan Tray
                                FRU Model Number:  SRX5600-HC-FAN
```

show chassis hardware clei-models

show chassis hardware clei-models (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE (SCB2) and SRX5K-RE-1800X4 (RE2)

```
user@host> show chassis hardware clei-models node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   710-024803                     SRX5800-BP-A
FPM Board        REV 01   710-024632                     SRX5800-CRAFT-A
PEM 0            Rev 04   740-034724                     SRX5800-PWR-4100-AC
PEM 1            Rev 05   740-034724                     SRX5800-PWR-4100-AC
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
```

```
CB 1             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
CB 2             REV 01   750-056587   COUCATSBAA        SRX5K-SCBE
FPC 0            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 1            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 2            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 3            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
FPC 4            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 7            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 8            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 0          REV 05   750-049486   COUIBCYBAA        SRX-MIC-1X100G-CFP
FPC 9            REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 10           REV 18   750-054877   COUCATLBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
Fan Tray 0       REV 04   740-035409                     SRX5800-HC-FAN
Fan Tray 1       REV 04   740-035409                     SRX5800-HC-FAN
```

show chassis hardware clei-models (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 (SCB3) with enhanced midplanes and SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```
user@host> show chassis hardware clei-models
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   760-063936   CLEI-CODE         SRX5600X-CHAS
FPM Board        REV 02   710-017254                     CRAFT-MX480-S
PEM 0            Rev 03   740-034701                     SRX5600-PWR-2520-AC-S
PEM 1            Rev 03   740-034701                     SRX5600-PWR-2520-AC-S
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-062257   CLEI-CODE         SRX5K-SCB3
FPC 0            REV 10   750-056758   COUCATLBAB        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 2            REV 01   750-062243
FPC 4            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
FPC 5            REV 05   750-044175   PROTOXCLEI        750-044175
  CPU                     BUILTIN
Fan Tray                                                 SRX5600-HC-FAN

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   760-063936   CLEI-CODE         SRX5600X-CHAS
PEM 0            Rev 01   740-038514                     SRX5600-PWR-2400-DC-S
PEM 1            Rev 01   740-038514                     SRX5600-PWR-2400-DC-S
Routing Engine 0 REV 01   740-056658   COUCATTBAA        SRX5K-RE-1800X4
CB 0             REV 01   750-062257   CLEI-CODE         SRX5K-SCB3
FPC 0            REV 07   750-044175   COUCASFBAA        SRX5K-SPC-4-15-320
  CPU                     BUILTIN
FPC 4            REV 11   750-043157   COUIBCWBAA        SRX5K-MPC
  MIC 1          REV 04   750-049488   COUIBCXBAA        SRX-MIC-10XG-SFPP
```


Copyright © 2016, Juniper Networks, Inc.

```
FPC 5            REV 10   750-056758   COUCATLBAB       SRX5K-SPC-4-15-320
  CPU                     BUILTIN
Fan Tray                                                 SRX5600-HC-FAN
```

Copyright © 2016, Juniper Networks, Inc.

show ethernet-switching mac-learning-log (View)

Syntax **show ethernet-switching mac-learning-log**

Release Information Command introduced in Junos OS Release 9.5.

Description Displays the event log of learned MAC addresses.

Required Privilege Level view

Related Documentation
- Understanding MAC Limiting on page 2999

Output Fields Table 279 lists the output fields for the show ethernet-switching mac-learning-log command. Output fields are listed in the approximate order in which they appear.

Table 279: show interfaces Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
VLAN-IDX	VLAN index. An internal value assigned by Junos OS for each VLAN.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log

```
user@host> show ethernet-switching mac-learning-log
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
```


Copyright © 2016, Juniper Networks, Inc.

```
vlan_idx 15 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:05:00:00:05 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:30:48:90:54:89 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5e:00:01:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5e:00:01:08 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5e:00:01:09 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:19:e2:50:ac:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5e:00:01:04 was learned
[output truncated]
```

Copyright © 2016, Juniper Networks, Inc.

show ethernet-switching table (View)

Syntax **show ethernet-switching table (brief |detail |extensive) interface** ***interface-name***

Release Information Command introduced in Junos OS Release 9.5.

Description Displays the Ethernet switching table.

Options
- **none**—(Optional) Display brief information about the Ethernet switching table.
- **brief** | **detail** | **extensive**—(Optional) Display the specified level of output.
- **interface-name**—(Optional) Display the Ethernet switching table for a specific interface.

Required Privilege Level view

Related Documentation
- Port Security Overview on page 2999
- Understanding MAC Limiting on page 2999

Output Fields Table 280 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 280: show ethernet-switching table Output Fields

Field Name	Field Description
VLAN	The name of a VLAN.
MAC address	The MAC address associated with the VLAN.
Type	The type of MAC address. Values are: • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.

Sample Output

show ethernet-switching table

```
user@host> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
```


Copyright © 2016, Juniper Networks, Inc.

```
F2 * Flood - All-members
F2 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
[output truncated]
```

Sample Output

show ethernet-switching table brief

```
user@host> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
F2 * Flood - All-members
F2 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
```

Copyright © 2016, Juniper Networks, Inc.

```
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
[output truncated]
```

Sample Output

show ethernet-switching table detail

```
user@host> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
```


Copyright © 2016, Juniper Networks, Inc.

```
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show ethernet-switching table interface ge-0/0/1

```
user@host> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN     MAC address         Type     Age Interfaces
V1       *                   Flood    - All-members
V1       00:00:05:00:00:05   Learn    0 ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

show igmp-snooping route (View)

Syntax show igmp-snooping route (brief | detail | ethernet-switching | inet | vlan)

Release Information Command introduced in Junos OS Release 9.5.

Description Display IGMP snooping route information.

Options

- **none**—Display general parameters.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **ethernet-switching**—(Optional) Display Ethernet switching information.
- **inet**—(Optional) Display inet information.
- **vlan *vlan-id* |*vlan-name***—(Optional) Display route information for the specified VLAN.

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

Output Fields Table 281 lists the output fields for the **show igmp-snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 281: show igmp-snooping route Output Fields

Field Name	Field Description
VLAN	Name of the VLAN.
Group	Multicast group address.
Next-hop	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@host> show igmp-snooping route
VLAN       Group          Next-hop
v11         224.1.1.1, *   533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12         224.1.1.3, *   534
Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@host> show igmp-snooping route vlan v1
Table: 0
VLAN      Group              Next-hop
v1        224.1.1.1, *       1266
Interfaces: ge-0/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
v1        224.1.1.3, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.5, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.7, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.9, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.11, *      1266
Interfaces: ge-0/0/0.0
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces (SRX Series)

Syntax

```
show interfaces {
  <brief | detail | extensive | terse>
  controller interface-name
  descriptions interface-name
  destination-class (all | destination-class-name logical-interface-name)
  diagnostics optics interface-name
  far-end-interval interface-fpc/pic/port
  filters interface-name
  flow-statistics interface-name
  interval interface-name
  load-balancing (detail | interface-name)
  mac-database mac-address mac-address
  mc-ae id identifier unit number revertive-info
  media interface-name
  policers interface-name
  queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics
  redundancy (detail | interface-name)
  routing brief detail summary interface-name
  routing-instance (all | instance-name)
  snmp-index snmp-index
  source-class (all | destination-class-name logical-interface-name)
  statistics interface-name
  switch-port switch-port number
  transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all |
    interface-name)
  zone interface-name
}
```

Release Information Command modified in Junos OS Release 9.5.

Description Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

On SRX Series appliance, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim*/0/*port***—E1 interface.
  - **e3-*pim*/0/*port***—E3 interface.

Copyright © 2016, Juniper Networks, Inc.

  - **fe-*pim*/0/*port***—Fast Ethernet interface.
  - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
  - **se-*pim*/0/*port***—Serial interface.
  - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
  - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
  - **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

- **brief | detail | extensive | terse**—(Optional) Display the specified level of output.
- **controller**—(Optional) Show controller information.
- **descriptions**—(Optional) Display interface description strings.
- **destination-class**—(Optional) Show statistics for destination class.
- **diagnostics**—(Optional) Show interface diagnostics information.
- **far-end-interval**—(Optional) Show far end interval statistics.
- **filters**—(Optional) Show interface filters information.
- **flow-statistics**—(Optional) Show security flow counters and errors.
- **interval**—(Optional) Show interval statistics.
- **load-balancing**—(Optional) Show load-balancing status.
- **mac-database**—(Optional) Show media access control database information.
- **mc-ae**—(Optional) Show MC-AE configured interface information.
- **media**—(Optional) Display media information.
- **policers**—(Optional) Show interface policers information.
- **queue**—(Optional) Show queue statistics for this interface.
- **redundancy**—(Optional) Show redundancy status.
- **routing**—(Optional) Show routing status.
- **routing-instance**—(Optional) Name of routing instance.
- **snmp-index**—(Optional) SNMP index of interface.
- **source-class**—(Optional) Show statistics for source class.
- **statistics**—(Optional) Display statistics and detailed output.
- **switch-port**—(Optional) Front end port number (0..15).
- **transport**—(Optional) Show interface transport information.
- **zone**—(Optional) Interface's zone.

Required Privilege Level view

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Interfaces on page 2441

List of Sample Output

show interfaces Gigabit Ethernet on page 2824
show interfaces brief (Gigabit Ethernet) on page 2825
show interfaces detail (Gigabit Ethernet) on page 2825
show interfaces extensive (Gigabit Ethernet) on page 2827
show interfaces terse on page 2830
show interfaces controller (Channelized E1 IQ with Logical E1) on page 2830
show interfaces controller (Channelized E1 IQ with Logical DS0) on page 2830
show interfaces descriptions on page 2831
show interfaces destination-class all on page 2831
show interfaces diagnostics optics on page 2831
show interfaces far-end-interval coc12-5/2/0 on page 2832
show interfaces far-end-interval coc1-5/2/1:1 on page 2832
show interfaces filters on page 2833
show interfaces flow-statistics (Gigabit Ethernet) on page 2833
show interfaces interval (Channelized OC12) on page 2834
show interfaces interval (E3) on page 2834
show interfaces interval (SONET/SDH) on page 2835
show interfaces load-balancing on page 2835
show interfaces load-balancing detail on page 2835
show interfaces mac-database (All MAC Addresses on a Port) on page 2836
show interfaces mac-database (All MAC Addresses on a Service) on page 2836
show interfaces mac-database mac-address on page 2837
show interfaces mc-ae on page 2837
show interfaces media (SONET/SDH) on page 2837
show interfaces policers on page 2838
show interfaces policers interface-name on page 2838
show interfaces queue on page 2838
show interfaces redundancy on page 2839
show interfaces redundancy (Aggregated Ethernet) on page 2839
show interfaces redundancy detail on page 2840
show interfaces routing brief on page 2840
show interfaces routing detail on page 2840
show interfaces routing-instance all on page 2841
show interfaces snmp-index on page 2841
show interfaces source-class all on page 2841
show interfaces statistics (Fast Ethernet) on page 2842
show interfaces switch-port on page 2842
show interfaces transport pm on page 2843
show security zones on page 2844

Output Fields

Table 194 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—There are two possible conditions regarding the number of oversized frames: ◦ Packet length exceeds 1518 octets, or ◦ Packet length exceeds MRU • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 282: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
```


Copyright © 2016, Juniper Networks, Inc.

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

Sample Output

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
```

Copyright © 2016, Juniper Networks, Inc.

```
Disabled,
 Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
 Device flags   : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags     : None
 CoS queues     : 8 supported, 8 maximum usable queues
 Hold-times     : Up 0 ms, Down 0 ms
 Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
 Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
 Statistics last cleared: Never
 Traffic statistics:
  Input  bytes  :                    0                    0 bps
  Output bytes  :                    0                    0 bps
  Input  packets:                    0                    0 pps
  Output packets:                    0                    0 pps
 Egress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort                   0                    0                    0

   1 expedited-fo                  0                    0                    0

   2 assured-forw                  0                    0                    0

   3 network-cont                  0                    0                    0

 Queue number:         Mapped forwarding classes
   0                   best-effort
   1                   expedited-forwarding
   2                   assured-forwarding
   3                   network-control
 Active alarms  : LINK
 Active defects : LINK
 Interface transmit statistics: Disabled

 Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
   Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Transit statistics:
    Input  bytes  :                   0                    0 bps
    Output bytes  :                   0                    0 bps
    Input  packets:                   0                    0 pps
    Output packets:                   0                    0 pps
   Security: Zone: public
   Flow Statistics :
   Flow Input statistics :
     Self packets :                    0
     ICMP packets :                    0
     VPN packets :                     0
     Multicast packets :               0
     Bytes permitted by policy :       0
     Connections established :         0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

Sample Output

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
```

Copyright © 2016, Juniper Networks, Inc.

```
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %           bps     %           usec
    0 best-effort            95     950000000    95              0      low
none
    3 network-control         5      50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
```


Copyright © 2016, Juniper Networks, Inc.

```
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Local statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Transit statistics:
 Input  bytes  :                   0                   0 bps
 Output bytes  :                   0                   0 bps
 Input  packets:                   0                   0 pps
 Output packets:                   0                   0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                    0
  ICMP packets :                    0
  VPN packets :                     0
  Multicast packets :               0
  Bytes permitted by policy :       0
  Connections established :         0
Flow Output statistics:
  Multicast packets :               0
  Bytes permitted by policy :       0
Flow error statistics (Packets dropped due to):
  Address spoofing:                 0
  Authentication failed:            0
  Incoming NAT errors:              0
  Invalid zone received packet:     0
  Multiple user authentications:    0
  Multiple incoming NAT:            0
  No parent for a gate:             0
  No one interested in self packets: 0
  No minor session:                 0
  No more sessions:                 0
  No NAT gate:                      0
  No route present:                 0
  No SA for incoming SPI:           0
  No tunnel found:                  0
  No session for a gate:            0
  No zone or NULL zone binding      0
  Policy denied:                    0
  Security association not active:  0
  TCP sequence number out of window: 0
  Syn-attack protection:            0
  User authentication errors:       0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
    Generation: 150
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1              --> 0/0
                                                 10.0.0.16             --> 0/0
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

Sample Output

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                      Admin Link
ce1-1/2/6                                                       up    up
e1-1/2/6                                                        up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                      Admin Link
ce1-1/2/3                                                       up    up

ds-1/2/3:1                                                      up    up

ds-1/2/3:2                                                      up    up
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

Sample Output

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                                Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                        0                    0
                                  (                  0) (                 0)
                       silver                        0                    0
                                  (                  0) (                 0)
  Logical interface so-0/1/3.0
                                                Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                        0                    0
                                  (                  0) (                 0)
                       silver                        0                    0
                                  (                  0) (                 0)
```

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
```

```
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

Sample Output

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
```

```
    Bytes permitted by policy :        3478
    Connections established :          1
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        16994
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
```


Copyright © 2016, Juniper Networks, Inc.

```
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State           Last change  Member count
ams0       Up              1d 00:50     2
ams1       Up              00:00:59     2
```

show interfaces load-balancing detail

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
```


Copyright © 2016, Juniper Networks, Inc.

```
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
```

Copyright © 2016, Juniper Networks, Inc.

```
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues      : 8 supported
  Last flapped    : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  SONET alarms    : None
  SONET defects   : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :            118750000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :              6250000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
```

Sample Output

show interfaces redundancy

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present               sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

show interfaces redundancy (Aggregated Ethernet)

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up

ae0
ae1
```

Copyright © 2016, Juniper Networks, Inc.

```
ae2
ae3
ae4
```

show interfaces redundancy detail

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

Sample Output

show interfaces routing brief

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1    up      up     inet     10.0.0.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0  up      up     inet     10.0.0.1/32
```

Sample Output

show interfaces snmp-index

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

Sample Output

show interfaces source-class all

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                                Packets                Bytes
      Source class                  (packet-per-second)    (bits-per-second)
                              gold              1928095            161959980
                                   (                889) (            597762)
                            bronze                    0                    0
```

```
                                        (                  0) (                  0)
                              silver                       0                     0
                                        (                  0) (                  0)
  Logical interface so-0/1/3.0
                                                      Packets                 Bytes
      Source class                     (packet-per-second)    (bits-per-second)
                                gold                       0                     0
                                        (                  0) (                  0)
                              bronze                       0                     0
                                        (                  0) (                  0)
                              silver                  116113               9753492
                                        (                939) (             631616)
```

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                                      Packets                 Bytes
      Destination class               (packet-per-second)    (bits-per-second)
                              silver1                      0                     0
                                        (                  0) (                  0)
                              silver2                      0                     0
                                        (                  0) (                  0)
                              silver3                      0                     0
                                        (                  0) (                  0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

Sample Output

show interfaces switch-port

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                          Receive         Transmit
          Total bytes                       28437086         21792250
```


Copyright © 2016, Juniper Networks, Inc.

```
          Total packets                        409145            88008
          Unicast packets                        9987            83817
          Multicast packets                    145002                0
          Broadcast packets                    254156             4191
          Multiple collisions                      23               10
          FIFO/CRC/Align errors                     0                0
          MAC pause frames                          0                0
          Oversized frames                          0
          Runt frames                               0
          Jabber frames                             0
          Fragment frames                           0
          Discarded frames                          0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

Sample Output

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr         2008544300       0                NA              NA
  FEC-UncorrectedWords     0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
```

Copyright © 2016, Juniper Networks, Inc.

```
    PM                  MIN       MAX       AVG       THRESHOLD       TCA-ENABLED
     TCA-RAISED
  BER                   3.6e-5    5.8e-5    3.6e-5     10.0e-3        No
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                           CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                                 (MIN)
(MAX)    (MIN) (MAX)       (MIN) (MAX)
  Lane chromatic dispersion          0         0         0         0         0
 0        NA    NA         NA     NA
  Lane differential group delay      0         0         0         0         0
 0        NA    NA         NA     NA
    q Value                             120       120       120       120       0
   0        NA    NA         NA     NA
  SNR                                  28        28        29        28        0
 0        NA    NA         NA     NA
  Tx output power(0.01dBm)           -5000     -5000     -5000     -5000     -300
 -100     No    No         No     No
  Rx input power(0.01dBm)            -3642     -3665     -3626     -3637     -1800
 -500     No    No         No     No
  Module temperature(Celsius)        46        46        46        46        -5
 75       No    No         No     No
  Tx laser bias current(0.1mA)       0         0         0         0         0
 0        NA    NA         NA     NA
  Rx laser bias current(0.1mA)       1270      1270      1270      1270      0
 0        NA    NA         NA     NA
  Carrier frequency offset(MHz)      -186      -186      -186      -186      -5000
 5000     No    No         No     No
```

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
```


Copyright © 2016, Juniper Networks, Inc.

```
Interfaces bound: 1
Interfaces:
  ge-0/0/2.0
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces diagnostics optics

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP) installed in SRX Series Services Gateways. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options *interface-name*—Name of the interface associated with the port in which the transceiver is installed: **ge-*fpc/pic/port*** .

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

List of Sample Output show interfaces diagnostics optics on page 2849

Output Fields Table 195 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 283: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is **On** or **Off**.

Copyright © 2016, Juniper Networks, Inc.

Table 283: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is **On** or **Off**.
Laser bias current high warning	Displays whether the laser bias power setting high warning is **On** or **Off**.
Laser bias current low warning	Displays whether the laser bias power setting low warning is **On** or **Off**.
Laser output power high alarm	Displays whether the laser output power high alarm is **On** or **Off**.
Laser output power low alarm	Displays whether the laser output power low alarm is **On** or **Off**.
Laser output power high warning	Displays whether the laser output power high warning is **On** or **Off**.
Laser output power low warning	Displays whether the laser output power low warning is **On** or **Off**.
Module temperature high alarm	Displays whether the module temperature high alarm is **On** or **Off**.
Module temperature low alarm	Displays whether the module temperature low alarm is **On** or **Off**.
Module temperature high warning	Displays whether the module temperature high warning is **On** or **Off**.
Module temperature low warning	Displays whether the module temperature low warning is **On** or **Off**.
Module voltage high alarm	Displays whether the module voltage high alarm is **On** or **Off**.
Module voltage low alarm	Displays whether the module voltage low alarm is **On** or **Off**.
Module voltage high warning	Displays whether the module voltage high warning is **On** or **Off**.
Module voltage low warning	Displays whether the module voltage low warning is **On** or **Off**.
Laser rx power high alarm	Displays whether the receive laser power high alarm is **On** or **Off**.
Laser rx power low alarm	Displays whether the receive laser power low alarm is **On** or **Off**.

Copyright © 2016, Juniper Networks, Inc.

Table 283: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Laser rx power high warning	Displays whether the receive laser power high warning is **On** or **Off**.
Laser rx power low warning	Displays whether the receive laser power low warning is **On** or **Off**.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.

Copyright © 2016, Juniper Networks, Inc.

Table 283: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
```

Copyright © 2016, Juniper Networks, Inc.

```
Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
Module temperature high warning threshold :  90 degrees C / 194 degrees F
Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
Module voltage high alarm threshold       :  3.900 V
Module voltage low alarm threshold        :  2.700 V
Module voltage high warning threshold     :  3.700 V
Module voltage low warning threshold      :  2.900 V
Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces flow-statistics

Syntax **show interfaces flow-statistics** *<interface-name>*

Release Information Command introduced in Junos OS Release 9.2.

Description Display interfaces flow statistics.

Options *Interface-name*—(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the "Interface Naming Conventions" on page 2446.

- **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
- **br-*pim*/0/*port***—Basic Rate Interface for establishing ISDN connections.
- **ce1-*pim*/0/*port***—Channelized E1 interface.
- **ct1-*pim*/0/*port***—Channelized T1 interface.
- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
- **e1-*pim*/0/*port***—E1 interface.
- **e3-*pim*/0/*port***—E3 interface.
- **fe-*pim*/0/ *port***—Fast Ethernet interface.
- **ge-*pim*/0/*port***—Gigabit Ethernet interface.
- **se-*pim*/0/*port***—Serial interface.
- **t1-*pim*/0/*port***—T1 (also called DS1) interface.
- **t3-*pim*/0/ *port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

List of Sample Output show interfaces flow-statistics (Gigabit Ethernet) on page 2854

Output Fields Table 196 lists the output fields for the **show interfaces flow-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 284: show interfaces flow-statistics Output Fields

Field Name	Field Description
Traffic statistics	Number of packets and bytes transmitted and received on the physical interface.

Copyright © 2016, Juniper Networks, Inc.

Table 284: show interfaces flow-statistics Output Fields *(continued)*

Field Name	Field Description
Local statistics	Number of packets and bytes transmitted and received on the physical interface.
Transit statistics	Number of packets and bytes transiting the physical interface.
Flow input statistics	Statistics on packets received by flow module.
Flow output statistics	Statistics on packets sent by flow module.
Flow error statistics	Packet drop statistics for the flow module. For further details, see Table 197.

Table 285: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.
Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.
User authentication errors	Packet was dropped because policy requires authentication; however: • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.

Copyright © 2016, Juniper Networks, Inc.

Table 285: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Flow:	
No one interested in self packets	This counter is incremented for one of the following reasons: • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	The packet was dropped because a valid route was not available to forward the packet. For new sessions, the counter is incremented for one of the following reasons: • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because: • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found
No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.
Policy denied	The error counter is incremented for one of the following reasons: • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.
TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-

Copyright © 2016, Juniper Networks, Inc.

Table 285: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Invalid zone received packet	-
No NAT gate	-

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
      Bytes permitted by policy :       3478
      Connections established :         1
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 0
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
```


Copyright © 2016, Juniper Networks, Inc.

```
Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces queue

Syntax

show interfaces queue
<both-ingress-egress>
<egress>
<forwarding-class *forwarding-class*>
<ingress>
<interface-name *interface-name*>
<l2-statistics>

Release Information

Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

both-ingress-egress—Display both ingress and egress queue statistics.

egress—Display egress queue statistics.

forwarding-class *forwarding-class*—(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue that is associated with the specified forwarding class.

ingress—Display ingress queue statistics.

interface-name *interface-name*—(Optional) Show detailed CoS queue statistics for the specified interface.

l2-statistics—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles.

Required Privilege Level

view

Related Documentation

- Understanding Class of Service

List of Sample Output

show interfaces queue on page 2858

Output Fields

Table 286 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 286: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.
Interface index	Index number of the physical interface. The number reflects the interface's initialization sequence.

Copyright © 2016, Juniper Networks, Inc.

Table 286: show interfaces queue Output Fields *(continued)*

Field Name	Field Description
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
The following output fields are applicable to both the interface component and Packet Forwarding Engine component in the **show interfaces queue** command:	
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets in this queue.
Queued Bytes	Number of bytes in this queue.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Number of packets dropped because of tail drop.
RL-dropped bytes	Number of bytes dropped because of rate limiting.
RED-dropped packets	Number of packets dropped because of random early detection (RED).
RED-dropped bytes	Number of bytes dropped because of RED. • **Low, non-TCP**—Number of low-loss priority, non-TCP bytes dropped because of RED. • **Low, TCP**—Number of low-loss priority, TCP bytes dropped because of RED. • **High, non-TCP**—Number of high-loss priority, non-TCP bytes dropped because of RED. • **High, TCP**—Number of high-loss priority, TCP bytes dropped because of RED.
Queue Buffer Usage:	• **Reserved buffer**—The size of the memory buffer that is allocated for storing packets • **Current**—The amount of buffer memory that is currently in use on this queue.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
  Queue Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces statistics (View)

Syntax show interfaces statistics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level view

Related Documentation
- Understanding Interfaces on page 2441

List of Sample Output show interfaces statistics on page 2859

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
     Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
     Input packets : 2743333
     Output packets: 6790470992
     Security: Zone: untrust
     Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset
 http https ike netconf ping reverse-telnet
    reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl lsping ntp sip
     Protocol inet, MTU: 9192
       Addresses, Flags: Is-Preferred Is-Primary
         Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces terse zone

Syntax show interfaces terse zone

Release Information Command introduced in Junos OS Release 12.3X48-D20.

Description Display summary information about zone interfaces.

Options This command has no options.

Required Privilege Level view

Sample Output

show interface terse zone

```
user@host> show interface terse zone
Interface     Admin    Link   Proto     Local           Remote     Zone
ge-0/0/0.0      up      up     inet     1.4.253.251/16             trust
```


Copyright © 2016, Juniper Networks, Inc.

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information Command introduced in Junos OS Release 12.1X45-D10.

Description Display information about the IPv6 neighbor cache.

Options This command has no options.

Required Privilege Level view

Related Documentation
- clear ipv6 neighbors on page 2764

List of Sample Output show ipv6 neighbors on page 2861

Output Fields Table 287 lists the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 287: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up, down, incomplete, reachable, stale, or unreachable.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address     Linklayer Address   State       Exp Rtr Secure  Interface
   10:1::2          00:00:0a:00:00:00   reachable   17   yes  no     reth0.0

   11:11::2         00:19:e2:4b:61:83   stale       1197 yes  no     at-1/0/0.0

   12:12::2         00:19:e2:4b:61:83   stale       1188 yes  no     at-3/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show lacp interfaces (View)

Syntax **show lacp interfaces** ***interface-name***

Release Information Command modified in Junos OS Release 10.2.

Description Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet interface, redundant Ethernet interface, Gigabit Ethernet interface, or 10-Gigabit Ethernet interface. If you do not specify an interface name, LACP information for all interfaces is displayed.

Options **none**—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Redundant Ethernet—**reth*number***
- Gigabit Ethernet—**ge-*fpc/pic/port***
- 10-Gigabit Ethernet—**xe-*fpc/pic/port***

NOTE: The **show lacp interfaces** command returns the following error message if your system is not configured in either active or passive LACP mode:

"Warning: lacp subsystem not running – not needed by configuration"

Required Privilege Level view

Related Documentation

- Verifying LACP on Redundant Ethernet Interfaces on page 2579

List of Sample Output

show lacp interfaces (Aggregated Ethernet) on page 2865
show lacp interfaces (Redundant Ethernet) on page 2866
show lacp interfaces (Gigabit Ethernet) on page 2866

Output Fields Table 288 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 288: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Copyright © 2016, Juniper Networks, Inc.

Table 288: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates the actor or partner is in an expired state. **No** indicates the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. **No** indicates the operational partner information in use has been received in a link aggregation control protocol data unit (PDU).
	• **Dist**—Distribution of outgoing frames. **No** indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is **No**, the link is not synchronized. It is currently not in the right aggregation.
	• **Aggr**—Ability of aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• **Timeout**—LACP timeout preference. Periodic transmissions of link aggregation control PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (**Long Timeout** or **Short Timeout**).
	• **Activity**—Actor or partner's port activity. **Passive** indicates the port's preference for not transmitting link aggregation control PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.

Copyright © 2016, Juniper Networks, Inc.

Table 288: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  • **Current**—The state machine receives a link aggregation control PDU and enters the **Current** state.
	  • **Defaulted**—If no link aggregation control PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  • **Expired**—If no link aggregation control PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state.
	  • **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state.
	  • **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port.
	  • **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state.
	• **Transmit State**—Transmit state of state machine. One of the following values:
	  • **Fast Periodic**—Periodic transmissions are enabled at a fast transmission rate.
	  • **No Periodic**—Periodic transmissions are disabled.
	  • **Periodic Timer**—Transitory state entered when the periodic timer expires.
	  • **Slow Periodic**—Periodic transmissions are enabled at a slow transmission rate.
	• **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  • **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
	  • **Collecting Distributing**—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.
	  • **Detached**—Process of detaching the port from the aggregator is in progress.
	  • **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/0                  Current   Fast periodic Collecting distributing
```

Copyright © 2016, Juniper Networks, Inc.

```
      ge-2/0/1                Current   Fast periodic Collecting distributing
      ge-2/2/0                Current   Fast periodic Collecting distributing
      ge-2/2/1                Current   Fast periodic Collecting distributing
```

show lacp interfaces (Redundant Ethernet)

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-11/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-11/0/0                Current   Fast periodic Collecting distributing
      ge-11/0/1                Current   Fast periodic Collecting distributing
      ge-11/0/2                Current   Fast periodic Collecting distributing
      ge-11/0/3                Current   Fast periodic Collecting distributing
      ge-3/0/0                 Current   Fast periodic Collecting distributing
      ge-3/0/1                 Current   Fast periodic Collecting distributing
      ge-3/0/2                 Current   Fast periodic Collecting distributing
      ge-3/0/3                 Current   Fast periodic Collecting distributing
{primary:node1}
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/3/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/3/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State          Mux State
      ge-0/3/0             Current     Fast periodic  Collecting distributing
```


Copyright © 2016, Juniper Networks, Inc.

show lacp statistics interfaces (View)

Syntax **show lacp statistics interfaces** *interface-name*

Release Information Command modified in Junos OS Release 10.2.

Description Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level view

Related Documentation
- Verifying LACP on Redundant Ethernet Interfaces on page 2579
- *Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets*

List of Sample Output show lacp statistics interfaces on page 2867

Output Fields Table 289 lists the output fields for the **show lacp statistics interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 289: show lacp statistics interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	LACP statistics provide the following information: • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.

Sample Output

show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-2/0/0                1352        2035            0            0
      ge-2/0/1                1352        2056            0            0
      ge-2/2/0                1352        2045            0            0
      ge-2/2/1                1352        2043            0            0
```

Copyright © 2016, Juniper Networks, Inc.

show oam ethernet link-fault-management

Syntax
show oam ethernet link-fault-management
<brief | detail>
<interface-name>

Release Information
Statement for branch SRX Series devices introduced in Junos OS Release 9.5.

Description
Display Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options
brief | detail—(Optional) Display the specified level of output.

interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level
view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097
- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 2603
- Example: Configuring Ethernet OAM Link Fault Management on page 2605

List of Sample Output
show oam ethernet link-fault-management brief on page 2872
show oam ethernet link-fault-management detail on page 2872

Output Fields
Table 290 lists the output fields for the **show oam ethernet link-fault-management** command. Output fields are listed in the approximate order in which they appear.

Table 290: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Status of the established link. • **Fail**—A link fault condition exists. • **Running**—A link fault condition does not exist.	All levels
Discovery state	State of the discovery mechanism: • **Passive Wait** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok**	All levels
Peer address	Address of the OAM peer.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 290: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. • **Remote-Stable**—Indicates remote OAM client acknowledgment of, and satisfaction with, local OAM state information. **False** indicates that remote DTE has either not seen or is unsatisfied with local state information. **True** indicates that remote DTE has seen and is satisfied with local state information. • **Local-Stable**—Indicates local OAM client acknowledgment of, and satisfaction with, remote OAM state information. **False** indicates that local DTE either has not seen or is unsatisfied with remote state information. **True** indicates that local DTE has seen and is satisfied with remote state information. • **Remote-State-Valid**—Indicates the OAM client has received remote state information found within local information TLVs (type, length, values) of received Information OAM PDUs. **False** indicates that the OAM client has not seen remote state information. **True** indicates that the OAM client has seen remote state information.	All levels
Remote loopback status	An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	Remote entity information. • **Remote MUX action**—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • **Remote parser action**—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to the higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • **Discovery mode**—Indicates whether discovery mode is active or inactive. • **Unidirectional mode**—Indicates the ability to operate a link in unidirectional mode for diagnostic purposes. • **Remote loopback mode**—Indicates whether remote loopback is supported or not supported. • **Link events**—Indicates whether interpreting link events is supported or not supported on the remote peer. • **Variable requests**—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.	All levels
OAM Receive Statistics		
Information	Number of information PDUs received.	**detail**
Event	Number of loopback control PDUs received.	**detail**
Variable request	Number of variable request PDUs received.	**detail**
Variable response	Number of variable response PDUs received.	**detail**
Loopback control	Number of loopback control PDUs received.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 290: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Organization specific	Number of vendor organization specific PDUs received.	**detail**
OAM Transmit Statistics		
Information	Number of information PDUs transmitted.	**detail**
Event	Number of event notification PDUs transmitted.	**detail**
Variable request	Number of variable request PDUs transmitted.	**detail**
Variable response	Number of variable response PDUs transmitted.	**detail**
Loopback control	Number of loopback control PDUs transmitted.	**detail**
Organization specific	Number of vendor organization specific PDUs transmitted.	**detail**
OAM Received Symbol Error Event information		
Events	Number of symbol error event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	**detail**
Threshold	Number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	Number of symbol errors in the period reported in the received event PDU.	**detail**
Total errors	Number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	**detail**
OAM Received Frame Error Event Information		
Events	Number of errored frame event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Duration of the window in terms of the number of 100 ms period intervals.	**detail**
Threshold	Number of detected errored frames required for the event to be generated.	**detail**
Errors in period	Number of detected errored frames in the period.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 290: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Total errors	Number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	**detail**
OAM Received Frame Period Error Event Information		
Events	Number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Duration of the frame seconds window.	**detail**
Threshold	Number of frame seconds errors in the period.	**detail**
Errors in period	Number of frame seconds errors in the period.	**detail**
Total errors	Number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	**detail**
OAM Transmitted Symbol Error Event Information		
Events	Number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the transmitted PDU.	**detail**
Threshold	Number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	Number of symbol errors in the period reported in the transmitted event PDU.	**detail**
Total errors	Number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
OAM Transmitted Frame Error Event Information		
Events	Number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	Duration of the window in terms of the number of 100-ms period intervals.	**detail**
Threshold	Number of detected errored frames required for the event to be generated.	**detail**
Errors in period	Number of detected errored frames in the period.	**detail**
Total errors	Number of errored frames that have been detected after the OAM sublayer was reset.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
  Status: Running, Discovery state: Send Any
  Peer address: 00:90:69:72:2c:83
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  Remote loopback status: Disabled on local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-0/0/1
    Status: Running, Discovery state: Send Any
    Peer address: 00:90:69:0a:07:14
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 186365, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM transmit statistics:
      Information: 186347, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM received symbol error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: supported, Link events: supported
      Variable requests: unsupported
```


Copyright © 2016, Juniper Networks, Inc.

show poe controller (View)

Syntax show poe controller

Release Information Command introduced in Junos OS Release 9.5.

Description Display the status of the Power over Ethernet (PoE) controller.

Options **none**—Display general parameters of the PoE software module controller.

Required Privilege Level View

Related Documentation
- Example: Configuring PoE on All Interfaces on page 2613

Output Fields Table 291 lists the output fields for the **show poe controller** command. Output fields are listed in the approximate order in which they appear.

Table 291: show poe controller Output Fields

Field name	Field Description
Controller-index	Identifies the controller.
Maximum-power	Specifies the maximum power that can be provided by the SRX Series device to PoE ports.
Power-consumption	Specifies the total amount of power allocated to the PoE ports.
Guard-band	Shows the guard band configured on the controller.
Management	Shows the power management mode.

Sample Output

show poe controller

user@host>show poe controller

```
Controller  Maximum   Power              Guard band  Management
index       power     consumption
   0        150.0 W    0.0 W               0 W         Static
```

Copyright © 2016, Juniper Networks, Inc.

show pppoe interfaces

Syntax
show pppoe interfaces
<brief | detail | extensive>
<pp0.*logical*>

Release Information
Command introduced in Junos OS Release 9.5.

Description
Display session-specific information about PPPoE interfaces.

Options
none—Display interface information for all PPPoE interfaces.

brief | detail—(Optional) Display the specified level of output.

extensive—(Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.

pp0.*logical*—(Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16,385. The logical unit number for dynamic interfaces can be a value from 1,073,741,824 through the maximum number of logical interfaces supported on your SRX300, SRX320, and SRX340, and SRX550 devices.

Required Privilege Level
view

Related Documentation
- Understanding Ethernet Interfaces on page 2541

List of Sample Output
show pppoe interfaces on page 2876
show pppoe interfaces brief on page 2876
show pppoe interfaces detail on page 2876
show pppoe interfaces extensive on page 2876

Output Fields
Table 292 lists the output fields for the **show pppoe interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 292: show pppoe interfaces Output Fields

Field Name	Field Description
Index	Index number of the logical interface, which reflects its initialization sequence.
State	State of the logical interface: **up** or **down**.
Session ID	Session ID.
Service name	Type of service required (can be used to indicate an ISP name, a class, or quality of service).
Configured AC name	Configured access concentrator name.
Session AC name	Name of the access concentrator.

Copyright © 2016, Juniper Networks, Inc.

Table 292: show pppoe interfaces Output Fields *(continued)*

Field Name	Field Description
Remote MAC address or **Remote MAC**	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.
Auto-reconnect timeout	Timeout value for reconnecting after a PPPoE session is terminated (in seconds).
Idle timeout	Length of time (in seconds) that a connection can be idle before disconnecting.
Session uptime	Length of time the session has been up, in *hh:mm:ss*.
Ignore End-Of-List tag	Disables the **End-of-List** tag to continue processing of other tags after the **End-of-List** tag in a PPPoE Active Discovery Offer (PADO) packet.
Underlying interface	Interface on which PPPoE is running.
Packet Type	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.
Timeout	Timeouts that occur during the PPPoE session: • **PADI**—No PADI packets received within the timeout period. • **PADO**—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • **PADR**—No PADR packets received within the timeout period.
Receive Error Counters	Error counters received during the PPPoE session: • **PADI**—No PADI error counters received during the session. • **PADO**—No PADO error counters received during the session. • **PADR**—No PADR error counters received during the session. • **PADS**—No PADS error counters received during the session.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
Interface       Underlying            State       Session    Remote
                interface                         ID         MAC
pp0.0           ge-0/0/1.0            Session up  4          b0:c6:9a:74:5e:c1
```

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  Ignore End-Of-List tag: Enable
```

show pppoe interfaces extensive

```
user@host> show pppoe interfaces extensive
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:22 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PacketType                      Sent         Received
    PADI                             1                0
    PADO                             0                1
    PADR                             1                0
    PADS                             0                1
    PADT                             0                0
    Service name error               0                0
    AC system error                  0                0
    Generic error                    0                0
    Malformed packets                0                0
    Unknown packets                  0                0
  Timeout
    PADI                             0
    PADO                             0
    PADR                             0
  Receive Error Counters
```


Copyright © 2016, Juniper Networks, Inc.

PADI	0
PADO	0
PADR	0
PADS	0

Copyright © 2016, Juniper Networks, Inc.

show pppoe statistics

Syntax show pppoe statistics
<*logical-interface-name*>

Release Information Command is t introduced in Junos OS Release 9.5.

Description Display statistics information about PPPoE interfaces.

Options **none**—Display PPPoE statistics for all interfaces.

logical-interface-name—(Optional) Name of an underlying PPPoE logical interface.

Required Privilege Level view

Related Documentation
- show pppoe interfaces on page 2874
- Understanding Ethernet Interfaces on page 2541

List of Sample Output show pppoe statistics on page 2879

Output Fields Table 293 lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 293: show pppoe statistics Output Fields

Field Name	Field Description
Active PPPoE sessions	Total number of active PPPoE sessions.
Packet Type	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.

Copyright © 2016, Juniper Networks, Inc.

Table 293: show pppoe statistics Output Fields *(continued)*

Field Name	Field Description
Timeout	Timeouts that occur during the PPPoE session: • **PADI**—No PADI packets received within the timeout period. • **PADO**—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • **PADR**—No PADR packets received within the timeout period.
Receive Error Counters	Error counters received during the PPPoE session: • **PADI**—No PADI error counters received during the session. • **PADO**—No PADO error counters received during the session. • **PADR**—No PADR error counters received during the session. • **PADS**—No PADS error counters received during the session.

Sample Output

show pppoe statistics

```
user@host> show pppoe statistics
Active PPPoE sessions: 0

  PacketType                      Sent          Received
    PADI                             0                 0
    PADO                             0                 0
    PADR                             0                 0
    PADS                             0                 0
    PADT                             0                 0
    Service name error               0                 0
    AC system error                  0                 0
    Generic error                    0                 0
    Malformed packets                0                 0
    Unknown packets                  0                 0
  Timeout
    PADI                             0
    PADO                             0
    PADR                             0
  Receive Error Counters
    PADI                             0
    PADO                             0
    PADR                             0
    PADS                             0
```

Copyright © 2016, Juniper Networks, Inc.

show poe telemetries

Syntax show poe telemetries
<interface *interface-name* count *number*>
<count *number* interface *interface-name*>

Release Information Command modified in Junos OS Release 12.3X48-D10.

Description Display a history of power consumption on the specified interface. Telemetries must be enabled on the interface before you can display a history of power consumption.

Options
- **Interface *interface-name***—Display telemetries for the specified PoE interface.
- **count *number***—Display the specified number of telemetries records for the specified PoE interface.

Required Privilege Level View

Related Documentation
- Example: Configuring PoE on All Interfaces on page 2613

Output Fields Table 294 lists the output fields for the **show poe telemetries interface** command. Output fields are listed in the approximate order in which they appear.

Table 294: show poe telemetries interface Output Fields

Field name	Field Description
Sl No	Number of the record for the specified port. The last record is the most is the most recent.
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Voltage on the specified port at the time the data was gathered.

Sample Output

show poe telemetries interface

user@host>show poe telemetries interface ge-0/0/1 count 8

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:41:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:40:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:39:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   5     Fri Jan 04 11:37:15 2009 6.6 W    47.2 V
   6     Fri Jan 04 11:36:15 2009 6.6 W    47.2 V
   7     Fri Jan 04 11:35:15 2009 6.6 W    47.2 V
```


Copyright © 2016, Juniper Networks, Inc.

```
   8     Fri Jan 04 11:34:15 2009 6.6 W    47.2 V
```

user@host>show poe telemetries count 5 interface ge-0/0/1

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:47:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:29:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:11:15 2009 6.6 W    47.2 V
   5     Fri Jan 04 11:10:15 2009 6.6 W    47.2 V
```

Copyright © 2016, Juniper Networks, Inc.

show services accounting

Syntax

show services accounting
aggregation
errors
<inline-jflow | inline-jflow fpc-slot *slot number*>
flow
<inline-jflow | inline-jflow fpc-slot *slot number*>
flow-detail
memory
packet-size-distribution
status
<inline-jflow | inline-jflow fpc-slot *slot number*>
usage

Release Information

Command introduced in Junos OS Release 10.4. The **inline-jflow** and **fpc-slot** options are added in Junos OS Release 12.1X45-D10.

Description

Display sampled accounting service.

Options

- aggregation—Display aggregation information.
- errors —Display error statistics.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow—Display flow information.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow-detail—Display flow detail.
- memory—Display memory information.
- packet-size-distribution—Display packet size distribution.
- status–Display service accounting parameters.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- usage–Display CPU usage.

Required Privilege Level

view

Related Documentation

- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461

Copyright © 2016, Juniper Networks, Inc.

List of Sample Output
show services accounting status inline-jflow on page 2883
show services accounting errors inline-jflow on page 2883
show service accounting flow inline-jflow on page 2883

Output Fields Lists the output fields for the **show services accounting** command.

Sample Output

show services accounting status inline-jflow

```
user@host> show services accounting status inline-jflow
Status information
    FPC Slot: 5
    Export format: IP-FIX(V9)
    IPv4 Route Record Count: 16, IPv6 Route Record Count: 5
    Route Record Count: 21, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
```

show services accounting errors inline-jflow

```
user@host> show services accounting errors inline-jflow
Error Information
  FPC Slot:  5
  PIC Slot:  0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0
  AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

  IPv4 Errors:
  IPv4 Flow Creation Failures: 0
  IPv4 Route Record Lookup Failures: 0
  IPv4 AS Lookup Failures: 0
  IPv4 Export Packet Failures: 0

  IPv6 Errors:
  IPv6 Flow Creation Failures: 0
  IPv6 Route Record Lookup Failures: 0
  IPv6 AS Lookup Failures: 0
  IPv6 Export Packet Failures: 0
```

show service accounting flow inline-jflow

```
user@host> show service accounting flow inline-jflow
Flow Information
FPC Slot:  5
PIC Slot:  0
Flow Packets: 2  Flow Bytes: 0
Active Flows: 1  Total Flows: 2
Flows Exported: 0 Flow Packets Exported: 231
Flows Inactive Timed Out: 1 Flows Active Timed Out: 2

IPv4 Flows:
IPv4 Flow Packets: 1  IPv4 Flow Bytes: 0
IPv4 Active Flows: 1  IPv4 Total Flows: 1
IPv4 Flows Exported: 0 IPv4 Flow Packets Exported: 132
IPv4 Flows Inactive Timed Out: 0 IPv4 Flows Active Timed Out: 1
```

Copyright © 2016, Juniper Networks, Inc.

```
IPv6 Flows:
IPv6 Flow Packets: 1  IPv6 Flow Bytes: 0
IPv6 Active Flows: 0  IPv6 Total Flows: 1
IPv6 Flows Exported: 0 IPv6 Flow Packets Exported: 99
IPv6 Flows Inactive Timed Out: 1 IPv6 Flows Active Timed Out: 1
```


Copyright © 2016, Juniper Networks, Inc.

show services accounting aggregation (View)

Syntax show services accounting aggregation

Release Information Command introduced in Junos OS Release 10.4.

Description Display aggregation information for the accounting service.

Options

- as—Display aggregation type AS.
- destination-prefix—Display aggregation type destination-prefix.
- protocol-port—Display aggregation type protocol-port.
- source-destination-prefix—Display aggregation type source-destination-prefix.
- source-prefix—Display aggregation type source-prefix.
- template—Display aggregation type template.

Required Privilege Level view

Related Documentation

- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461

Copyright © 2016, Juniper Networks, Inc.

show services accounting aggregation template (View)

Syntax show services accounting aggregation template

Release Information Command introduced in Junos OS Release 10.4.

Description Display aggregation type template.

Options
- detail—Display detailed output.
- extensive—Display extensive output.
- template-name—Display name of the template.
- terse—Display terse output (default).

Required Privilege Level view

Related Documentation
- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461

Copyright © 2016, Juniper Networks, Inc.

show services accounting flow-detail (View)

Syntax show services accounting flow-detail

Release Information Command introduced in Junos OS Release 10.4.

Description Display flow detail

Options

- destination-as—Filter term destination AS.
- destination-port—Filter term destination port.
- destination-prefix—Filter term destination prefix.
- detail—Display detailed output.
- extensive—Display extensive output.
- input-snmp-interface-index—Filter term input SNMP interface index.
- limit—Display maximum number of flows to display.
- name-Display name of the service, wildcard, or “all”.
- order—Display order for displaying flows.
- output-snmp-interface-index—Filter term output SNMP interface index.
- proto—Filter term protocol.
- source-as—Filter term source AS.

Required Privilege Level view

Related Documentation

- Configuring Flow Aggregation to Use Version 9 Flow Templates on page 2461

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Layer 2 Bridging and Transparent Mode for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 40

Overview

- Introduction to Layer 2 Bridging and Switching on page 2893

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 126

Introduction to Layer 2 Bridging and Switching

- Layer 2 Transparent and Switching Overview on page 2893

Layer 2 Transparent and Switching Overview

Layer 2 transparent mode provides the ability to deploy the firewall without making changes to the existing routing infrastructure. The firewall is deployed as a Layer 2 bridge with multiple VLAN segments and provides security services within VLAN segments. Secure-wire is a special version of Layer 2 transparent mode that allows bump-in-wire deployment.

Ethernet switching forwards the Ethernet frames within or across the LAN segment (or VLAN) using the Ethernet MAC address information. Ethernet switching on the SRX1500 device is performed in the hardware using ASICs.

Use the **set protocols l2-learning global-mode(transparent-bridge | switching)** command to switch between Layer 2 transparent bridge mode and switching mode. After switching the mode, you must reboot the device for the configuration to take effect.

NOTE: **The default mode for Layer 2 is transparent mode.**

The Layer 2 protocol supported in switching mode is Link Aggregation Control Protocol (LACP).

- LACP is not supported on SRX300 and SRX320 devices.
- LACP is not supported in transparent bridge mode.

You can configure Layer 2 transparent mode on a redundant Ethernet interface. Use the following commands to define a redundant Ethernet interface:

- **set interfaces *interface-name* ether-options redundant-parent *reth-interface-name***
- **set interfaces *reth-interface-name* redundant-ether-options redundancy-group *number***

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- **global-mode (Protocols) on page 3056**
- **l2-learning (Protocols) on page 3063**

Copyright © 2016, Juniper Networks, Inc.

PART 41

Configuring Layer 2 Bridging and Transparent Mode

- Configuring Bridging and Transparent Mode on page 2897
- Configuring Interfaces on page 2905
- Configuring Security Zones and Security Policies on page 2921
- Configuring Layer 2 Forwarding Tables on page 2927
- Configuring Layer 2 Transparent Mode Chassis Clusters on page 2931
- Configuring IP Spoofing in Layer 2 Transparent Mode on page 2935
- Configuring Class of Service in Transparent Mode on page 2939
- Configuring IPv6 Flows on page 2947
- Configuring Secure Wire on page 2955

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 127

Configuring Bridging and Transparent Mode

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding VLANs on page 2900
- Example: Configuring VLANs on page 2901
- Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 2902

Layer 2 Bridging and Transparent Mode Overview

For SRX Series devices, transparent mode provides full security services for Layer 2 bridging capabilities. On these SRX Series devices, you can configure one or more bridge domains to perform Layer 2 bridging. A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics. Like a virtual LAN (VLAN), a bridge domain spans one or more ports of multiple devices. Thus, the SRX Series device can function as a Layer 2 switch with multiple bridge domains that participate in the same Layer 2 network.

In transparent mode, the SRX Series device filters packets that traverse the device without modifying any of the source or destination information in the IP packet headers. Transparent mode is useful for protecting servers that mainly receive traffic from untrusted sources because there is no need to reconfigure the IP settings of routers or protected servers.

In transparent mode, all physical ports on the device are assigned to Layer 2 interfaces. Do not route Layer 3 traffic through the device. Layer 2 zones can be configured to host Layer 2 interfaces, and security policies can be defined between Layer 2 zones. When packets travel between Layer 2 zones, security policies can be enforced on these packets.

Table 295 lists the security features that are supported and are not supported in transparent mode for Layer 2 bridging.

Copyright © 2016, Juniper Networks, Inc.

Table 295: Security Features Supported in Transparent Mode

Mode Type	Supported	Not Supported
Transparent mode	• Application Layer Gateways (ALGs) • Firewall User Authentication (FWAUTH) • Intrusion Detection and Prevention (IDP) • Screen • AppSecure	• Network Address Translation (NAT) • VPN

NOTE:

- On all SRX Series devices, bridging and transparent mode are not supported on mPIMs.
- On all branch SRX Series devices, the DHCP server propagation is not supported in Layer 2 transparent mode.

Layer 2 Bridging Exceptions on SRX Series Devices

The bridging functions on the SRX Series devices are similar to the bridging features on Juniper Networks MX Series routers. However, the following Layer 2 networking features on MX Series routers are not supported on SRX Series devices:

- Layer 2 control protocols—These protocols are used on MX Series routers for Rapid Spanning Tree Protocol (RSTP) or Multiple Spanning Tree Protocol (MSTP) in customer edge interfaces of a VPLS routing instance.
- Virtual switch routing instance—The virtual switching routing instance is used on MX Series routers to group one or more VLANs.
- Virtual private LAN services (VPLS) routing instance—The VPLS routing instance is used on MX Series routers for point-to-multipoint LAN implementations between a set of sites in a VPN.

In addition, the SRX Series devices do not support the following Layer 2 features:

- Spanning Tree Protocol (STP), RSTP, or MSTP—It is the user's responsibility to ensure that no flooding loops exist in the network topology.
- Internet Group Management Protocol (IGMP) snooping—Host-to-router signaling protocol for IPv4 used to report their multicast group memberships to neighboring routers and determine whether group members are present during IP multicasting.

Copyright © 2016, Juniper Networks, Inc.

- Double-tagged VLANs or IEEE 802.1Q VLAN identifiers encapsulated within 802.1Q packets (also called “Q in Q” VLAN tagging)—Only untagged or single-tagged VLAN identifiers are supported on SRX Series devices.
- Nonqualified VLAN learning, where only the MAC address is used for learning within the VLAN—VLAN learning on SRX Series devices is qualified; that is, both the VLAN identifier and MAC address are used.

Also, on SRX300, SRX320, SRX340, and SRX345 devices, the following features are not supported for Layer 2 transparent mode:

- G-ARP on the Layer 2 interface
- IP address monitoring on any interface
- Transit traffic through IRB
- IRB interface in a routing instance
- IRB interface handling of Layer 3 traffic

NOTE: **The IRB interface is a pseudointerface and does not belong to the reth interface and redundancy group.**

Layer 2 Transparent Mode on the SRX5000 Line Module Port Concentrator

The SRX5000 line Module Port Concentrator (SRX5K-MPC) supports Layer 2 transparent mode and processes the traffic when the SRX Series device is configured in Layer 2 transparent mode.

When the SRX5K-MPC is operating in Layer 2 mode, you can configure all interfaces on the SRX5K-MPC as Layer 2 bridging ports to support Layer 2 traffic.

The security processing unit (SPU) supports all security services for Layer 2 bridging functions, and the MPC delivers the ingress packets to the SPU and forwards the egress packets that are encapsulated by the SPU to the outgoing interfaces.

When the SRX Series device is configured in Layer 2 transparent mode, you can enable the interfaces on the MPC to work in Layer 2 mode by defining one or more logical units on a physical interface with the family address type as **bridge**. Later you can proceed with configuring Layer 2 security zones and configuring security policies in transparent mode. Once this is done, next-hop topologies are set up to process ingress and egress packets.

Related Documentation

- Understanding VLANs on page 2900
- Understanding Transparent Mode Conditions on page 2905
- Understanding Layer 2 Interfaces on page 2906
- Understanding Layer 2 Security Zones on page 2921
- Understanding Security Policies in Transparent Mode on page 2923

Copyright © 2016, Juniper Networks, Inc.

Understanding VLANs

The packets that are forwarded within a VLAN are determined by the VLAN ID of the packets and the VLAN ID of the VLAN. Only the packets with VLAN IDs that match the VLAN ID configured for a VLAN are forwarded within the VLAN.

When configuring VLANs, you can specify either a single VLAN ID or a list of specific VLAN IDs. If you specify a list of VLAN IDs, a VLAN is created for each VLAN ID in the list. Certain VLAN properties, such as the integrated routing and bridging interface (IRB), are not configurable if VLANs are created in this manner.

Each Layer 2 logical interface configured on the device is implicitly assigned to a VLAN based on the VLAN ID of the packets accepted by the interface. You do not need to explicitly define the logical interfaces when configuring a VLAN.

You can configure one or more static MAC addresses for a logical interface in a VLAN; this is only applicable if you specified a single VLAN ID when creating the VLAN.

NOTE: **If a static MAC address you configure for a logical interface appears on a different logical interface, packets sent to that interface are dropped.**

You can configure the following properties that apply to all VLANs on the SRX Series device:

- Layer 2 address learning—Layer 2 address learning is enabled by default. A bridge domain learns unicast media access control (MAC) addresses to avoid flooding packets to all interfaces in the bridge domain. Each bridge domain creates a source MAC entry in its forwarding tables for each source MAC address learned from packets received on interfaces that belong to the bridge domain. When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into a bridge domain.
- Maximum number of MAC addresses learned from all logical interfaces on the SRX Series device—After the MAC address limit is reached, the default is for any incoming packets with a new source MAC address to be forwarded. You can specify that the packets be dropped instead. The default limits of MAC addresses for the SRX Series devices are shown in Table 296.

Table 296: MAC Addresses Default Limits

SRX Series Devices	Default Limit for MAC Addresses
SRX300	16383
SRX320	16383
SRX340	16383
SRX345	16383

Copyright © 2016, Juniper Networks, Inc.

Table 296: MAC Addresses Default Limits *(continued)*

SRX Series Devices	Default Limit for MAC Addresses
SRX5600	131,071
SRX5800	131,071

- Timeout interval for MAC table entries. By default, the timeout interval for MAC table entries is 300 seconds. The minimum you can configure is 10 seconds and the maximum is 64,000 seconds. The timeout interval applies only to dynamically learned MAC addresses. This value does not apply to configured static MAC addresses, which never time out.

NOTE: **SRX300, SRX320, SRX340, and SRX345 devices support only 16,000 MAC entries.**

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Example: Configuring VLANs on page 2901
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Understanding Layer 2 Interfaces on page 2906
- Understanding Layer 2 Forwarding Tables on page 2927

Example: Configuring VLANs

This example shows how to configure VLANs.

- Requirements on page 2901
- Overview on page 2901
- Configuration on page 2902
- Verification on page 2902

Requirements

Before you begin, determine the properties you want to configure for the VLAN. See "Understanding VLANs" on page 2900.

Overview

In this example, you configure bridge domain vlan-a for VLANs 1 and 10, and bridge domain vlan-b for VLAN 2. You then limit the number of MAC addresses learned on all logical interfaces on the device to 64,000. When this limit is reached, incoming packets with a new source MAC address will be dropped.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To configure VLANs:

1. Configure the domain type and VLANs.

   ```
   [edit]
   user@host# set vlans vlan2 vlan-id-list 110-120
   ```

2. Limit the number of MAC addresses.

   ```
   [edit]
   user@host# set protocols l2-learning global-mac-limit 64000 packet-action drop
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show vlans** and **show protocols l2-learning** commands.

Related Documentation

- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Understanding Layer 2 Interfaces on page 2906
- Understanding Layer 2 Forwarding Tables on page 2927
- Understanding VLANs on page 2900

Enhanced Layer 2 CLI Configuration Statement and Command Changes

Table 297 and Table 298 provide lists of existing commands that have been moved to new hierarchies or changed on SRX Series devices as part of this CLI enhancement effort. The tables are provided as a high-level reference only. For detailed information about these commands, see CLI Explorer.

Table 297: Enhanced Layer 2 Configuration Statement Changes

Original Hierarchy	Changed Hierarchy	Hierarchy Level	Change Description

bridge-domains
  bridge-domain--name **{**
  ...
  }
}
	vlans *vlans-name* {
	  ...
	  }
	}
		[edit]	Hierarchy renamed.

bridge-domains
  bridge-domain--name {
  vlan-id-list [*vlan-id*] ;
}
	vlans *vlans-name* {
	  vlan members [*vlan-id*] ;
	}
		[edit vlans *vlans-name*]	Statement renamed.

Copyright © 2016, Juniper Networks, Inc.

Table 297: Enhanced Layer 2 Configuration Statement Changes *(continued)*

Original Hierarchy	Changed Hierarchy	Hierarchy Level	Change Description

bridge-options {
interface *interface-name* {
encapsulation-type;
ignore-encapsulation-mismatch;
pseudowire-status-tlv;
static-mac *mac-address* {
vlan-id *vlan-id*;
}
}
mac-table-aging-time
seconds;
mac-table-size {
number;
packet-action drop;
}
}
	switch-options {
	interface *interface-name* {
	action-priority;
	encapsulation-type;
	ignore-encapsulation-mismatch;
	pseudowire-status-tlv;
	static-mac *mac-address* {
	vlan-id *vlan-id*;
	}
	}
	mac-table-aging-time
	seconds;
	mac-table-size {
	number;
	packet-action drop;
	}
	}
		[edit vlans *vlans-name*]	Statement renamed.

bridge {
block-non-ip-all;
bpdu-vlan-flooding;
bypass-non-ip-unicast;
no-packet-flooding {
no-trace-route;
}
}
	ethernet-switching {
	block-non-ip-all;
	bpdu-vlan-flooding;
	bypass-non-ip-unicast;
	no-packet-flooding {
	no-trace-route;
	}
	}
		[edit security flow]	Statement renamed.

family {
bridge {
bridge-domain-type (svlan|
bvlan);
...
	family {
	ethernet-switching {
	...
		[edit interfaces *interface-name*]
		unit *unit-number*
			Hierarchy renamed.

... **routing-interface** irb.0; ...	... **l3-interface** irb.0; ...	[edit vlans *vlans-name*]	Statement renamed.

Table 298: Enhanced Layer 2 Operational Command Changes

Original Operational Command	Modified Operational Command
clear bridge mac-table	clear ethernet-switching table
clear bridge mac-table persistent-learning	clear ethernet-switching table persistent-learning
show bridge domain	show vlans
show bridge mac-table	show ethernet-switching table
show l2-learning interface	show ethernet-switching interface

Copyright © 2016, Juniper Networks, Inc.

NOTE: There is no fxp0 out-of-band management interface on the SRX300, SRX320, and SRX500HM devices.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 128

Configuring Interfaces

- Understanding Transparent Mode Conditions on page 2905
- Understanding Layer 2 Interfaces on page 2906
- Example: Configuring Layer 2 Logical Interfaces on page 2906
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Example: Configuring an IRB Interface on page 2909
- Understanding Mixed Mode (Layer 2 and Layer 3) on page 2910
- Example: Improving Security Services by Configuring an SRX Series Device Using Mixed Mode (Layer 2 and Layer 3) on page 2913

Understanding Transparent Mode Conditions

A device operates in Layer 2 transparent mode when all physical interfaces on the device are configured as Layer 2 interfaces. A physical interface is a Layer 2 interface if its logical interface is configured with the **bridge** family.

There is no command to define or enable transparent mode on the device. The device operates in transparent mode when there are interfaces defined as Layer 2 interfaces. The device operates in route mode (the default mode) if there are no physical interfaces configured as Layer 2 interfaces.

NOTE: In mixed mode, which is the default mode, you can configure an SRX Series device using both transparent mode (Layer 2) and route mode (Layer 3) simultaneously, with no reboot required.

You can configure the **fxp0** out-of-band management interface on the SRX Series device as a Layer 3 interface, even if Layer 2 interfaces are defined on the device. With the exception of the **fxp0** interface, you can define Layer 2 and Layer 3 interfaces on the device's network ports.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Example: Configuring Layer 2 Logical Interfaces on page 2906
- Understanding Layer 2 Interfaces on page 2906
- Understanding Mixed Mode (Layer 2 and Layer 3) on page 2910

Copyright © 2016, Juniper Networks, Inc.

Understanding Layer 2 Interfaces

Layer 2 logical interfaces are created by defining one or more logical units on a physical interface with the family address type **bridge**. If a physical interface has a **bridge** family logical interface, it cannot have any other family type in its logical interfaces. A logical interface can be configured in one of the following modes:

- Access mode—Interface accepts untagged packets, assigns the specified VLAN identifier to the packet, and forwards the packet within the bridge domain that is configured with the matching VLAN identifier.
- Trunk mode—Interface accepts any packet tagged with a VLAN identifier that matches a specified list of VLAN identifiers. Trunk mode interfaces are generally used to interconnect switches. To configure a VLAN identifier for untagged packets received on the physical interface, use the **native-vlan-id** option. If the **native-vlan-id** option is not configured, untagged packets are dropped.

  Tagged packets arriving on a trunk mode interface can be rewritten or "retagged" with a different VLAN identifier. This allows incoming packets to be selectively redirected to a firewall or other security device.

NOTE: **Multiple trunk mode logical interfaces can be defined, as long as the VLAN identifiers of a trunk interface do not overlap with those of another trunk interface. The native-vlan-id must belong to a VLAN identifier list configured for a trunk interface.**

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Example: Configuring Layer 2 Logical Interfaces on page 2906
- Understanding Transparent Mode Conditions on page 2905

Example: Configuring Layer 2 Logical Interfaces

This example shows how to configure a Layer 2 logical interface as a trunk port so that the incoming packets can be selectively redirected to a firewall or other security device.

- Requirements on page 2906
- Overview on page 2907
- Configuration on page 2907
- Verification on page 2907

Requirements

Before you begin, configure the VLANs. See "Example: Configuring VLANs" on page 2901.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you configure logical interface ge-3/0/0.0 as a trunk port that carries traffic for packets tagged with VLAN identifiers 1 through 10; this interface is implicitly assigned to the previously configured VLANs vlan-a and vlan-b. Then you assign a VLAN ID of 10 to any untagged packets received on physical interface ge-3/0/0.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a Layer 2 logical interface as a trunk port:

1. Configure the logical interface.

   ```
   [edit interfaces ge-3/0/0]
   user@host# set unit 0 family ethernet-switching interface-mode trunk vlan members
     1–10
   ```

2. Specify a VLAN ID for untagged packets.

   ```
   [edit interfaces ge-3/0/0]
   user@host# set vlan-tagging native-vlan-id 10
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces ge-3/0/0** and **show interfaces ge-3/0/0.0** commands.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Layer 2 Interfaces on page 2906
- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Layer 2 Security Zones on page 2922

Understanding Integrated Routing and Bridging Interfaces

For VLANs configured with a single VLAN identifier, you can optionally configure an integrated routing and bridging (IRB) interface for management traffic in the bridge domain. An IRB interface acts as a Layer 3 routing interface for a bridge domain.

NOTE: If you specify a VLAN identifier list in the bridge domain configuration, you cannot configure an IRB interface for the bridge domain.

Copyright © 2016, Juniper Networks, Inc.

Currently the IRB interface on the SRX Series device does not support traffic forwarding or routing. In transparent mode, packets arriving on a Layer 2 interface that are destined for the device's MAC address are classified as Layer 3 traffic while packets that are not destined for the device's MAC address are classified as Layer 2 traffic. Packets destined for the device's MAC address are sent to the IRB interface. Packets from the device's routing engine are sent out the IRB interface.

You create an IRB logical interface in a similar manner as a Layer 3 interface, but the IRB interface does not support traffic forwarding or routing. The IRB interface cannot be assigned to a security zone; however, you can configure certain services on a per-zone basis to allow host-inbound traffic for management of the device. This allows you to control the type of traffic that can reach the device from interfaces bound to a specific zone.

NOTE:

- On SRX1500, SRX5600, and SRX5800 devices, we support an IRB interface that allows you to terminate management connections in transparent mode. However, you cannot route traffic on that interface or terminate IPsec VPNs.
- You can configure only one IRB logical interface for each bridge domain.
- On SRX300 devices, the multicast data traffic is not supported on IRB interfaces.
- On SRX320 devices, IGMPv2 JOINS messages are dropped on an IRB interface. As a workaround, enable IGMP snooping to use IGMP over IRB interfaces.

NOTE: On SRX300, SRX320, SRX340, and SRX345 devices, on the routed VLAN interface (RVI), the following features are not supported:

- IS-IS (family ISO)
- Encapsulations (Ether CCC, VLAN CCC, VPLS, PPPoE, and so on) on VLAN interfaces
- CLNS
- DVMRP
- VLAN interface MAC change
- G-ARP
- Change VLAN-Id for VLAN interface

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Example: Configuring an IRB Interface on page 2909
- Understanding VLANs on page 2900

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring VLANs on page 2901

Example: Configuring an IRB Interface

This example shows how to configure an IRB interface so it can act as a Layer 3 routing interface for a bridge domain.

- Requirements on page 2909
- Overview on page 2909
- Configuration on page 2909
- Verification on page 2910

Requirements

Before you begin, configure a bridge domain with a single VLAN identifier. See "Example: Configuring VLANs" on page 2901.

Overview

In this example, you configure the IRB logical interface unit 0 with the family type inet and IP address 10.1.1.1/24, and then reference the IRB interface irb.0 in the vlan-2 bridge domain configuration. Then you enable Web authentication on the IRB interface and activate the webserver on the device.

NOTE: To complete the Web authentication configuration, you must perform the following tasks:

- Define the access profile and password for a Web authentication client.
- Define the security policy that enables Web authentication for the client.

Either a local database or an external authentication server can be used as the Web authentication server.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an IRB interface:

1. Create a Layer 2 trunk interface.

```
[edit]
user@host# set interfaces ge-0/0/3 unit 0 family ethernet-switching interface-mode
  trunk
```

2. Create an IRB logical interface.

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface irb unit 0 family inet address 10.1.1.1/24
```

3. Create a Layer 2 VLAN.

```
[edit]
user@host# set vlans vlan-2 vlan-id 1
```

4. Associate the IRB interface with the VLAN.

```
[edit]
user@host#set vlans vlan-2 l3-interface irb.0
```

5. Activate the webserver.

```
[edit]
user@host# set system services web-management http
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interface irb** , and **show vlans** commands.

Related Documentation

- *Building Blocks Feature Guide for Security Devices*
- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Example: Configuring Layer 2 Security Zones on page 2922
- Understanding VLANs on page 2900

Understanding Mixed Mode (Layer 2 and Layer 3)

Mixed mode supports both Layer 2 and Layer 3 interfaces; it is the default mode. You can configure both Layer 2 and Layer 3 interfaces simultaneously using separate security zones.

NOTE: **For the mixed mode configuration, you must reboot the device after you commit the changes. However, for SRX5000 line devices, reboot is not required.**

In mixed mode (Layer 2 and Layer 3):

- There is no routing among IRB interfaces and between IRB interfaces and Layer 3 interfaces.
- The user logical system is not supported for Layer 2 traffic. However, you can configure Layer 2 traffic using the root logical system.

Copyright © 2016, Juniper Networks, Inc.

- You can configure Layer 3 interfaces using both the user logical system and the root logical system.

The device in Figure 122 looks like two separate devices. One device runs in Layer 2 mode and the other device runs in Layer 3 mode. But both devices run independently. Packets cannot be transferred between the Layer 2 and Layer 3 interfaces, because there is no routing among IRB interfaces and between IRB interfaces and Layer 3 interfaces.

Figure 122: Architecture of Mixed Layer 2 and Layer 3 Mode

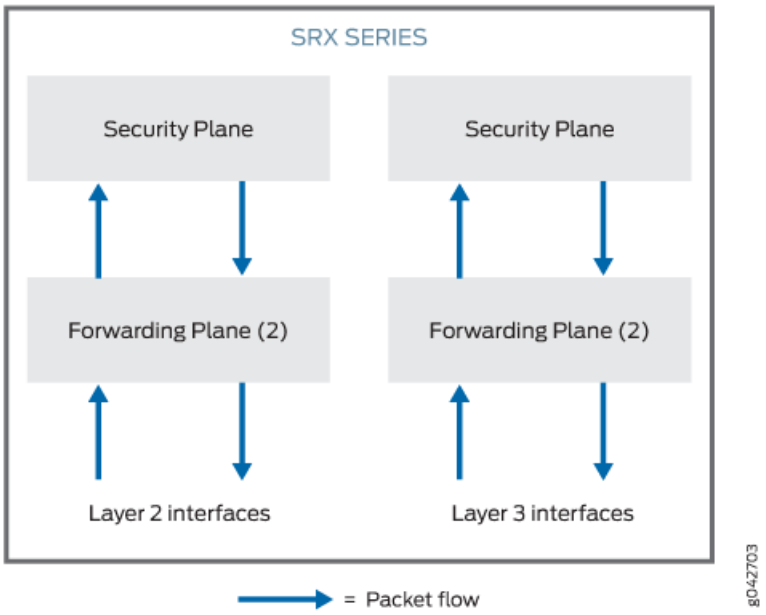

In mixed mode, the Ethernet physical interface can be either a Layer 2 interface or a Layer 3 interface, but the Ethernet physical interface cannot be both simultaneously. However, Layer 2 and Layer 3 families can exist on separate physical interfaces on the same device.

Table 299 lists the Ethernet physical interface types and supported family types.

Table 299: Ethernet Physical Interface and Supported Family Types

Ethernet Physical Interface Type	Supported Family Type
Layer 2 Interface	**bridge**
Layer 3 Interface	**inet** and **inet6**

NOTE: Multiple routing instances are supported.

Copyright © 2016, Juniper Networks, Inc.

You can configure both the pseudointerface **irb.x** and the Layer 3 interface under the same default routing instance using either a default routing instance or a user-defined routing instance. See Figure 123.

Figure 123: Mixed Layer 2 and Layer 3 Mode

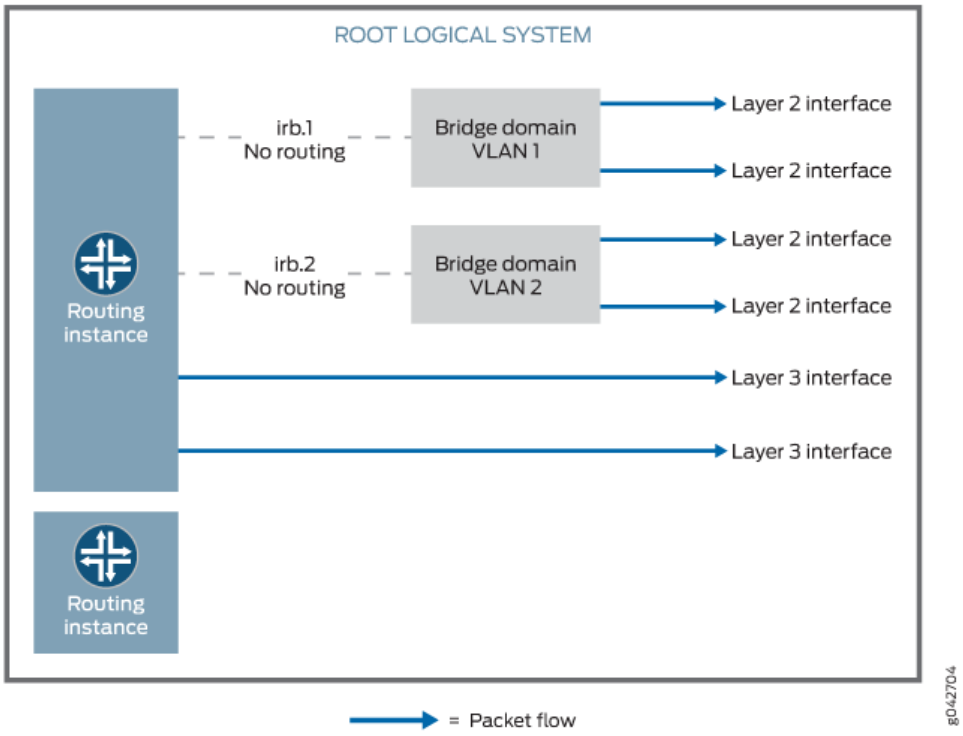

Packets from the Layer 2 interface are switched within the same bridge domain, or they connect to the host through the IRB interface. Packets cannot be routed to another IRB interface or a Layer 3 interface through their own IRB interface.

Packets from the Layer 3 interface are routed to another Layer 3 interface. Packets cannot be routed to a Layer 2 interface through an IRB interface.

Table 300 lists the the security features that are supported in mixed mode and the features that are not supported in transparent mode for Layer 2 bridging..

Copyright © 2016, Juniper Networks, Inc.

Table 300: Security Features Supported in Mixed Mode (Layer 2 and Layer 3)

Mode Type	Supported	Not Supported
Mixed mode	• Application Layer Gateways (ALGs) • Firewall User Authentication (FWAUTH) • Intrusion Detection and Prevention (IDP) • Screen • AppSecure	—
Layer 3 interface of mixed mode	• Network Address Translation (NAT) • VPN	—
Layer 2 mode (transparent mode)		• Network Address Translation (NAT) • VPN

- On all branch SRX Series devices, you cannot configure Ethernet switching and virtual private LAN service (VPLS) using mixed mode (Layer 2 and Layer 3).
- On all branch SRX Series devices, you must reboot the device when you configure bridge domain if the bridge domain was not already configured on the device.
- On all high-end SRX Series devices, you do not have to reboot the device when you configure bridge domain.

Related Documentation

- Example: Improving Security Services by Configuring an SRX Series Device Using Mixed Mode (Layer 2 and Layer 3) on page 2913
- Understanding Secure Wire on page 2955

Example: Improving Security Services by Configuring an SRX Series Device Using Mixed Mode (Layer 2 and Layer 3)

You can configure an SRX Series device using both Layer 2 and Layer 3 interfaces simultaneously to simplify deployments and to improve security services.

This example shows how to pass the Layer 2 traffic from interface ge-0/0/1.0 to interface ge-0/0/0.0 and Layer 3 traffic from interface ge-0/0/2.0 to interface ge-0/0/3.0.

- Requirements on page 2914
- Overview on page 2914
- Configuration on page 2916
- Verification on page 2919

Copyright © 2016, Juniper Networks, Inc.

Requirements

This example uses the following hardware and software components:

- An SRX Series device
- Four PCs

Before you begin:

- Create a separate security zone for Layer 2 and Layer 3 interfaces. See "Understanding Layer 2 Security Zones" on page 2921.

Overview

In enterprises where different business groups have either Layer 2 or Layer 3 based security solutions, using a single mixed mode configuration simplifies their deployments. In a mixed mode configuration, you can also provide security services with integrated switching and routing.

In addition, you can configure an SRX Series device in both standalone and chassis cluster mode using mixed mode.

In mixed mode(default mode), you can configure both Layer 2 and Layer 3 interfaces simultaneously using separate security zones.

NOTE: **For the mixed mode configuration, you must reboot the device after you commit the changes. However, for SRX5000 line devices, reboot is not required.**

In this example, first you configure a Layer 2 family type called bridge to identify Layer 2 interfaces. You set the IP address 10.10.10.1/24 to IRB interface. Then you create zone L2 and add Layer 2 interfaces ge-0/0/1.0 and ge-0/0/0.0 to it.

Next you configure a Layer 3 family type inet to identify Layer 3 interfaces. You set the IP address 192.0.2.1/24 to interface ge-0/0/2.0 and the IP address 192.0.2.3/24 to interface ge-0/0/3. Then you create zone L3 and add Layer 3 interfaces ge-0/0/2.0 and ge-0/0/3.0 to it.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 124 shows a mixed mode topology.

Figure 124: Mixed Mode Topology

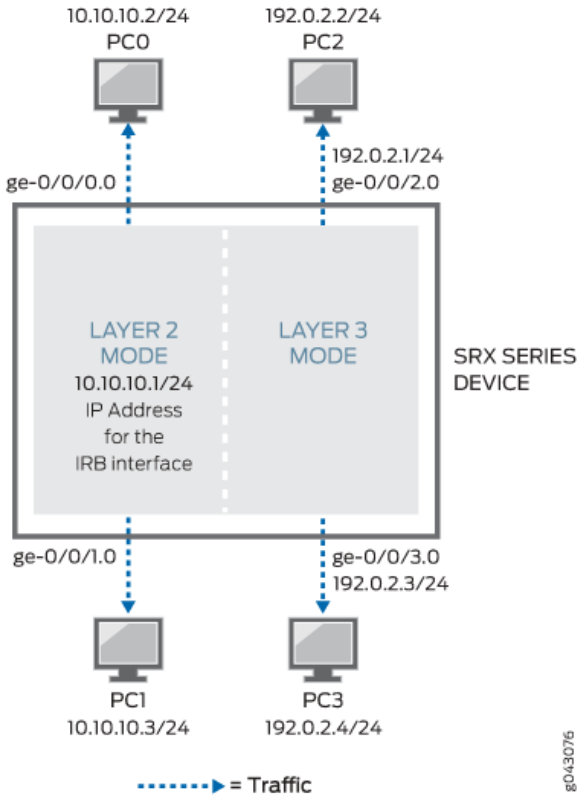

Table 301 shows the parameters configured in this example.

Table 301: Layer 2 and Layer 3 Parameters

Parameter	Description
L2	Layer 2 zone.
ge-0/0/1.0 and ge-0/0/0.0	Layer 2 interfaces added to the Layer 2 zone.
L3	Layer 3 zone.
ge-0/0/2.0 and ge-0/0/3.0	Layer 3 interfaces added to the Layer 3 zone.
10.10.10.1/24	IP address for the IRB interface.
192.0.2.1/24 and 192.0.2.3/24	IP addresses for the Layer 3 interface.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members 10
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 10
set interfaces irb unit 10 family inet address 10.10.10.1/24
set security zones security-zone L2 interfaces ge-0/0/1.0
set security zones security-zone L2 interfaces ge-0/0/0.0
set vlans vlan-10 vlan-id 10
set vlans vlan-10 l3-interface irb.10
set interfaces ge-0/0/2 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/3 unit 0 family inet address 192.0.2.3/24
set security policies default-policy permit-all
set security zones security-zone L2 host-inbound-traffic system-services any-service
set security zones security-zone L2 host-inbound-traffic protocols all
set security zones security-zone L3 host-inbound-traffic system-services any-service
set security zones security-zone L3 host-inbound-traffic protocols all
set security zones security-zone L3 interfaces ge-0/0/2.0
set security zones security-zone L3 interfaces ge-0/0/3.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Layer 2 and Layer 3 interfaces:

1. Create a Layer 2 family type to configure Layer 2 interfaces.

   ```
   [edit interfaces]
   user@host#set ge-0/0/0 unit 0 family ethernet-switching interface-mode access
   user@host#set ge-0/0/0 unit 0 family ethernet-switching vlan members 10
   user@host#set ge-0/0/1 unit 0 family ethernet-switching interface-mode access
   user@host#set ge-0/0/1 unit 0 family ethernet-switching vlan members 10
   ```

2. Configure an IP address for the IRB interface.

   ```
   [edit interfaces]
   user@host# set irb unit 10 family inet address 10.10.10.1/24
   ```

3. Configure Layer 2 interfaces.

   ```
   [edit security zones security-zone L2 interfaces]
   user@host# set ge-0/0/1.0
   user@host# set ge-0/0/0.0
   ```

4. Configure bridge domain.

   ```
   [edit vlans vlan-10]
   user@host# set vlan-id 10
   user@host# set l3-interface irb.10
   ```

5. Configure IP addresses for Layer 3 interfaces.

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set ge-0/0/2.0 unit 0 family inet address 192.0.2.1/24
user@host# set ge-0/0/3.0 unit 0 family inet address 192.0.2.3/24
```

6. Configure the policy to permit the traffic.

```
[edit security policies]
user@host#set default-policy permit-all
```

7. Configure Layer 3 interfaces.

```
[edit security zones security-zone]
user@host# set L2 host-inbound-traffic system-services any-service
user@host# set L2 host-inbound-traffic protocols all
user@host# set L3 host-inbound-traffic system-services any-service
user@host# set L3 host-inbound-traffic protocols all
user@host# set L3 interfaces ge-0/0/2.0
user@host# set L3 interfaces ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security policies**, **show vlans**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.0.2.1/24;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.0.2.2/24;
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    irb {
      unit 10 {
        family inet {
          address 10.10.10.1/24;
        }
      }
    }
[edit]
user@host#  show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show vlans
  vlan-10 {
    vlan-id 10;
    l3-interface irb.10;
  }
[edit]
user@host#  show security zones
  security-zone L2 {
    host-inbound-traffic {
      system-services {
        any-service;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/1.0;
      ge-0/0/0.0;
    }
  }
  security-zone L3 {
    host-inbound-traffic {
      system-services {
        any-service;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/2.0;
      ge-0/0/3.0;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying the Layer 2 and Layer 3 Interfaces and Zones on page 2919
- Verifying the Layer 2 and Layer 3 Session on page 2920

Verifying the Layer 2 and Layer 3 Interfaces and Zones

Purpose Verify that the Layer 2 and Layer 3 interfaces and Layer 2 and Layer 3 zones are created.

Action From operational mode, enter the **show security zones** command.

```
user@host>show security zones
 Security zone: HOST
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:

Security zone: L2
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 2
  Interfaces:
    ge-0/0/0.0
    ge-0/0/1.0

Security zone: L3
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 2
  Interfaces:
    ge-0/0/2.0
    ge-0/0/3.0

Security zone: junos-host
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 0
  Interfaces:
```

Meaning The output shows the Layer 2 (L2) and Layer 3 (L3) zone names and the number and names of Layer 2 and Layer 3 interfaces bound to the L2 and L3 zones.

Copyright © 2016, Juniper Networks, Inc.

Verifying the Layer 2 and Layer 3 Session

Purpose Verify that the Layer 2 and Layer 3 sessions are established on the device.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
Session ID: 130000050, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 10.10.10.2/22 --> 10.10.10.3/28;icmp, If: ge-0/0/0.0, Pkts: 1, Bytes: 98
  Out: 10.10.10.3/245 --> 10.10.10.2/248;icmp, If: ge-0/0/1.0, Pkts: 1, Bytes:
98

Session ID: 130000051, Policy name: default-policy-02/2, Timeout: 4, Valid
  In: 192.0.2.1/17 --> 192.0.2.2/19;icmp, If: ge-0/0/2.0, Pkts: 1, Bytes: 84
Out: 192.0.2.2/212 --> 192.0.2.1/218;icmp, If: ge-0/0/3.0, Pkts: 1, Bytes: 84
```

Meaning The output shows active sessions on the device and each session's associated security policy.

- **Session ID 130000050**—Number that identifies the Layer 2 session. Use this ID to get more information about the Layer 2 session such as policy name or number of packets in and out.
- **default-policy-00/2**—Default policy name that permitted the Layer 2 traffic.
- **In**—Incoming flow (source and destination Layer 2 IP addresses with their respective source and destination port numbers, session is ICMP, and the source interface for this session is ge-0/0/0.0).
- **Out**—Reverse flow (source and destination Layer 2 IP addresses with their respective source and destination port numbers, session is ICMP, and destination interface for this session is ge-0/0/1.0).
- **Session ID 130000051**—Number that identifies the Layer 3 session. Use this ID to get more information about the Layer 3 session such as policy name or number of packets in and out.
- **default-policy-02/2**—Default policy name that permitted the Layer 3 traffic.
- **In**—Incoming flow (source and destination Layer 3 IP addresses with their respective source and destination port numbers, session is ICMP, and the source interface for this session is ge-0/0/2.0).
- **Out**—Reverse flow (source and destination Layer 3 IP addresses with their respective source and destination port numbers, session is ICMP, and destination interface for this session is ge-0/0/3.0).

Related Documentation

- Understanding Mixed Mode (Layer 2 and Layer 3) on page 2910
- Understanding Secure Wire on page 2955

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 129

Configuring Security Zones and Security Policies

- Understanding Layer 2 Security Zones on page 2921
- Example: Configuring Layer 2 Security Zones on page 2922
- Understanding Security Policies in Transparent Mode on page 2923
- Example: Configuring Security Policies in Transparent Mode on page 2924
- Understanding Firewall User Authentication in Transparent Mode on page 2926

Understanding Layer 2 Security Zones

A Layer 2 security zone is a zone that hosts Layer 2 interfaces. A security zone can be either a Layer 2 or Layer 3 zone; it can host either all Layer 2 interfaces or all Layer 3 interfaces, but it cannot contain a mix of Layer 2 and Layer 3 interfaces.

The security zone type—Layer 2 or Layer 3—is implicitly set from the first interface configured for the security zone. Subsequent interfaces configured for the same security zone must be the same type as the first interface.

NOTE: **You cannot configure a device with both Layer 2 and Layer 3 security zones.**

You can configure the following properties for Layer 2 security zones:

- Interfaces—List of interfaces in the zone.
- Policies—Active security policies that enforce rules for the transit traffic, in terms of what traffic can pass through the firewall, and the actions that need to take place on the traffic as it passes through the firewall.
- Screens—A Juniper Networks stateful firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another. For every security zone, and the MGT zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.

Copyright © 2016, Juniper Networks, Inc.

NOTE: You can configure the same screen options for a Layer 2 security zone as for a Layer 3 security zone.

- Address books—IP addresses and address sets that make up an address book to identify its members so that you can apply policies to them.
- TCP-RST—When this feature is enabled, the system sends a TCP segment with the reset flag set when traffic arrives that does not match an existing session and does not have the synchronize flag set.

In addition, you can configure a Layer 2 zone for host-inbound traffic. This allows you to specify the kinds of traffic that can reach the device from systems that are directly connected to the interfaces in the zone. You must specify all expected host-inbound traffic because inbound traffic from devices directly connected to the device's interfaces is dropped by default.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Layer 2 Interfaces on page 2906
- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Layer 2 Security Zones on page 2922
- Example: Configuring Layer 2 Logical Interfaces on page 2906

Example: Configuring Layer 2 Security Zones

This example shows how to configure Layer 2 security zones.

- Requirements on page 2922
- Overview on page 2922
- Configuration on page 2923
- Verification on page 2923

Requirements

Before you begin, determine the properties you want to configure for the Layer 2 security zone. See “Understanding Layer 2 Security Zones” on page 2921.

Overview

In this example, you configure security zone l2-zone1 to include a Layer 2 logical interface called ge-3/0/0.0 and security zone l2-zone2 to include a Layer 2 logical interface called ge-3/0/1.0. Then you configure l2-zone2 to allow all supported application services (such as SSH, Telnet, and SNMP) as host-inbound traffic.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Layer 2 security zones:

1. Create a Layer 2 security zone and assign interfaces to it.

```
[edit security zones]
user@host# set security-zone l2-zone1 interfaces ge-3/0/0.0
user@host# set security-zone l2-zone2 interfaces ge-3/0/1.0
```

2. Configure one of the Layer 2 security zones.

```
[edit security zones]
user@host# set security-zone l2–zone2 host-inbound-traffic system-services all
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security zones** command.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Example: Configuring Security Policies in Transparent Mode on page 2924
- Example: Configuring Layer 2 Logical Interfaces on page 2906

Understanding Security Policies in Transparent Mode

In transparent mode, security policies can be configured only between Layer 2 zones. When packets are forwarded through the bridge domain, the security policies are applied between security zones. A security policy for transparent mode is similar to a policy configured for Layer 3 zones, with the following exceptions:

- NAT is not supported.
- IPsec VPN is not supported.
- Application ANY is used.

Layer 2 forwarding does not permit any interzone traffic unless there is a policy explicitly configured on the device. By default, Layer 2 forwarding performs the following actions:

- Allows or denies traffic specified by the configured policy.
- Allows Address Resolution Protocol (ARP) and Layer 2 non-IP multicast and broadcast traffic. The device can receive and pass Layer 2 broadcast traffic for STP.
- Continues to block all non-IP and non-ARP unicast traffic.

Copyright © 2016, Juniper Networks, Inc.

This default behavior can be changed for bridge packet flow by using either J-Web or the CLI configuration editor:

- Configure the **block-non-ip-all** option to block all Layer 2 non-IP and non-ARP traffic, including multicast and broadcast traffic.
- Configure the **bypass-non-ip-unicast** option to allow all Layer 2 non-IP traffic to pass through the device.

NOTE: **You cannot configure both options at the same time.**

In mixed mode (default mode), you can create a separate security zone for Layer 2 and Layer 3 interfaces. However, there is no routing among IRB interfaces and between IRB interfaces and Layer 3 interfaces. Hence, you cannot configure security policies between Layer 2 and Layer 3 zones. You can only configure security policies between the Layer 2 zones or between Layer 3 zones.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Security Policies in Transparent Mode on page 2924
- Example: Configuring Layer 2 Security Zones on page 2922
- Understanding Mixed Mode (Layer 2 and Layer 3) on page 2910

Example: Configuring Security Policies in Transparent Mode

This example shows how to configure security policies in transparent mode between Layer 2 zones.

- Requirements on page 2924
- Overview on page 2924
- Configuration on page 2925
- Verification on page 2926

Requirements

Before you begin, determine the policy behavior you want to include in the Layer 2 security zone. See “Understanding Security Policies in Transparent Mode” on page 2923.

Overview

In this example, you configure a security policy to allow HTTP traffic from the 10.1.1.1/24 subnetwork in the l2–zone1 security zone to the server at 20.1.1.1/32 in the l2–zone2 security zone.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone l2-zone1 to-zone l2-zone2 policy p1 match source-address
  10.1.1.1/24
set security policies from-zone l2-zone1 to-zone l2-zone2 policy p1 match
  destination-address 20.1.1.1/32
set security policies from-zone l2-zone1 to-zone l2-zone2 policy p1 match application http
set security policies from-zone l2-zone1 to-zone l2-zone2 policy p1 then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies in transparent mode:

1. Create policies and assign addresses to the interfaces for the zones.

```
[edit security policies]
user@host# set from-zone l2-zone1 to-zone l2-zone2 policy p1 match source-address
  10.1.1.1/24
user@host# set from-zone l2-zone1 to-zone l2-zone2 policy p1 match
  destination-address 20.1.1.1/32
```

2. Set policies for the application.

```
[edit security policies]
user@host# set from-zone l2-zone1 to-zone l2-zone2 policy p1 match application
  http
user@host# set from-zone l2-zone1 to-zone l2-zone2 policy p1 then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host#show security policies
  from-zone l2-zone1 to-zone l2-zone2
{
  policy p1 {
    match {
    source-address 10.1.1.1/24;
    destination-address 20.1.1.1/32;
    application junos-http;
  }
  then {
    permit;
  }
}
}
```

Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

Verifying Layer 2 Security Policies

Purpose Verify that the Layer 2 security policies are configured properly.

Action From configuration mode, enter the **show security policies** command.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Layer 2 Security Zones on page 2922

Understanding Firewall User Authentication in Transparent Mode

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Firewall user authentication enables administrators to restrict and permit users accessing protected resources behind a firewall based on their source IP address and other credentials. Junos OS supports the following types of firewall user authentication for transparent mode on the SRX Series device:

- Pass-through authentication—A host or a user from one zone tries to access resources on another zone. You must use an FTP, Telnet, or HTTP client to access the IP address of the protected resource and be authenticated by the firewall. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.
- Web authentication—Users try to connect, by using HTTP, to an IP address on the IRB interface that is enabled for Web authentication. You are prompted for the username and password that are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

Related Documentation

- *Authentication and Integrated User Firewalls Feature Guide for Security Devices*
- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Example: Configuring an IRB Interface on page 2909

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 130

Configuring Layer 2 Forwarding Tables

- Understanding Layer 2 Forwarding Tables on page 2927
- Example: Configuring the Default Learning for Unknown MAC Addresses on page 2929

Understanding Layer 2 Forwarding Tables

The SRX Series device maintains forwarding tables that contain MAC addresses and associated interfaces for each Layer 2 bridge domain. When a packet arrives with a new source MAC address in its frame header, the device adds the MAC address to its forwarding table and tracks the interface at which the packet arrived. The table also contains the corresponding interface through which the device can forward traffic for a particular MAC address.

If the destination MAC address of a packet is unknown to the device (that is, the destination MAC address in the packet does not have an entry in the forwarding table), the device duplicates the packet and floods it on all interfaces in the bridge domain other than the interface on which the packet arrived. This is known as *packet flooding* and is the default behavior for the device to determine the outgoing interface for an unknown destination MAC address. Packet flooding is performed at two levels: packets are flooded to different zones as permitted by configured Layer 2 security policies, and packets are also flooded to different interfaces with the same VLAN identifier within the same zone. The device learns the forwarding interface for the MAC address when a reply with that MAC address arrives at one of its interfaces.

You can specify that the SRX Series device use ARP queries and traceroute requests (which are ICMP echo requests with the time-to-live values set to 1) instead of packet flooding to locate an unknown destination MAC address. This method is considered more secure than packet flooding because the device floods ARP queries and traceroute packets—not the initial packet—on all interfaces. When ARP or traceroute flooding is used, the original packet is dropped. The device broadcasts an ARP or ICMP query to all other devices on the same subnetwork, requesting the device at the specified destination IP address to send back a reply. Only the device with the specified IP address replies, which provides the requestor with the MAC address of the responder.

ARP allows the device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnetwork as the ingress IP address. (The ingress IP address refers to the IP address of the last device to send the packet to the device. The device might be the source that sent the packet or a router forwarding the

Copyright © 2016, Juniper Networks, Inc.

packet.) Traceroute allows the device to discover the destination MAC address even if the destination IP address belongs to a device in a subnetwork beyond that of the ingress IP address.

When you enable ARP queries to locate an unknown destination MAC address, traceroute requests are also enabled. You can also optionally specify that traceroute requests not be used; however, the device can then discover destination MAC addresses for unicast packets only if the destination IP address is in the same subnetwork as the ingress IP address.

Whether you enable ARP queries and traceroute requests or ARP-only queries to locate unknown destination MAC addresses, the SRX Series device performs the following series of actions:

1. The device notes the destination MAC address in the initial packet. The device adds the source MAC address and its corresponding interface to its forwarding table, if they are not already there.
2. The device drops the initial packet.
3. The device generates an ARP query packet and optionally a traceroute packet and floods those packets out all interfaces except the interface on which the initial packet arrived.

   ARP packets are sent out with the following field values:

   - Source IP address set to the IP address of the IRB
   - Destination IP address set to the destination IP address of the original packet
   - Source MAC address set to the MAC address of the IRB
   - Destination MAC address set to the broadcast MAC address (all **0xf**)

   Traceroute (ICMP echo request or ping) packets are sent out with the following field values:

   - Source IP address set to the IP address of the original packet
   - Destination IP address set to the destination IP address of the original packet
   - Source MAC address set to the source MAC address of the original packet
   - Destination MAC address set to the destination MAC address of the original packet
   - Time-to-live (TTL) set to **1**
4. Combining the destination MAC address from the initial packet with the interface leading to that MAC address, the device adds a new entry to its forwarding table.
5. The device forwards all subsequent packets it receives for the destination MAC address out the correct interface to the destination.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Example: Configuring an IRB Interface on page 2909

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring the Default Learning for Unknown MAC Addresses on page 2929

Example: Configuring the Default Learning for Unknown MAC Addresses

This example shows how to configure the device to use only ARP requests to learn the outgoing interfaces for unknown destination MAC addresses.

- Requirements on page 2929
- Overview on page 2929
- Configuration on page 2929
- Verification on page 2929

Requirements

Before you begin, determine the MAC addresses and associated interfaces of the forwarding table. See "Understanding Layer 2 Forwarding Tables" on page 2927.

Overview

In this example, you configure the device to use only ARP queries without traceroute requests.

Configuration

Step-by-Step Procedure

To configure the device to use only ARP requests to learn unknown destination MAC addresses:

1. Enable the device.

   ```
   [edit]
   user@host# set security flow ethernet-switching no-packet-flooding no-trace-route
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security flow** command.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Integrated Routing and Bridging Interfaces on page 2907
- Example: Configuring an IRB Interface on page 2909

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 131

Configuring Layer 2 Transparent Mode Chassis Clusters

- Understanding Layer 2 Transparent Mode Chassis Clusters on page 2931
- Example: Configuring Redundant Ethernet Interfaces for Layer 2 Transparent Mode Chassis Clusters on page 2933

Understanding Layer 2 Transparent Mode Chassis Clusters

A pair of SRX Series devices in Layer 2 transparent mode can be connected in a chassis cluster to provide network node redundancy. When configured in a chassis cluster, one node acts as the primary device and the other as the secondary device, ensuring stateful failover of processes and services in the event of system or hardware failure. If the primary device fails, the secondary device takes over processing of traffic.

NOTE: **If the primary device fails in a Layer 2 transparent mode chassis cluster, the physical ports in the failed device become inactive (go down) for a few seconds before they become active (come up) again.**

To form a chassis cluster, a pair of the same kind of supported SRX Series devices combines to act as a single system that enforces the same overall security.

Devices in Layer 2 transparent mode can be deployed in active/backup and active/active chassis cluster configurations.

The following chassis cluster features are not supported for devices in Layer 2 transparent mode:

- Gratuitous ARP—The newly elected master in a redundancy group cannot send gratuitous ARP requests to notify network devices of a change in mastership on the redundant Ethernet interface links.
- IP address monitoring—Failure of an upstream device cannot be detected.

A redundancy group is a construct that includes a collection of objects on both nodes. A redundancy group is primary on one node and backup on the other. When a redundancy group is primary on a node, its objects on that node are active. When a redundancy group fails over, all its objects fail over together.

Copyright © 2016, Juniper Networks, Inc.

You can create one or more redundancy groups numbered 1 through 128 for an active/active chassis cluster configuration. Each redundancy group contains one or more redundant Ethernet interfaces. A redundant Ethernet interface is a pseudointerface that contains physical interfaces from each node of the cluster. The physical interfaces in a redundant Ethernet interface must be the same kind—either Fast Ethernet or Gigabit Ethernet. If a redundancy group is active on node 0, then the child links of all associated redundant Ethernet interfaces on node 0 are active. If the redundancy group fails over to the node 1, then the child links of all redundant Ethernet interfaces on node 1 become active.

NOTE: **In the active/active chassis cluster configuration, the maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces that you configure. In the active/backup chassis cluster configuration, the maximum number of redundancy groups supported is two.**

Configuring redundant Ethernet interfaces on a device in Layer 2 transparent mode is similar to configuring redundant Ethernet interfaces on a device in Layer 3 route mode, with the following difference: the redundant Ethernet interface on a device in Layer 2 transparent mode is configured as a Layer 2 logical interface.

The redundant Ethernet interface may be configured as either an access interface (with a single VLAN ID assigned to untagged packets received on the interface) or as a trunk interface (with a list of VLAN IDs accepted on the interface and, optionally, a native-vlan-id for untagged packets received on the interface). Physical interfaces (one from each node in the chassis cluster) are bound as child interfaces to the parent redundant Ethernet interface.

In Layer 2 transparent mode, MAC learning is based on the redundant Ethernet interface. The MAC table is synchronized across redundant Ethernet interfaces and Services Processing Units (SPUs) between the pair of chassis cluster devices.

The IRB interface is used only for management traffic, and it cannot be assigned to any redundant Ethernet interface or redundancy group.

All Junos OS screen options that are available for a single, nonclustered device are available for devices in Layer 2 transparent mode chassis clusters.

NOTE: **Spanning-tree protocols are not supported for Layer 2 transparent mode. You should ensure that there are no loop connections in the deployment topology.**

Related Documentation

- *Chassis Cluster Feature Guide for Branch SRX Series Devices*
- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Layer 2 Interfaces on page 2906
- Example: Configuring Layer 2 Logical Interfaces on page 2906

Copyright © 2016, Juniper Networks, Inc.

- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Redundant Ethernet Interfaces for Layer 2 Transparent Mode Chassis Clusters on page 2933
- Understanding Layer 2 Forwarding Tables on page 2927

Example: Configuring Redundant Ethernet Interfaces for Layer 2 Transparent Mode Chassis Clusters

This example shows how to configure a redundant Ethernet interface on a device as a Layer 2 logical interface for a Layer 2 transparent mode chassis cluster.

- Requirements on page 2933
- Overview on page 2933
- Configuration on page 2933
- Verification on page 2934

Requirements

Before you begin, determine the devices you want to connect in a chassis cluster. See "Understanding Layer 2 Transparent Mode Chassis Clusters" on page 2931.

Overview

This example shows you how to configure the redundant Ethernet interface as a Layer 2 logical interface and how to bind the physical interfaces (one from each node in the chassis cluster) to the redundant Ethernet interface. In this example, you create redundant Ethernet interface reth0 for redundancy group 1 and configure reth0 as an access interface with the VLAN identifier 1. Then you assign physical interface ge-2/0/2 on a chassis cluster node to the redundant Ethernet interface reth0.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a redundant Ethernet interface as a Layer 2 logical interface:

1. Configure the interfaces and redundancy group.

   ```
   [edit interfaces]
   user@host# set reth0 redundant-ether-options redundancy-group 1
   user@host# set reth0 unit 0 family ethernet-switching interface-mode access
     vlan-id 1
   ```

2. Assign a physical interface on a chassis cluster node.

   ```
   [edit interfaces]
   user@host# set ge-2/0/2 gigether-options redundant-parent reth0
   ```

3. If you are done configuring the device, commit the configuration.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces reth0** and **show interfaces ge-2/0/2** commands.

Related Documentation

- *Chassis Cluster Feature Guide for Branch SRX Series Devices*
- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Understanding Layer 2 Transparent Mode Chassis Clusters on page 2931
- Understanding Layer 2 Forwarding Tables on page 2927

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 132

Configuring IP Spoofing in Layer 2 Transparent Mode

- Understanding IP Spoofing in Layer 2 Transparent Mode on page 2935
- Configuring IP Spoofing in Layer 2 Transparent Mode on page 2936

Understanding IP Spoofing in Layer 2 Transparent Mode

In an IP spoofing attack, the attacker gains access to a restricted area of the network and inserts a false source address in the packet header to make the packet appear to come from a trusted source. IP spoofing is most frequently used in denial-of-service (DoS) attacks. When SRX Series devices are operating in transparent mode, the IP spoof-checking mechanism makes use of address book entries. Address books only exist on the Routing Engine. IP spoofing in Layer 2 transparent mode is performed on the Packet Forwarding Engine. Address book information cannot be obtained from the Routing Engine each time a packet is received by the Packet Forwarding Engine. Therefore, address books attached to the Layer 2 zones must be pushed to the Packet Forwarding Engine.

NOTE: **IP spoofing in Layer 2 transparent mode does not support DNS and wildcard addresses.**

When a packet is received by the Packet Forwarding Engine, the packet's source IP address is checked to determine if it is in the incoming zone's address-book. If the packet's source IP address is in the incoming zone's address book, then this IP address is allowed on the interface, and traffic is passed.

If the source IP address is not present in the incoming zone's address-book, but exists in other zones', then the IP address is considered a spoofed IP. Accordingly, actions such as drop and logging can be taken depending on the screen configuration (alarm-without-drop).

NOTE: **If the alarm-without-drop option is configured, the Layer 2 spoofing packet only triggers an alarm message, but the packet is not dropped.**

Copyright © 2016, Juniper Networks, Inc.

If a packet's source IP address is not present in the incoming zone's address book or other zones', then you cannot determine if the IP is spoofed or not. In such instances, the packet is passed.

Junos OS takes into account the following match conditions while it searches for source IP addresses in the address book:

- **Host-match**—The IP address match found in the address-book is an address without a prefix.
- **Prefix-match**—The IP address match found in the address-book is an address with a prefix.
- **Any-match**—The IP address match found in the address-book is "any", "any-IPv4", or "any-IPv6".
- **No-match**—No IP address match is found.

Related Documentation

- Configuring IP Spoofing in Layer 2 Transparent Mode on page 954

Configuring IP Spoofing in Layer 2 Transparent Mode

You can configure the IP spoof-checking mechanism to determine whether or not an IP is being spoofed.

To configure IP spoofing in Layer 2 transparent mode:

1. Set the interface in Layer 2 transparent mode.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching
   ```

2. (Optional) Set the zone in Layer 2 transparent mode.

   ```
   [edit]
   user@host# set security zones security-zone untrust interfaces ge-0/0/1.0
   ```

3. Configure the address book.

   ```
   [edit]
   user@host# set security address-book my-book address myadd1 10.1.1.0/24
   user@host# set security address-book my-book address myadd2 10.1.2.0/24
   ```

4. Apply the address book to the zone.

   ```
   [edit]
   user@host# set security address-book my-book attach zone untrust
   ```

5. Configure screen IP spoofing.

   ```
   [edit]
   user@host# set security screen ids-option my-screen ip spoofing
   ```

6. Apply the screen to the zone.

   ```
   [edit]
   user@host# set security zones security-zone untrust screen my-screen
   ```


Copyright © 2016, Juniper Networks, Inc.

7. (Optional) Configure the **alarm-without-drop** option.

```
[edit]
user@host# set security screen ids-option my-screen alarm-without-drop
```

NOTE: If the **alarm-without-drop** option is configured, the Layer 2 spoofing packet only triggers an alarm message, but the packet is not dropped.

Related Documentation

- Understanding IP Spoofing in Layer 2 Transparent Mode on page 953

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 133

Configuring Class of Service in Transparent Mode

- Class of Service Functions in Transparent Mode Overview on page 2939
- Understanding BA Traffic Classification on Transparent Mode Devices on page 2940
- Example: Configuring BA Classifiers on Transparent Mode Devices on page 2941
- Understanding Rewrite of Packet Headers on Transparent Mode Devices on page 2943
- Example: Configuring Rewrite Rules on Transparent Mode Devices on page 2944

Class of Service Functions in Transparent Mode Overview

Devices operating in Layer 2 transparent mode support the following class-of-service (CoS) functions:

- IEEE 802.1p behavior aggregate (BA) classifiers to determine the forwarding treatment for packets entering the device

  NOTE: **Only IEEE 802.1p BA classifier types are supported on devices operating in transparent mode.**

- Rewrite rules to redefine IEEE 802.1 CoS values in outgoing packets

  NOTE: **Rewrite rules that redefine IP precedence CoS values and Differentiated Services Code Point (DSCP) CoS values are not supported on devices operating in transparent mode.**

- Shapers to apply rate limiting to an interface
- Schedulers that define the properties of an output queue

You configure BA classifiers and rewrite rules on transparent mode devices in the same way as on devices operating in Layer 3 mode. For transparent mode devices, however, you apply BA classifiers and rewrite rules only to logical interfaces configured with the **family ethernet-switching** configuration statement.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *Class of Service Feature Guide for Security Devices*
- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Understanding BA Traffic Classification on Transparent Mode Devices on page 2940
- Example: Configuring BA Classifiers on Transparent Mode Devices on page 2941

Understanding BA Traffic Classification on Transparent Mode Devices

A BA classifier checks the header information of an ingress packet. The resulting traffic classification consists of a forwarding class (FC) and packet loss priority (PLP). The FC and PLP associated with a packet specify the CoS behavior of a hop within the system. For example, a hop can place a packet into a priority queue according to its FC, and manage queues by checking the packet's PLP. Junos OS supports up to eight FCs and four PLPs.

NOTE: **MPLS EXP bit-based traffic classification is not supported.**

BA classification can be applied within one DiffServ domain. BA classification can also be applied between two domains, where each domain honors the CoS results generated by the other domain. Junos OS performs BA classification for a packet by examining its Layer 2 and Layer 3 CoS-related parameters. Those parameters include the following:

- Layer 2—IEEE 802.1p: User Priority
- Layer 3—IPv4 Precedence, IPv4 DSCP, IPv6 DSCP

On SRX Series devices in transparent mode, a BA classifier evaluates only Layer 2 parameters. On SRX Series devices in Layer 3 mode, a BA classifier can evaluate Layer 2 and Layer 3 parameters; in that case, classification resulting from Layer 3 parameters overrides that of Layer 2 parameters.

On SRX Series devices in transparent mode, you specify one of four PLP levels—high, medium-high, medium-low, or low—when configuring a BA classifier.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Class of Service Functions in Transparent Mode Overview on page 2939
- Example: Configuring BA Classifiers on Transparent Mode Devices on page 2941

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring BA Classifiers on Transparent Mode Devices

This example shows how to configure BA classifiers on transparent mode devices to determine the forwarding treatment of packets entering the devices.

- Requirements on page 2941
- Overview on page 2941
- Configuration on page 2941
- Verification on page 2943

Requirements

Before you begin, configure a Layer 2 logical interface. See "Example: Configuring Layer 2 Logical Interfaces" on page 2906.

Overview

In this example, you configure logical interface ge-0/0/4.0 as a trunk port that carries traffic for packets tagged with VLAN identifiers 200 through 390. You then configure forwarding classes and create BA classifier c1 for IEEE 802.1 traffic where incoming packets with IEEE 802.1p priority bits 110 are assigned to the forwarding class fc1 with a low loss priority. Finally, you apply the BA classifier c1 to interface ge-0/0/4.0.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/4 vlan-tagging unit 0 family ethernet-switching interface-mode
    trunk vlan members 200-390
set class-of-service forwarding-classes queue 0 fc1
set class-of-service forwarding-classes queue 1 fc2
set class-of-service forwarding-classes queue 3 fc4
set class-of-service forwarding-classes queue 4 fc5
set class-of-service forwarding-classes queue 5 fc6
set class-of-service forwarding-classes queue 6 fc7
set class-of-service forwarding-classes queue 7 fc8
set class-of-service forwarding-classes queue 2 fc3
set class-of-service classifiers ieee-802.1 c1 forwarding-class fc1 loss-priority low
    code-point 110
set class-of-service interfaces ge-0/0/4 unit 0 classifiers ieee-802.1 c1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BA classifiers on transparent mode devices:

1. Configure the logical interface as a Layer 2 trunk port.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces ge-0/0/4 vlan-tagging unit 0 family ethernet-switching
  interface-mode trunk vlan members 200-390
```

2. Configure the class of service.

```
[edit]
user@host# edit class-of-service
```

3. Configure the forwarding classes.

```
[edit class-of-service]
user@host# set forwarding-classes queue 0 fc1
user@host# set forwarding-classes queue 1 fc2
user@host# set forwarding-classes queue 3 fc4
user@host# set forwarding-classes queue 4 fc5
user@host# set forwarding-classes queue 5 fc6
user@host# set forwarding-classes queue 6 fc7
user@host# set forwarding-classes queue 7 fc8
user@host# set forwarding-classes queue 2 fc3
```

4. Configure a BA classifier.

```
[edit class-of-service]
user@host# set classifiers ieee-802.1 c1 forwarding-class fc1 loss-priority low
  code-points 110
```

5. Apply the BA classifier to the interface.

```
[edit class-of-service]
user@host# set interfaces ge-0/0/4 unit 0 classifiers ieee-802.1 c1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces ge-0/0/4** and **show class-of-service** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces ge-0/0/4
  vlan-tagging;
  unit 0 {
    family ethernet-switching {
    interface-mode trunk;
    vlan members 200-390;
  }
}
[edit]
  user@host# show class-of-service
  classifiers {
  ieee-802.1 c1 {
    forwarding-class fc1 {
    loss-priority low code-points 110;
    }
  }
  }
  forwarding-classes {
    queue 0 fc1;
    queue 1 fc2;
```

 Copyright © 2016, Juniper Networks, Inc.

```
        queue 3 fc4;
        queue 4 fc5;
        queue 5 fc6;
        queue 6 fc7;
        queue 7 fc8;
        queue 2 fc3;
        }
        interfaces {
          ge-0/0/4 {
        unit 0 {
        classifiers {
          ieee-802.1 c1;
        }
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying BA Classifiers on Transparent Mode Devices on page 2943

Verifying BA Classifiers on Transparent Mode Devices

Purpose Verify that the BA classifier was configured on the transparent mode devices properly.

Action From configuration mode, enter the **show interfaces ge-0/0/4** and **show class-of-service** commands.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Class of Service Functions in Transparent Mode Overview on page 2939
- Understanding BA Traffic Classification on Transparent Mode Devices on page 2940

Understanding Rewrite of Packet Headers on Transparent Mode Devices

Before a packet is transmitted from an interface, the CoS fields in the packet's header can be rewritten for the forwarding class (FC) and packet loss priority (PLP) of the packet. The rewriting function converts a packet's FC and PLP into corresponding CoS fields in the packet header. In Layer 2 transparent mode, the CoS fields are the IEEE 802.1p priority bits.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Example: Configuring Rewrite Rules on Transparent Mode Devices on page 2944

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Rewrite Rules on Transparent Mode Devices

This example shows how to configure rewrite rules on transparent mode devices to redefine IEEE 802.1 CoS values in outgoing packets.

- Requirements on page 2944
- Overview on page 2944
- Configuration on page 2944
- Verification on page 2946

Requirements

Before you begin, configure a Layer 2 logical interface. See "Example: Configuring Layer 2 Logical Interfaces" on page 2906.

Overview

In this example, you configure logical interface ge-1/0/3.0 as a trunk port that carries traffic for packets tagged with VLAN identifiers 200 through 390. You then configure the forwarding classes and create rewrite rule rw1 for IEEE 802.1 traffic. For outgoing packets in the forwarding class fc1 with low loss priority, the IEEE 802.1p priority bits are rewritten as 011. Finally, you apply the rewrite rule rw1 to interface ge-1/0/3.0.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-1/0/3 vlan-tagging unit 0 family ethernet-switching interface-mode
    trunk vlan members 200-390
set class-of-service forwarding-classes queue 0 fc1
set class-of-service forwarding-classes queue 1 fc2
set class-of-service forwarding-classes queue 3 fc4
set class-of-service forwarding-classes queue 4 fc5
set class-of-service forwarding-classes queue 5 fc6
set class-of-service forwarding-classes queue 6 fc7
set class-of-service forwarding-classes queue 7 fc8
set class-of-service forwarding-classes queue 2 fc3
set class-of-service rewrite-rules ieee-802.1 rw1 forwarding-class fc1 loss-priority low
    code-point 011
set class-of-service interfaces ge-1/0/3 unit 0 rewrite-rules ieee-802.1 rw1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure rewrite rules on transparent mode devices:

1. Configure the logical interface as a Layer 2 trunk port.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces ge-1/0/3 vlan-tagging unit 0 family ethernet-switching
  interface-mode trunk vlan members 200-390
```

2. Configure the class of service.

```
[edit]
user@host# edit class-of-service
```

3. Configure the forwarding classes.

```
[edit class-of-service]
user@host# set forwarding-classes queue 0 fc1
user@host# set forwarding-classes queue 1 fc2
user@host# set forwarding-classes queue 3 fc4
user@host# set forwarding-classes queue 4 fc5
user@host# set forwarding-classes queue 5 fc6
user@host# set forwarding-classes queue 6 fc7
user@host# set forwarding-classes queue 7 fc8
user@host# set forwarding-classes queue 2 fc3
```

4. Configure a rewrite rule.

```
[edit class-of-service]
user@host# set rewrite-rules ieee-802.1 rw1 forwarding-class fc1 loss-priority low
  code-point 011
```

5. Apply the rewrite rule to the interface.

```
[edit class-of-service]
user@host# set interfaces ge-1/0/3 unit 0 rewrite-rules ieee-802.1 rw1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces ge-1/0/3** and **show class-of-service** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces ge-1/0/3
  vlan-tagging;
  unit 0 {
    family ethernet-switching {
    interface-mode trunk;
    vlan members 200-390;
  }
}
[edit]
  user@host# show class-of-service
  forwarding-classes {
  queue 0 fc1;
    queue 1 fc2;
    queue 3 fc4;
    queue 4 fc5;
    queue 5 fc6;
    queue 6 fc7;
    queue 7 fc8;
  queue 2 fc3;
  }
  interfaces {
```

Copyright © 2016, Juniper Networks, Inc.

```
        ge-1/0/3 {
        unit 0 {
        rewrite-rules {
          ieee-802.1 rw1;
          }
        }
      }
      }
      rewrite-rules {
        ieee-802.1 rw1 {
        forwarding-class fc1 {
        loss-priority low code-point 011;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying Rewrite Rules on Transparent Mode Devices on page 2946

Verifying Rewrite Rules on Transparent Mode Devices

Purpose Verify that the rewrite rule was configured on the transparent mode devices properly.

Action From configuration mode, enter the **show interfaces ge-1/0/3** and **show class-of-service** commands.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897
- Understanding Transparent Mode Conditions on page 2905
- Understanding Rewrite of Packet Headers on Transparent Mode Devices on page 2943

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 134

Configuring IPv6 Flows

- Understanding IPv6 Flows in Transparent Mode on page 2947
- Enabling Flow-Based Processing for IPv6 Traffic on page 2948
- Example: Configuring Transparent Mode for IPv6 Flows on page 2950

Understanding IPv6 Flows in Transparent Mode

In transparent mode, the SRX Series device filters packets that traverse the device without modifying any of the source or destination information in the packet MAC headers. Transparent mode is useful for protecting servers that mainly receive traffic from untrusted sources because there is no need to reconfigure the IP settings of routers or protected servers.

A device operates in transparent mode when all physical interfaces on the device are configured as Layer 2 interfaces. A physical interface is a Layer 2 interface if its logical interface is configured with the **bridge** option at the [**edit interfaces *interface-name* unit *unit-number* family**] hierarchy level. There is no command to define or enable transparent mode on the device. The device operates in transparent mode when there are interfaces defined as Layer 2 interfaces. The device operates in route mode (the default mode) if all physical interfaces are configured as Layer 3 interfaces.

By default, IPv6 flows are dropped on security devices. To enable processing by security features such as zones, screens, and firewall policies, you must enable flow-based forwarding for IPv6 traffic with the **mode flow-based** configuration option at the [**edit security forwarding-options family inet6**] hierarchy level. You must reboot the device when you change the mode.

In transparent mode, you can configure Layer 2 zones to host Layer 2 interfaces, and you can define security policies between Layer 2 zones. When packets travel between Layer 2 zones, security policies can be enforced on these packets. The following security features are supported for IPv6 traffic in transparent mode:

- Layer 2 security zones and security policies. See "Understanding Layer 2 Security Zones" on page 2921 and "Understanding Security Policies in Transparent Mode" on page 2923.
- Firewall user authentication. See "Understanding Firewall User Authentication in Transparent Mode" on page 2926.

Copyright © 2016, Juniper Networks, Inc.

- Layer 2 transparent mode chassis clusters. See "Understanding Layer 2 Transparent Mode Chassis Clusters" on page 2931.
- Class of service functions. See "Class of Service Functions in Transparent Mode Overview" on page 2939.

The following security features are *not* supported for IPv6 flows in transparent mode:

- Logical systems
- IPv6 GTPv2
- J-Web interface
- NAT
- IPsec VPN
- With the exception of DNS, FTP, and TFTP ALGs, all other ALGs are not supported.

Configuring VLANs and Layer 2 logical interfaces for IPv6 flows is the same as configuring VLANs and Layer 2 logical interfaces for IPv4 flows. You can optionally configure an integrated routing and bridging (IRB) interface for management traffic in a bridge domain. The IRB interface is the only Layer 3 interface allowed in transparent mode. The IRB interface on the SRX Series device does not support traffic forwarding or routing. The IRB interface can be configured with both IPv4 and IPv6 addresses. You can assign an IPv6 address for the IRB interface with the **address** configuration statement at the [**edit interfaces irb unit *number* family inet6**] hierarchy level. You can assign an IPv4 address for the IRB interface with the **address** configuration statement at the [**edit interfaces irb unit *number* family inet**] hierarchy level.

The bridging functions on SRX Series devices are similar to the bridging features on Juniper Networks MX Series routers. However, not all Layer 2 networking features supported on MX Series routers are supported on SRX Series devices. See "Layer 2 Bridging and Transparent Mode Overview" on page 2897.

The SRX Series device maintains forwarding tables that contain MAC addresses and associated interfaces for each Layer 2 bridge domain. The IPv6 flow processing is similar to IPv4 flows. See "Understanding Layer 2 Forwarding Tables" on page 2927.

Related Documentation

- Enabling Flow-Based Processing for IPv6 Traffic on page 1756
- Example: Configuring Transparent Mode for IPv6 Flows on page 2950

Enabling Flow-Based Processing for IPv6 Traffic

By default, the SRX Series device drops IP version 6 (IPv6) traffic. To enable processing by security features such as zones, screens, and firewall policies, you must enable flow-based forwarding for IPv6 traffic.

To enable flow-based forwarding for IPv6 traffic, modify the **mode** statement at the [**edit security forwarding-options family inet6**] hierarchy level:

```
security {
```


Copyright © 2016, Juniper Networks, Inc.

```
forwarding-options {
  family {
    inet6 {
      mode flow-based;
    }
  }
}
}
```

The following example shows the CLI commands you use to configure forwarding for IPv6 traffic.

1. Use the **set** command to change the forwarding option mode for IPv6 to flow-based.

```
[edit]
user@host# set security forwarding-options family inet6 mode flow-based
```

2. Use the **show** command to review your configuration.

```
[edit]
user@host# show security forwarding-options
family {
    inet6 {
        mode flow-based;
    }
}
```

3. Check your changes to the configuration before committing.

```
[edit]
user@host# commit check
warning: You have enabled/disabled inet6 flow.
You must reboot the system for your change to take effect.
If you have deployed a cluster, be sure to reboot all nodes.
configuration check succeeds
```

4. Commit the configuration.

```
[edit]
user@host# commit
warning: You have enabled/disabled inet6 flow.
You must reboot the system for your change to take effect.
If you have deployed a cluster, be sure to reboot all nodes.
commit complete
```

5. At an appropriate time, reboot the device.

NOTE: SRX Series devices only process IPv6 Routing Header 0 (RH0) to-self packets, the segleft field of which is zero. Other packets will be dropped.

Table 172 summarizes device status upon forwarding option configuration change.

Table 302: Device Status Upon Configuration Change

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Drop to flow-based	Yes	Yes	Dropped	Dropped

Copyright © 2016, Juniper Networks, Inc.

Table 302: Device Status Upon Configuration Change *(continued)*

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Drop to packet-based	No	No	Packet-based	Packet-based
Flow-based to packet-based	Yes	Yes	None	Flow sessions created
Flow-based to drop	Yes	Yes	None	Flow sessions created
Packet-based to flow-based	Yes	Yes	Packet-based	Packet-based
Packet-based to drop	No	No	Dropped	Dropped

To process IPv6 traffic, you also need to configure IPv6 addresses for the transit interfaces that receive and forward the traffic. For information on the inet6 protocol family and procedures for configuring IPv6 addresses for interfaces, see the *Interfaces Feature Guide for Security Devices*.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Using Filters to Display IPv6 Session and Flow Information for SRX Series Services Gateways on page 1757

Example: Configuring Transparent Mode for IPv6 Flows

This example shows how to configure VLANs, a Layer 2 interface, and an IRB interface that supports both IPv4 and IPv6 addresses. This example also shows how to configure the device to use only ARP requests to learn the outgoing interfaces for unknown destination MAC addresses.

- Requirements on page 2950
- Overview on page 2950
- Configuration on page 2951
- Verification on page 2953

Requirements

The device must be enabled for IPv6 flow processing. See "Enabling Flow-Based Processing for IPv6 Traffic" on page 1756.

Overview

This example creates the configuration described in Table 303.

Copyright © 2016, Juniper Networks, Inc.

Table 303: IPv6 Transparent Mode Configuration for IPv6 Flows

Feature	Name	Configuration Parameters
VLANs	vlan-a	VLAN 2
	vlan-b	VLAN 10
Logical interface	ge-0/0/0.0	Trunk port for packets tagged with VLAN IDs 1 through 10
Physical interface	ge-0/0/0	VLAN ID 30 assigned to untagged packets
IRB interface	irb.0	Addresses: • IPv4 address 10.1.1.1/24 • IPv6 address 2:10::1/64 Referenced in vlan-b bridge domain
Learn the outgoing interfaces for unknown destination MAC addresses		Use only ARP queries without traceroute requests

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-a vlan-id 2
set vlans vlan2 vlan-id-list 1-10
set interfaces ge-0/0/0 vlan-tagging native-vlan-id 30
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode trunk vlan
  members 1-10
set interfaces irb unit 0 family inet address 10.1.1.1/24
set interfaces irb unit 0 family inet6 address 2:10::1/64
set vlans vlan-b l3-interface irb.0
set security flow ethernet-switching no-packet-flooding no-trace-route
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure transparent mode for IPv6 flows:

1. Configure VLANs.

   ```
   [edit vlans]
   user@host# set vlan-a vlan-id 2
   user@host# set vlans vlan2 vlan-id-list 1-10
   ```

2. Configure the Layer 2 interface.

   ```
   [edit interfaces ge-0/0/0]
   user@host# set vlan-tagging native-vlan-id 30
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set unit 0 family ethernet-switching interface-mode trunk vlan members
  1-10
```

3. Configure the IRB interface.

```
[edit interfaces irb unit 0]
user@host# set family inet address 10.1.1.1/24
user@host# set family inet6 address 2:10::1/64
```

4. Configure the IRB interface for the bridge domain.

```
[edit vlans]
user@host# set vlan-b l3-interface irb.0
```

5. Configure learning for unknown destination MAC addresses.

```
[edit security flow ethernet-switching]
user@host# set no-packet-flooding no-trace-route
```

Results

From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, and **show security flow ethernet-switching** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show vlans
vlan-a {
  vlan-id 2;
}
  vlan-b {
    vlan members 1-10;
    l3-interface irb.0;
  }
user@host# show interfaces
ge-0/0/0 {
  vlan-tagging;
  native-vlan-id 30;
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan members 1-10;
    }
  }
}
user@host# show security flow ethernet-switching
no-packet-flooding {
  no-trace-route;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying IPv6 Sessions on page 2953
- Verifying IPv6 Gates on page 2953
- Verifying IPv6 IP-action Settings on page 2953

Verifying IPv6 Sessions

Purpose Verify IPv6 sessions on the device.

Action From operational mode, enter the **show security flow session family inet6** command.

Verifying IPv6 Gates

Purpose Verify IPv6 gates on the device.

Action From operational mode, enter the **show security flow gate family inet6** command.

Verifying IPv6 IP-action Settings

Purpose Verify IPv6 IP-action settings on the device.

Action From operational mode, enter the **show security flow ip-action family inet6** command.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Understanding IPv6 Flows in Transparent Mode on page 2947

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 135

Configuring Secure Wire

- Understanding Secure Wire on page 2955
- Example: Simplifying SRX Series Device Deployment with Secure Wire over Access Mode Interfaces on page 2957
- Example: Simplifying SRX Series Device Deployment with Secure Wire over Trunk Mode Interfaces on page 2960
- Example: Simplifying SRX Series Device Deployment with Secure Wire over Aggregated Interface Member Links on page 2964
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Redundant Ethernet Interfaces on page 2968
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Aggregated Redundant Ethernet Interfaces on page 2972

Understanding Secure Wire

Traffic that arrives on a specific interface can be forwarded unchanged through another interface. This mapping of interfaces, called secure wire, allows an SRX Series to be deployed in the path of network traffic without requiring a change to routing tables or a reconfiguration of neighboring devices. Figure 125 shows a typical in-path deployment of an SRX Series with secure wire.

Copyright © 2016, Juniper Networks, Inc.

Figure 125: SRX Series In-Path Deployment with Secure Wire

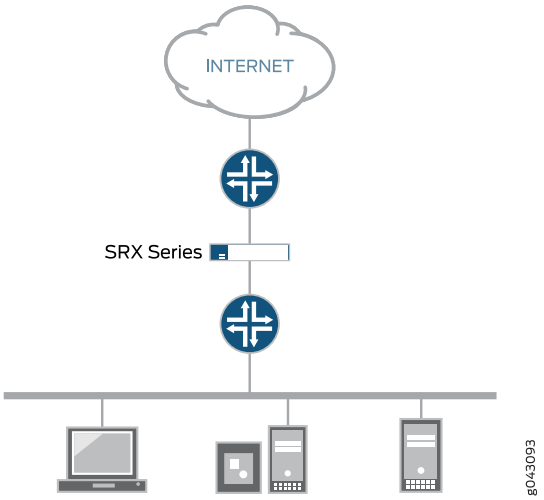

Secure wire maps two peer interfaces. It differs from transparent and route modes in that there is no switching or routing lookup to forward traffic. As long as the traffic is permitted by a security policy, a packet arriving on one peer interface is immediately forwarded unchanged out of the other peer interface. There is no routing or switching decision made on the packet. Return traffic is also forwarded unchanged.

Secure wire mapping is configured with the **secure-wire** statement at the [**edit security forwarding-options**] hierarchy level; two Ethernet logical interfaces must be specified. The Ethernet logical interfaces must be configured with **family ethernet-switching** and each pair of interfaces must belong to the same bridge domain(s). The interfaces must be bound to security zones and a security policy configured to permit traffic between the zones.

This feature is available on Ethernet logical interfaces only; both IPv4 and IPv6 traffic are supported. You can configure interfaces for access or trunk mode. Secure wire supports chassis cluster redundant Ethernet interfaces. This feature does not support security features not supported in transparent mode, including NAT and IPsec VPN. Layer 7 features, including AppSecure, and IPS are supported.

Secure wire is a special case of Layer 2 transparent mode on SRX Series devices that provide point-to-point connections. This means that the two interfaces of a secure wire should ideally be directly connected to Layer 3 entities, such as routers or hosts. Secure wire interfaces can be connected to switches. However, note that a secure wire interface forwards all arriving traffic to the peer interface only if the traffic is permitted by a security policy.

Secure wire can coexist with Layer 3 mode. While you can configure Layer 2 and Layer 3 interfaces at the same time, traffic forwarding occurs independently on Layer 2 and Layer 3 interfaces.

Copyright © 2016, Juniper Networks, Inc.

Secure wire can coexist with Layer 2 transparent mode. If both features exist on the same SRX Series device, you need to configure them in different VLANs.

NOTE: Integrated routing and bridging (IRB) interfaces are not supported with secure wire.

Related Documentation

- Example: Simplifying SRX Series Device Deployment with Secure Wire over Access Mode Interfaces on page 2957
- Example: Simplifying SRX Series Device Deployment with Secure Wire over Trunk Mode Interfaces on page 2960
- Example: Simplifying SRX Series Device Deployment with Secure Wire over Aggregated Interface Member Links on page 2964
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Redundant Ethernet Interfaces on page 2968
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Aggregated Redundant Ethernet Interfaces on page 2972
- Understanding Mixed Mode (Layer 2 and Layer 3) on page 2910

Example: Simplifying SRX Series Device Deployment with Secure Wire over Access Mode Interfaces

If you are connecting an SRX Series device to other network devices, you can use secure wire to simplify the device deployment in the network. No changes to routing or forwarding tables on the SRX Series device and no reconfiguration of neighboring devices is needed. Secure wire allows traffic to be forwarded unchanged between specified access mode interfaces on an SRX Series device as long as it is permitted by security policies or other security features. Follow this example if you are connecting an SRX Series device to other network devices through access mode interfaces.

This example shows how to configure a secure wire mapping for two access mode interfaces. This configuration applies to scenarios where user traffic is not VLAN tagged.

- Requirements on page 2957
- Overview on page 2958
- Configuration on page 2958
- Verification on page 2960

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Copyright © 2016, Juniper Networks, Inc.

Overview

This example configures the secure wire access-sw that maps interface ge-0/0/0.0 to interface ge-0/0/1.0. The two peer interfaces are configured for access mode. The VLAN ID 10 is configured for the bridge domain vlan-10 and the access mode interfaces.

NOTE: A specific VLAN ID must be configured for a bridge domain.

Topology

Figure 126 shows the access mode interfaces that are mapped in secure wire access-sw.

Figure 126: Secure Wire Access Mode Interfaces

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-10 domain-type bridge vlan-id 10
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set security forwarding-options secure-wire access-sw interface [ge-0/0/0.0 ge-0/0/1.0]
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a secure wire mapping for access mode interfaces:

1. Configure the bridge domain.

   ```
   [edit vlans vlan-10]
   user@host# set domain-type bridge vlan-id 10
   ```

2. Configure the access mode interfaces.

   ```
   [edit interfaces ]
   user@host# set ge-0/0/0 unit 0 family ethernet-switching interface-mode access
     vlan-id 10
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ge-0/0/1 unit 0 family ethernet-switching interface-mode access
  vlan-id 10
```

3. Configure the secure wire mapping.

```
[edit security forwarding-options]
user@host# set secure-wire access-sw interface [ge-0/0/0.0 ge-0/0/1.0]
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/0/0.0
user@host# set security-zone untrust interfaces ge-0/0/1.0
```

5. Configure a security policy to permit traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, **show security forwarding-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show vlans
  vlan-10 {
    domain-type bridge;
    vlan-id 10;
  }
  user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  user@host# show security forwarding-options
  secure-wire access-sw {
    interfaces [ge-0/0/0.0 ge-0/0/1.0];
  }
user@host# show security zones
  security-zone trust {
    interfaces {
      ge-0/0/0.0;
    }
  }
  security-zone untrust {
```

Copyright © 2016, Juniper Networks, Inc.

```
    interfaces {
      ge-0/0/1.0;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Secure Wire Mapping on page 2960
- Verifying the Bridge Domain on page 2960

Verifying Secure Wire Mapping

Purpose Verify the secure wire mapping.

Action From operational mode, enter the **show security forwarding-options secure-wire** command.

```
user@host> show security forward-options secure-wire
Secure wire                    Interface       Link   Interface       Link

access-sw                      ge-0/0/0.0      up     ge-0/0/1.0      up
Total secure wires: 1
```

Verifying the Bridge Domain

Purpose Verify the bridge domain.

Action From operational mode, enter the **show vlans vlan-10** command.

```
user@host> show vlans vlan-10
Routing instance        VLAN domain             Tag     Interfaces
default-switch          vlan-10                      10          ge-0/0/0.0
                                                                 ge-0/0/1.0
```

Related Documentation

- Understanding Secure Wire on page 2955

Example: Simplifying SRX Series Device Deployment with Secure Wire over Trunk Mode Interfaces

If you are connecting an SRX Series device to other network devices, you can use secure wire to simplify the device deployment in the network. No changes to routing or forwarding tables on the SRX Series device and no reconfiguration of neighboring devices is needed. Secure wire allows traffic to be forwarded unchanged between specified trunk mode interfaces on an SRX Series device as long as it is permitted by security policies or other security features. Follow this example if you are connecting an SRX Series device to other network devices through trunk mode interfaces.

- Requirements on page 2961
- Overview on page 2961

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 2961
- Verification on page 2963

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

This example configures the secure wire trunk-sw that maps interface ge-0/1/0.0 to interface ge-0/1/1.0. The two peer interfaces are configured for trunk mode and carry user traffic tagged with VLAN IDs from 100 to 102. The VLAN ID list 100-102 is configured for the bridge domain vlan-100 and the trunk mode interfaces.

NOTE: A specific VLAN ID must be configured for a bridge domain.

Topology

Figure 127 shows the trunk mode interfaces that are mapped in secure wire trunk-sw.

Figure 127: Secure Wire Trunk Mode Interfaces

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-100 vlan members 100-102
set interfaces ge-0/1/0 unit 0 family ethernet-switching interface-mode trunk vlan
  members 100-102
set interfaces ge-0/1/1 unit 0 family ethernet-switching interface-mode trunk vlan
  members 100-102
set security forwarding-options secure-wire trunk-sw interface [ge-0/1/0.0 ge-0/1/1.0]
set security zones security-zone trust interfaces ge-0/1/0.0
set security zones security-zone untrust interfaces ge-0/1/1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a secure wire mapping for trunk mode interfaces:

1. Configure the bridge domain.

Copyright © 2016, Juniper Networks, Inc.

```
[edit vlans vlan-100]
user@host# set vlan members 100-102
```

2. Configure the trunk mode interfaces.

```
[edit interfaces]
user@host# set ge-0/1/0 unit 0 family ethernet-switching interface-mode trunk
  vlan members 100-102
user@host# set ge-0/1/1 unit 0 family ethernet-switching interface-mode trunk
  vlan members 100-102
```

3. Configure the secure wire mapping.

```
[edit security forwarding-options]
user@host# set secure-wire trunk-sw interface [ge-0/1/0.0 ge-0/1/1.0]
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/1/0.0
user@host# set security-zone untrust interfaces ge-0/1/1.0
```

5. Configure a security policy to permit traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, **show security forwarding-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show vlans
  vlan-100 {
    vlan members 100-102;
  }
  user@host# show interfaces
  ge-0/1/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan members 100-102;
      }
    }
  }
  ge-0/1/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan members 100-102;
      }
    }
  }
  user@host# show security forwarding-options
  secure-wire trunk-sw {
    interfaces [ge-0/1/0.0 ge-0/1/1.0];
  }
user@host# show security zones
```

 Copyright © 2016, Juniper Networks, Inc.

```
security-zone trust {
  interfaces {
    ge-0/1/0.0;
  }
}
security-zone untrust {
  interfaces {
    ge-0/1/1.0;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Secure Wire Mapping on page 2963
- Verifying the Bridge Domain on page 2963

Verifying Secure Wire Mapping

Purpose Verify the secure wire mapping.

Action From operational mode, enter the **show security forwarding-options secure-wire** command.

```
user@host> show security forward-options secure-wire
Secure wire                    Interface       Link   Interface       Link

trunk-sw                       ge-0/1/0.0      up     ge-0/1/1.0      up
Total secure wires: 1
```

Verifying the Bridge Domain

Purpose Verify the bridge domain.

Action From operational mode, enter the **show vlans** command.

```
user@host> show vlans
Routing instance        VLAN name            VLAN ID     Interfaces
default-switch          vlan-100-vlan-0100         100         ge-0/1/0.0
                                                               ge-0/1/1.0
default-switch          vlan-100-vlan-0101         101         ge-0/1/0.0
                                                               ge-0/1/1.0
default-switch          vlan-100-vlan-0102         102         ge-0/1/0.0
                                                               ge-0/1/1.0
```

NOTE: VLANs are automatically expanded, with one bridge domain for each VLAN ID in the VLAN ID list.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Secure Wire on page 2955

Example: Simplifying SRX Series Device Deployment with Secure Wire over Aggregated Interface Member Links

If you are connecting an SRX Series device to other network devices, you can use secure wire to simplify the device deployment in the network. No changes to routing or forwarding tables on the SRX Series device and no reconfiguration of neighboring devices is needed. Secure wire allows traffic to be forwarded unchanged between specified aggregated interface member links on an SRX Series device as long as it is permitted by security policies or other security features. Follow this example if you are connecting an SRX Series device to other network devices through aggregated interface member links.

NOTE: LACP is not supported. Secure wire mappings can be configured for member links of link bundles instead of directly mapping aggregated Ethernet interfaces.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

This example configures secure wires for two aggregated Ethernet interface link bundles with two links each. Two separate secure wires ae-link1 and ae-link2 are configured using one link from each aggregated Ethernet link bundle. This static mapping requires that the two link bundles have the same number of links.

For link bundles, all logical interfaces of the secure wire mappings must belong to the same bridge domain. VLAN ID 10 is configured for the bridge domain vlan-10 and the logical interfaces. All logical interfaces of a link bundle must belong to the same security zone.

NOTE: A specific VLAN ID or VLAN ID list must be configured for a bridge domain.

Topology

Figure 128 shows the aggregated interfaces that are mapped in secure wire configurations.

Copyright © 2016, Juniper Networks, Inc.

Figure 128: Secure Wire Aggregated Interfaces

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-10 domain-type bridge vlan-id 10
set interfaces ge-0/0/0 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set interfaces ge-0/1/0 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set interfaces ge-0/1/1 unit 0 family ethernet-switching interface-mode access vlan-id
  10
set security forwarding-options secure-wire ae-link1-sw interface [ge-0/1/0.0 ge-0/1/1.0]
set security forwarding-options secure-wire ae-link2-sw interface [ge-0/0/0.0 ge-0/0/1.0]
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone trust interfaces ge-0/1/0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces ge-0/1/1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a secure wire mapping for aggregated interface member links:

1. Configure the bridge domain.

   ```
   [edit vlans vlan-10]
   user@host# set domain-type bridge vlan-id10
   ```

2. Configure the interfaces.

   ```
   [edit interfaces ]
   user@host# set ge-0/0/0 unit 0 family ethernet-switching interface-mode access
     vlan-id 10
   user@host# set ge-0/0/1 unit 0 family ethernet-switching interface-mode access
     vlan-id 10
   user@host# set ge-0/1/0 unit 0 family ethernet-switching interface-mode access
     vlan-id 10
   user@host# set ge-0/1/1 unit 0 family ethernet-switching interface-mode access
     vlan-id 10
   ```

3. Configure the secure wire mappings.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security forwarding-options]
user@host# set secure-wire ae-link1-sw interface [ ge-0/1/0.0 ge-0/1/1.0 ]
user@host# set secure-wire ae-link2-sw interface [ ge-0/0/0.0 ge-0/0/1.0 ]
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces ge-0/0/0.0
user@host# set security-zone trust interfaces ge-0/1/0.0
user@host# set security-zone untrust interfaces ge-0/0/1.0
user@host# set security-zone untrust interfaces ge-0/1/1.0
```

5. Configure a security policy to permit traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, **show security forwarding-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show vlans
  vlan-10 {
    domain-type bridge;
    vlan-id 10;
  }
  user@host# show interfaces
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/1/0 {
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  ge-0/1/1{
    unit 0 {
      family ethernet-switching {
        interface-mode access;
```


Copyright © 2016, Juniper Networks, Inc.

```
            vlan-id 10;
          }
        }
      }
      user@host# show security forwarding-options
      secure-wire ae-link1-sw {
        interfaces [ge-0/1/0.0 ge-0/1/1.0];
      }
      secure-wire ae-link2-sw {
        interfaces [ge-0/0/0.0 ge-0/0/1.0];
      }
    user@host# show security zones
      security-zone trust {
        interfaces {
          ge-0/0/0.0;
          ge-0/1/0.0;
        }
      }
      security-zone untrust {
        interfaces {
          ge-0/0/1.0;
          ge-0/1/1.0;
        }
      }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Secure Wire Mapping on page 2967
- Verifying the Bridge Domain on page 2967

Verifying Secure Wire Mapping

Purpose Verify the secure wire mapping.

Action From operational mode, enter the **show security forwarding-options secure-wire** command.

```
user@host> show security forward-options secure-wire
Secure wire                    Interface       Link   Interface       Link

ae-link1-sw                    ge-0/1/0.0      up     ge-0/1/1.0      up
ae-link2-sw                    ge-0/0/0.0      up     ge-0/0/1.0      up
Total secure wires: 2
```

Verifying the Bridge Domain

Purpose Verify the bridge domain.

Action From operational mode, enter the **show vlans vlan-10** command.

```
user@host> show vlans vlan-10
```

Copyright © 2016, Juniper Networks, Inc.

```
Routing instance        VLAN name            VLAN ID     Interfaces
default-switch          vlan-10                   10          ge-0/0/0.0
                                                              ge-0/0/1.0
                                                              ge-0/1/0.0
                                                              ge-0/1/1.0
```

Related Documentation

- Understanding Secure Wire on page 2955

Example: Simplifying Chassis Cluster Deployment with Secure Wire over Redundant Ethernet Interfaces

If you are connecting an SRX Series chassis cluster to other network devices, you can use secure wire to simplify the cluster deployment in the network. No changes to routing or forwarding tables on the cluster and no reconfiguration of neighboring devices is needed. Secure wire allows traffic to be forwarded unchanged between specified redundant Ethernet interfaces on the SRX Series chassis cluster as long as it is permitted by security policies or other security features. Follow this example if you are connecting an SRX Series chassis cluster to other network devices through redundant Ethernet interfaces.

- Requirements on page 2968
- Overview on page 2968
- Configuration on page 2969
- Verification on page 2972

Requirements

Before you begin:

- Connect a pair of the same SRX Series devices in a chassis cluster.
- Configure the chassis cluster node ID and cluster ID.
- Set the number of redundant Ethernet interfaces in the chassis cluster.
- Configure the chassis cluster fabric.
- Configure chassis cluster redundancy group (in this example redundancy group 1 is used).

For more information, see the *Chassis Cluster Feature Guide for Branch SRX Series Devices*.

Overview

Secure wire is supported over redundant Ethernet interfaces in a chassis cluster. The two redundant Ethernet interfaces must be configured in the same redundancy group. If failover occurs, both redundant Ethernet interfaces should fail over together.

NOTE: Secure wire mapping of redundant Ethernet link aggregation groups (LAGs) are not supported. LACP is not supported.

Copyright © 2016, Juniper Networks, Inc.

This example configures the secure wire reth-sw that maps ingress interface reth0.0 to egress interface reth1.0. Each redundant Ethernet interface consists of two child interfaces, one on each node of the chassis cluster. The two redundant Ethernet interfaces are configured for access mode. VLAN ID 10 is configured for the bridge domain vlan-10 and the redundant Ethernet interfaces.

NOTE: A specific VLAN ID or VLAN ID list must be configured for a bridge domain.

Topology

Figure 129 shows the redundant Ethernet interfaces that are mapped in secure wire reth-sw.

Figure 129: Secure Wire Redundant Ethernet Interfaces

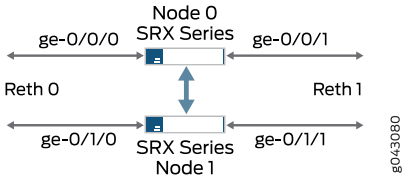

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-10 domain-type bridge vlan-id 10
set interfaces ge-0/0/0 gigether-options redundant-parent reth0
set interfaces ge-0/0/1 gigether-options redundant-parent reth1
set interfaces ge-0/1/0 gigether-options redundant-parent reth0
set interfaces ge-0/1/1 gigether-options redundant-parent reth1
set interfaces reth0 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth1 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth1 redundant-ether-options redundancy-group 1
set security forwarding-options secure-wire reth-sw interface [reth0.0 reth1.0]
set security zones security-zone trust interfaces reth0.0
set security zones security-zone untrust interfaces reth1.0
set security policies default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a secure wire mapping for chassis cluster redundant Ethernet interfaces:

1. Configure the bridge domain.

```
[edit vlans vlan-10]
user@host# set domain-type bridge vlan-id 10
```

2. Configure the redundant Ethernet interfaces.

```
[edit interfaces ]
user@host# set ge-0/0/0 gigether-options redundant-parent reth0
user@host# set ge-0/0/1 gigether-options redundant-parent reth1
user@host# set ge-0/1/0 gigether-options redundant-parent reth0
user@host# set ge-0/1/1 gigether-options redundant-parent reth1

user@host#set reth0 unit 0 family ethernet-switching interface-mode access vlan-id
  10
user@host#set reth1 unit 0 family ethernet-switching interface-mode access vlan-id
  10

user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth1 redundant-ether-options redundancy-group 1
```

3. Configure the secure wire mapping.

```
[edit security forwarding-options]
user@host# set secure-wire reth-sw interface [reth0.0 reth1.0]
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces reth0.0
user@host# set security-zone untrust interfaces reth1.0
```

5. Configure a security policy to permit traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results

From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, **show security forwarding-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show vlans
  vlan-10 {
    domain-type bridge;
    vlan-id 10;
  }
  user@host# show interfaces
  ge-0/0/0 {
    gigether-options {
      redundant-parent reth0;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  ge-0/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-0/1/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-0/1/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  reth0 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  reth1 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family ethernet-switching {
        interface-mode access;
        vlan-id 10;
      }
    }
  }
  user@host# show security forwarding-options
  secure-wire reth-sw {
    interfaces [reth0.0 reth1.0];
  }
user@host# show security zones
  security-zone trust {
    interfaces {
      reth0.0;
    }
  }
  security-zone untrust {
    interfaces {
      reth1.0;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Secure Wire Mapping on page 2972
- Verifying the Bridge Domain on page 2972

Verifying Secure Wire Mapping

Purpose Verify the secure wire mapping.

Action From operational mode, enter the **show security forwarding-options secure-wire** command.

```
user@host> show security forward-options secure-wire
node0:
--------------------------------------------------------------------------
Secure wire                       Interface       Link   Interface       Link

reth-sw                           reth0.0         up     reth1.0         up

Total secure wires: 1

node1:
--------------------------------------------------------------------------
Secure wire                       Interface       Link   Interface       Link

reth-sw                           reth0.0         up     reth1.0         up

Total secure wires: 1
```

Verifying the Bridge Domain

Purpose Verify the bridge domain.

Action From operational mode, enter the **show vlan vlan-10** command.

```
user@host> show vlan vlan-10
Routing instance        VLAN Name             VLAN ID     Interfaces
default-switch          vlan-10                    10         reth0.0
                                                              reth1.0
```

Related Documentation

- Understanding Secure Wire on page 2955
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Aggregated Redundant Ethernet Interfaces on page 2972

Example: Simplifying Chassis Cluster Deployment with Secure Wire over Aggregated Redundant Ethernet Interfaces

If you are connecting an SRX Series chassis cluster to other network devices, you can use secure wire to simplify the cluster deployment in the network. No changes to routing or forwarding tables on the cluster and no reconfiguration of neighboring devices is needed. Secure wire allows traffic to be forwarded unchanged between specified

Copyright © 2016, Juniper Networks, Inc.

redundant Ethernet interfaces on the SRX Series chassis cluster as long as it is permitted by security policies or other security features. Follow this example if you are connecting an SRX Series chassis cluster to other network devices through aggregated redundant Ethernet interfaces.

NOTE: Secure wires cannot be configured for redundant Ethernet interface link aggregation groups (LAGs). For the secure wire mapping shown in this example, there is no LAG configuration on the SRX Series chassis cluster. Each redundant Ethernet interface consists of two child interfaces, one on each node of the chassis cluster. Users on upstream or downstream devices connected to the SRX Series cluster can configure the redundant Ethernet interface child links in LAGs.

- Requirements on page 2973
- Overview on page 2973
- Configuration on page 2974
- Verification on page 2978

Requirements

Before you begin:

- Connect a pair of the same SRX Series devices in a chassis cluster.
- Configure the chassis cluster node ID and cluster ID.
- Set the number of redundant Ethernet interfaces in the chassis cluster.
- Configure the chassis cluster fabric.
- Configure the chassis cluster redundancy group (in this example, redundancy group 1 is used).

For more information, see the *Chassis Cluster Feature Guide for Branch SRX Series Devices*.

Overview

This example configures secure wires for four redundant Ethernet interfaces: reth0, reth1, reth2, and reth3. Each redundant Ethernet interface consists of two child interfaces, one on each node of the chassis cluster. All four redundant Ethernet interfaces must be in the same bridge domain—in this example, the bridge domain is vlan-0. Two of the redundant Ethernet interfaces, reth0.0 and reth2.0, are assigned to the trust zone, while the other two interfaces, reth1.0 and reth3.0, are assigned to the untrust zone.

This example configures the following secure wires:

- reth-sw1 maps interface reth0.0 to interface reth1.0
- reth-sw2 maps interface reth2.0 to reth3.0

All redundant Ethernet interfaces are configured for access mode. VLAN ID 10 is configured for the bridge domain vlan-0 and the redundant Ethernet interfaces.

Copyright © 2016, Juniper Networks, Inc.

NOTE: A specific VLAN ID or VLAN ID list must be configured for a bridge domain.

Topology

Figure 130 shows the redundant Ethernet interface child links that are mapped in secure wire configurations reth-sw1 and reth-sw2. Each redundant Ethernet interface consists of two child interfaces, one on each node of the chassis cluster.

Figure 130: Secure Wire Redundant Ethernet Interface Child Links

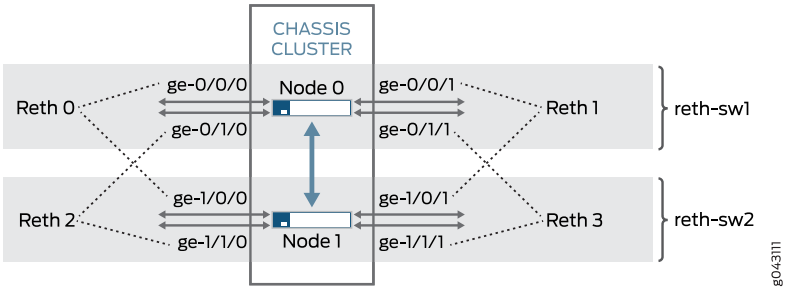

Users on upstream or downstream devices connected to the SRX Series cluster can configure redundant Ethernet interface child links in a LAG as long as the LAG does not span chassis cluster nodes. For example, ge-0/0/0 and ge-0/1/0 and ge-0/0/1 and ge-0/1/1 on node 0 can be configured as LAGs on connected devices. In the same way, ge-1/0/0 and ge-1/1/0 and ge-1/0/1 and ge-1/1/1 on node 1 can be configured as LAGs on connected devices.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set vlans vlan-0 domain-type bridge vlan-id 10
set interfaces ge-0/0/0 gigether-options redundant-parent reth0
set interfaces ge-0/0/1 gigether-options redundant-parent reth1
set interfaces ge-0/1/0 gigether-options redundant-parent reth2
set interfaces ge-0/1/1 gigether-options redundant-parent reth3
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/1/0 gigether-options redundant-parent reth2
set interfaces ge-1/1/1 gigether-options redundant-parent reth3
set interfaces reth0 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth1 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth2 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth3 unit 0 family ethernet-switching interface-mode access vlan-id 10
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth1 redundant-ether-options redundancy-group 1
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
set security forwarding-options secure-wire reth-sw1 interface [reth0.0 reth1.0]
set security forwarding-options secure-wire reth-sw2 interface [reth2.0 reth3.0]
set security zones security-zone trust interfaces reth0.0
set security zones security-zone trust interfaces reth2.0
set security zones security-zone untrust interfaces reth1.0
set security zones security-zone untrust interfaces reth3.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a secure wire mapping for aggregated interface member links:

1. Configure the bridge domain.

```
[edit vlans vlan-0]
user@host# set domain-type bridge vlan-id 10
```

2. Configure the redundant Ethernet interfaces.

```
[edit interfaces ]
user@host# set ge-0/0/0 gigether-options redundant-parent reth0
user@host# set ge-0/0/1 gigether-options redundant-parent reth1
user@host# set ge-0/1/0 gigether-options redundant-parent reth2
user@host# set ge-0/1/1 gigether-options redundant-parent reth3
user@host# set ge-1/0/0 gigether-options redundant-parent reth0
user@host# set ge-1/0/1 gigether-options redundant-parent reth1
user@host# set ge-1/1/0 gigether-options redundant-parent reth2
user@host# set ge-1/1/1 gigether-options redundant-parent reth3

user@host# set reth0 unit 0 family ethernet-switching interface-mode access
  vlan-id 10
user@host# set reth1 unit 0 family ethernet-switching interface-mode access vlan-id
  10
user@host# set reth2 unit 0 family ethernet-switching interface-mode access vlan-id
  10
user@host# set reth3 unit 0 family ethernet-switching interface-mode access vlan-id
  10

user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

3. Configure the secure wire mappings.

```
[edit security forwarding-options]
user@host# set secure-wire reth-sw1 interface [reth0.0 reth1.0]
user@host# set secure-wire reth-sw2 interface [reth2.0 reth3.0]
```

4. Configure security zones.

```
[edit security zones]
user@host# set security-zone trust interfaces reth0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security-zone trust interfaces reth2.0

user@host# set security-zone untrust interfaces reth1.0
user@host# set security-zone untrust interfaces reth3.0
```

5. Configure a security policy to permit traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show vlans**, **show interfaces**, **show security forwarding-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show vlans
  vlan-0 {
    domain-type bridge;
    vlan-id 10;
  }
  user@host# show interfaces
  ge-0/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-0/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-0/1/0 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-0/1/1 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  ge-1/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-1/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-1/1/0 {
    gigether-options {
      redundant-parent reth2;
    }
```

 Copyright © 2016, Juniper Networks, Inc.

```
}
ge-1/1/1 {
  gigether-options {
    redundant-parent reth3;
  }
}
reth0 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family ethernet-switching {
      interface-mode access;
      vlan-id 10;
    }
  }
}
reth1 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family ethernet-switching {
      interface-mode access;
      vlan-id 10;
    }
  }
}
reth2 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family ethernet-switching {
      interface-mode access;
      vlan-id 10;
    }
  }
}
reth3 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family ethernet-switching {
      interface-mode access;
      vlan-id 10;
    }
  }
}
user@host# show security forwarding-options
secure-wire reth-sw1 {
  interfaces [reth0.0 reth1.0];
}
secure-wire reth-sw2 {
  interfaces [reth2.0 reth3.0];
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  user@host# show security zones
    security-zone trust {
      interfaces {
        reth0.0;
        reth2.0;
      }
    }
    security-zone untrust {
      interfaces {
        reth1.0;
        reth3.0;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Secure Wire Mapping on page 2978
- Verifying Bridge Domain on page 2978

Verifying Secure Wire Mapping

Purpose Verify the secure wire mapping.

Action From operational mode, enter the **show security forwarding-options secure-wire** command.

```
user@host> show security forward-options secure-wire
node0:
--------------------------------------------------------------------------
Secure wire                       Interface       Link   Interface       Link

reth-sw1                          reth0.0         up     reth1.0         up
reth-sw2                          reth2.0         up     reth3.0         up

Total secure wires: 2

node1:
--------------------------------------------------------------------------
Secure wire                       Interface       Link   Interface       Link

reth-sw1                          reth0.0         up     reth1.0         up
reth-sw2                          reth2.0         up     reth3.0         up

Total secure wires: 2
```

Verifying Bridge Domain

Purpose Verify the bridge domain.

Action From operational mode, enter the **show vlans vlan-0** command.

```
user@host> show vlans vlan-0
```


Copyright © 2016, Juniper Networks, Inc.

```
Routing instance        VLAN name            VLAN ID     Interfaces
default-switch          vlan-0                    10
                                                          reth0.0
                                                          reth1.0
                                                          reth2.0
                                                          reth3.0
```

Related Documentation

- Understanding Secure Wire on page 2955
- Example: Simplifying Chassis Cluster Deployment with Secure Wire over Redundant Ethernet Interfaces on page 2968

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 42

Configuring Layer 2 Next generation Switching Mode

- Configuring Switching Modes on page 2983
- Configuring VLANs on page 2989
- Configuring Link Aggregation Control Protocol on page 2993
- Configuring Port Security on page 2999

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 136

Configuring Switching Modes

- Understanding Switching Modes on page 2983
- Ethernet Ports Switching Overview on page 2984
- Example: Configuring Switching Modes on page 2985
- Verifying Switching Mode Configuration on page 2986

Understanding Switching Modes

There are two types of switching modes:

- Switching Mode–The uPIM appears in the list of interfaces as a single interface, which is the first interface on the uPIM. For example, ge-2/0/0. You can optionally configure each uPIM port only for autonegotiation, speed, and duplex mode. A uPIM in switching mode can perform the following functions:
  - Layer 3 forwarding—Routes traffic destined for WAN interfaces and other PIMs present on the chassis.
  - Layer 2 forwarding—Switches intra-LAN traffic from one host on the LAN to another LAN host (one port of uPIM to another port of same uPIM).
- Enhanced Switching Mode–Each port can be configured for switching or routing mode. This usage differs from the routing and switching modes, in which all ports must be in either switching or routing mode. The uPIM in enhanced switching mode provides the following features:
  - Supports configuration of different types of VLANs and inter-VLAN routing.
  - Supports Layer 2 control plane protocols such as Link Aggregation Control Protocol (LACP).

You can set a multiport Gigabit Ethernet uPIM on a device to either switching or enhanced switching mode.

When you set a multiport uPIM to switching mode, the uPIM appears as a single entity for monitoring purposes. The only physical port settings that you can configure are autonegotiation, speed, and duplex mode on each uPIM port, and these settings are optional.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Example: Configuring Switching Modes on page 2985
- Ethernet Ports Switching Overview on page 2984

Ethernet Ports Switching Overview

Certain ports on Juniper Networks devices can function as Ethernet access switches that switch traffic at Layer 2 and route traffic at Layer 3.

You can deploy supported devices in branch offices as an access or desktop switch with integrated routing capability, thus eliminating intermediate access switch devices from your network topology. The Ethernet ports provide switching while the Routing Engine provides routing functionality, enabling you to use a single device to provide routing, access switching, and WAN interfaces.

This topic contains the following sections:

- Supported Devices and Ports on page 2984
- Integrated Bridging and Routing on page 2984
- Types of Switch Ports on page 2985
- uPIM in a Daisy Chain on page 2985

Supported Devices and Ports

Juniper Networks supports switching features on the onboard Ethernet ports (Gigabit Ethernet built-in ports) and devices shown in Table 304.

Table 304: Supported Devices and Ports for Switching Features

Device	Ports
SRX300	Onboard Gigabit Ethernet ports (ge-0/0/0 through ge-0/0/5) and 2-gigabit SFP ports
SRX320 and SRX320 PoE	Onboard Gigabit Ethernet ports (ge-0/0/0 through ge-0/0/5) and 2-gigabit SFP ports
SRX340	Onboard Gigabit Ethernet ports (ge-0/0/0 through ge-0/0/7) and 8-gigabit SFP ports
SRX345	Onboard Gigabit Ethernet ports (ge-0/0/0 through ge-0/0/7) and 8-gigabit SFP ports
SRX550HM	Onboard Gigabit Ethernet ports (ge-0/0/0 through ge-0/0/5) and 4-gigabit SPF ports

Integrated Bridging and Routing

Integrated bridging and routing (IRB) provides support for simultaneous Layer 2 bridging and Layer 3 routing within the same bridge domain. Packets arriving on an interface of the bridge domain are switched or routed based on the destination MAC address of the

Copyright © 2016, Juniper Networks, Inc.

packet. Packets with the router's MAC address as the destination are routed to other Layer 3 interfaces.

Types of Switch Ports

The ports, or interfaces, on a switch operate in either access mode or trunk mode.

An interface in access mode connects to a network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The interface itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames.

Trunk interfaces handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to one another.

uPIM in a Daisy Chain

You cannot combine multiple uPIMs to act as a single integrated switch. However, you can connect uPIMs on the same chassis externally by physically connecting a port on one uPIM to a port on another uPIM in a daisy-chain fashion.

Two or more uPIMs daisy-chained together create a single switch with a higher port count than either individual uPIM. One port on each uPIM is used solely for the connection. For example, if you daisy-chain a 6-port uPIM and an 8-port uPIM, the result operates as a 12-port uPIM. Any port of a uPIM can be used for daisy chaining.

Configure the IP address for only one of the daisy-chained uPIMs, making it the primary uPIM. The secondary uPIM routes traffic to the primary uPIM, which forwards it to the Routing Engine. This results in some increase in latency and packet drops due to oversubscription of the external link.

Only one link between the two uPIMs is supported. Connecting more than one link between uPIMs creates a loop topology, which is not supported.

NOTE: **uPIM is currently supported only on the SRX550HM device.**

Related Documentation

- Understanding Switching Modes on page 2983

Example: Configuring Switching Modes

This example shows how to configure a multiport Gigabit Ethernet uPIM to function in switching mode so the uPIM appears as a single entity for monitoring purposes.

- Requirements on page 2986
- Overview on page 2986
- Configuration on page 2986
- Verification on page 2986

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, see "Understanding Switching Modes" on page 2983.

Overview

In this example, you configure **chassis** and set the l2-learning protocol to global mode switching. You then set a physical port parameter on the l2-learning protocols.

Configuration

Step-by-Step Procedure

To configure a uPIM to function in switching mode:

1. Set l2-learning protocol to global mode switching.

   ```
   [edit protocols l2-learning]
   user@host# set protocols l2-learning global-mode switching
   ```

2. Set a physical port parameter on the l2-learning protocols.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching interface-mode
     access
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces ge-0/0/1 switch-options** and **show protocols l2-learning** commands.

Related Documentation

- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Verifying Switching Mode Configuration

Purpose

The operational mode command for checking the status and statistics for multiport uPIMs in switching mode is different from that in routing mode. For uPIMs in routing mode, the operational commands are the same as for other Gigabit Ethernet interfaces, such as the 1-port Gigabit Ethernet ePIM and built-in Gigabit Ethernet ports.

However, not all operational mode commands are supported for ports of a uPIM in switching mode. For example, the operational mode command for monitoring port statistics is not supported.

NOTE: To clear the statistics for the individual switch ports, use the **clear interfaces statistics ge-*pim*/0/0 switch-port *port-number*** command.

Copyright © 2016, Juniper Networks, Inc.

To verify the status and view statistics for a port on a uPIM in switching mode:

user@host# show interfaces ge-*slot*/0/0 switch-port *port-number*

```
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                          Receive         Transmit
          Total bytes                    28437086         21792250
          Total packets                    409145            88008
          Unicast packets                    9987            83817
          Multicast packets                145002                0
          Broadcast packets                254156             4191
          Multiple collisions                  23               10
          FIFO/CRC/Align errors                 0                0
          MAC pause frames                      0                0
          Oversized frames                      0
          Runt frames                           0
          Jabber frames                         0
          Fragment frames                       0
          Discarded frames                      0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 137

Configuring VLANs

- Understanding VLANs on page 2989
- Example: Configuring VLANs on page 2991

Understanding VLANs

Each VLAN is a collection of network nodes that are grouped together to form separate broadcast domains. On an Ethernet network that is a single LAN, all traffic is forwarded to all nodes on the LAN. On VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN. Frames that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within a VLAN and on the LAN as a whole.

On an Ethernet LAN, all network nodes must be physically connected to the same network. On VLANs, the physical location of the nodes is not important, so you can group network devices in any way that makes sense for your organization, such as by department or business function, by types of network nodes, or even by physical location. Each VLAN is identified by a single IP subnetwork and by standardized IEEE 802.1Q encapsulation.

To identify which VLAN the traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are tagged and are encapsulated with 802.1Q tags.

For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag. When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know to which VLAN a frame belongs. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For VLAN configuration details, see Table 305.

Table 305: VLAN Configuration Details

Field	Function	Action
General		

Copyright © 2016, Juniper Networks, Inc.

Table 305: VLAN Configuration Details *(continued)*

Field	Function	Action
VLAN Name	Specifies a unique name for the VLAN.	Enter a name. NOTE: VLAN text field is disabled when vlan-tagging is not enabled.
VLAN ID/Range	Specifies the identifier or range for the VLAN.	Select one: • **VLAN ID**—Type a unique identification number from **1** through **4094**. If no value is specified, it defaults to 1. • **VLAN Range**—Type a number range to create VLANs with IDs corresponding to the range. For example, the range 2–3 will create two VLANs with the ID 2 and 3.
Description	Describes the VLAN.	Enter a brief description for the VLAN.
Input Filter	Specifies the VLAN firewall filter that is applied to incoming packets.	To apply an input firewall filter, select the firewall filter from the list.
Output Filter	Specifies the VLAN firewall filter that is applied to outgoing packets.	To apply an output firewall filter, select the firewall filter from the list.
Ports		
Ports	Specifies the ports to be associated with this VLAN for data traffic. You can also remove the port association.	Click one: • **Add**—Select the ports from the available list. • **Remove**—Select the port that you do not want associated with the VLAN.
IP Address		
Layer 3 Information	Specifies IP address options for the VLAN.	Select to enable the IP address options.
IP Address	Specifies the IP address of the VLAN.	Enter the IP address.
Subnet Mask	Specifies the range of logical addresses within the address space that is assigned to an organization.	Enter the address, for example, **255.255.255.0**. You can also specify the address prefix.
Input Filter	Specifies the VLAN interface firewall filter that is applied to incoming packets.	To apply an input firewall filter to an interface, select the firewall filter from the list.
Output Filter	Specifies the VLAN interface firewall filter that is applied to outgoing packets.	To apply an output firewall filter to an interface, select the firewall filter from the list.
ARP/MAC Details	Specifies the details for configuring the static IP address and MAC.	Click the **ARP/MAC Details** button. Enter the static IP address and MAC address in the window that is displayed.
VoIP		

Copyright © 2016, Juniper Networks, Inc.

Table 305: VLAN Configuration Details *(continued)*

Field	Function	Action
Ports	Specifies the ports to be associated with this VLAN for voice traffic. You can also remove the port association.	Click one: • **Add**—Select the ports from the available list. • **Remove**—Select the port that you do not want associated with the VLAN.

NOTE: On SRX340 and SRX345 devices, the VLAN range from 3967 to 4094 falls under the reserved VLAN address range, and the user is not allowed any configured VLANs from this range.

Related Documentation

- Example: Configuring VLANs on page 2991
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Example: Configuring VLANs

This example shows you how to configure a VLAN.

Requirements

Before you begin:

- Determine which interfaces to use and verify that they are in switch mode. See "Example: Configuring Switching Modes" on page 2985.
- Determine what ports to use on the device and how to segment your network. See "Understanding Switching Modes" on page 2983.

Overview

In this example, you create a new VLAN and then configure attributes.

Configuration

GUI Step-by-Step Procedure

To access the VLAN:

1. In the J-Web user interface, select **Configure>Switching>VLAN**.

   The VLAN configuration page displays a list of existing VLANs. If you select a specific VLAN, the specific VLAN details are displayed in the details section.

2. Click one:

   - **Add**—Creates a VLAN.
   - **Edit**—Edits an existing VLAN configuration.

Copyright © 2016, Juniper Networks, Inc.

- **Delete**—Deletes an existing VLAN.

NOTE: If you delete a VLAN, the VLAN configuration for all the associated interfaces is also deleted.

Add or edit VLAN information.

3. Click one:
   - **OK**—Saves the configuration and returns to the main configuration page, then click **Commit Options**>**Commit**.
   - **Cancel**—Cancels your entries and returns to the main configuration page.

Related Documentation

- Understanding VLANs on page 2989
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 138

Configuring Link Aggregation Control Protocol

- Understanding Link Aggregation Control Protocol on page 2993
- Example: Configuring Link Aggregation Control Protocol on page 2997

Understanding Link Aggregation Control Protocol

LACP, a subcomponent of IEEE 802.3ad, provides additional functionality for link aggregation groups (LAGs). Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a logical point-to-point link, known as a LAG, virtual link, or bundle. The MAC client can treat this virtual link like a single link.

Starting with Junos OS Release 15.1X49-D40, Link Aggregation Control Protocol (LACP) is supported in Layer 2 transparent mode in addition to existing support in Layer 3 mode. When the SRX Series device uses LACP to bundle the member links, it creates high-speed connections, also known as *fat pipe*, with peer systems. Bandwidth can be increased by adding member links. Increased bandwidth is important especially for redundant Ethernet (reth) and aggregated Ethernet (ae) interfaces, for transmitting and receiving packets to and from the peer end for the whole system. LACP also provides automatic determination, configuration, and monitoring member links. LACP is compatible with other peers that run the 802.3ad LACP protocol. It automatically binds the member links without manually configuring the LAG, thereby avoiding errors.

NOTE: **Tentative sessions are created for all interfaces in a particular VLAN. If there is plenty of one-way traffic, numerous tentative sessions are created. When sessions reach the maximum limit, vector fails and packet loss might be seen.**

This topic contains the following sections:

- Link Aggregation Benefits on page 2994
- Link Aggregation Configuration Guidelines on page 2994

Copyright © 2016, Juniper Networks, Inc.

Link Aggregation Benefits

Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. It provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.

When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail. When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

A typical LAG deployment includes aggregate trunk links between an access switch and a distribution switch or customer edge (CE) device.

Link Aggregation Configuration Guidelines

When configuring link aggregation, note the following guidelines and restrictions:

- Link aggregation is supported only for Ethernet interfaces that are configured in switching mode (**family ethernet-switching**). Aggregating interfaces that are configured in routed mode (**family inet**) is not supported.
- You can configure a LAG by specifying the link number as a physical device and then associating a set of ports with the link. All the ports must have the same speed and be in full-duplex mode. Junos OS assigns a unique ID and port priority to each port. The ID and priority are not configurable.
- You can optionally configure LACP for link negotiation.
- You can optionally configure LACP for link protection.
- You can create up to eight Ethernet ports in each bundle.
- Each LAG must be configured on both sides of the link. The ports on either side of the link must be set to the same speed. At least one end of the LAG should be configured as active.
- LAGs are not supported on virtual chassis port links.
- By default, Ethernet links do not exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. The transmitting link is known as the actor and the receiving link is known as the partner.
- LAGs can only be used for a point-to-point connection.

For LACP configuration details, see Table 306 and Table 307.

Table 306: LACP (Link Aggregation Control Protocol) Configuration

Field	Function
Aggregated Interface	Indicates the name of the aggregated interface.

Copyright © 2016, Juniper Networks, Inc.

Table 306: LACP (Link Aggregation Control Protocol) Configuration *(continued)*

Field	Function
Link Status	Indicates whether the interface is linked (Up) or not linked (Down).
VLAN (VLAN ID)	Virtual LAN identifier value for IEEE 802.1Q VLAN tags (0.4094).
Description	The description for the LAG.

Table 307: Details of Aggregation

Field	Function
Administrative Status	Displays if the interface is enabled (Up) or disabled (Down).
Logical Interfaces	Shows the logical interface of the aggregated interface.
Member Interfaces	Member interfaces hold all the aggregated interfaces of the selected interfaces.
Port Mode	Specifies the mode of operation for the port: trunk or access.
Native VLAN (VLAN ID)	VLAN identifier to associate with untagged packets received on the interface.
IP Address/Subnet Mask	Specifies the address of the aggregated interfaces.
IPV6 Address/Subnet Mask	Specifies the IPV6 address of the aggregated interfaces.

For aggregated Ethernet interface options, see Table 308.

Table 308: Aggregated Ethernet Interface Options

Field	Function	Action
Aggregated Interface	Indicates the name of the aggregated interface.	Enter the aggregated interface name. If an aggregated interface already exists, then the field is displayed as read-only.
LACP Mode	Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are: • None—Indicates that no mode is applicable. • Active—Indicates that the interface initiates transmission of LACP packets • Passive—Indicates that the interface only responds to LACP packets.	Select from the list.
Description	The description for the LAG.	Enter the description.

Copyright © 2016, Juniper Networks, Inc.

Table 308: Aggregated Ethernet Interface Options *(continued)*

Field	Function	Action
Interface	Indicates that the interfaces available for aggregation.	Click **Add** to select the interfaces. NOTE: Only interfaces that are configured with the same speeds can be selected together for a LAG.
Speed	Indicates the speed of the interface.	
Enable Log	Specifies whether to enable generation of log entries for LAG.	Select to enable log generation.

NOTE: On SRX300, SRX320, SRX340, and SRX345 devices, the speed mode and link mode configuration are available for member interfaces of ae.

For VLAN options, see Table 309.

Table 309: Edit VLAN Options

Field	Function	Action
Port Mode	Specifies the mode of operation for the port: trunk or access.	If you select Trunk, you can: 1. Click **Add** to add a VLAN member. 2. Select the VLAN and click **OK**. 3. (Optional) Associate a native VLAN ID with the port. If you select Access, you can: 1. Select the VLAN member to be associated with the port. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. 3. Click **OK**.
VLAN Options	For trunk interfaces, the VLANs for which the interface can carry traffic.	Click **Add** to select VLAN members.
Native VLAN	VLAN identifier to associate with untagged packets received on the interface.	Select the VLAN identifier.

Related Documentation

- Example: Configuring Link Aggregation Control Protocol on page 2997
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Link Aggregation Control Protocol

This example shows how to configure LACP.

Requirements

Before you begin:

- Verify that the Ethernet interfaces are in switch mode. See "Example: Configuring Switching Modes" on page 2985.
- Link aggregation of one or more interfaces must be set up to form a virtual link or link aggregation group (LAG) before you can apply LACP. See "Understanding Switching Modes" on page 2983.

Overview

In this example, you configure link aggregation for switched Ethernet interfaces then apply LACP.

Configuration

GUI Step-by-Step Procedure

To access the LACP Configuration:

1. In the J-Web user interface, select **Configure>Interfaces>Link Aggregation**.

   The Aggregated Interfaces list is displayed.

2. Click one of the following:

   - **Device Count**—Creates an aggregated Ethernet interface, or LAG. You can choose the number of device that you want to create.
   - **Add**—Adds a new aggregated Ethernet Interface, or LAG.
   - **Edit**— Modifies a selected LAG
     - **Aggregation**—Modifies an selected LAG.
     - **VLAN**—Specifies VLAN options for the selected LAG.
     - **IP Option**—Configuring IP address to LAG is not supported and when you try to configure the IP address an error message is displayed.
   - **Delete**—Deletes the selected LAG.
   - **Disable Port** or **Enable Port**—Disables or enables the administrative status on the selected interface.

3. Click one:

   - Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options>Commit**.
   - Click **Cancel** to cancel the configuration without saving changes.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Link Aggregation Control Protocol on page 2993
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 139

Configuring Port Security

- Port Security Overview on page 2999
- Understanding MAC Limiting on page 2999
- Example: Configuring MAC Limiting on page 3000

Port Security Overview

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) attacks on network devices. Port security features help protect the access ports on your services gateway against the losses of information and productivity that can result from such attacks.

Junos OS on SRX Series devices provides features to help secure ports on a switching port on the services gateway. The ports can be categorized as either trusted or untrusted. You apply policies appropriate to those categories to protect against various types of attacks.

The MAC limit port security feature can be turned on to obtain the most robust port security level. Basic port security features are enabled in the services gateway's default configuration. You can configure additional features with minimal configuration steps.

Related Documentation

- Ethernet Ports Switching Overview on page 2984
- Understanding MAC Limiting on page 2999
- Verifying Switching Mode Configuration on page 2986

Understanding MAC Limiting

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

MAC limiting sets a limit on the number of MAC addresses that can be learned dynamically on a single Layer 2 access interface or on all the Layer 2 access interfaces on the services gateway.

Copyright © 2016, Juniper Networks, Inc.

You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses are treated as specified by the configuration.

You can choose to have one of the following actions performed when the MAC addresses limit is exceeded:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.

NOTE: MAC limit is only applied to new MAC learning requests. If you already have 10 learned MAC addresses and you configure the limit as 5, all the MACs will remain in the forwarding database (FDB) table. Once the learned MAC addresses age out (or are cleared by the user with the clear ethernet-switching command), they are not relearned.

MAC limiting does not apply to static MAC addresses. Users can configure any number of static MAC addresses independent of MAC limiting and all of them are added to FDB.

Related Documentation

- Example: Configuring MAC Limiting on page 3000
- Port Security Overview on page 2999
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Example: Configuring MAC Limiting

- Requirements on page 3000
- Overview on page 3000
- Configuration on page 3001
- Verification on page 3001

Requirements

Before you begin, verify that the interfaces that will be used are in switch mode. See "Example: Configuring Switching Modes" on page 2985 and "Understanding Switching Modes" on page 2983.

Overview

MAC limiting protects against flooding of the Ethernet switching table on the SRX Series Services Gateways. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

This example shows how to configure port security features by setting a MAC limit of 5.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

The action is not specified, so the switch performs the default action drop if the limit is exceeded:

1. On a single interface (here, the interface is ge-0/0/1):

```
[edit ethernet-switching-options secure-access-port]
user@host# set interface ge–0/0/1 mac-limit 5
```

2. On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@host# set interface all mac–limit 5
```

NOTE: Do not set the mac-limit to 1. The first learned MAC address is often inserted into the FDB automatically (for example, for routed VLAN interfaces the first MAC address inserted into the forwarding database is the MAC address of the RVI; for Aggregated Ethernet bundles using LACP, the first MAC address inserted into the FDB in the forwarding table is the source address of the protocol packet). The services gateway will therefore not learn MAC addresses other than the automatic addresses when the mac-limit is set to 1, and this will cause problems with MAC learning and forwarding.

3. For specifying specific allowed MAC addresses:

   - On a single interface (here, the interface is ge-0/0/2):

```
[edit ethernet-switching-options secure-access-port]
user@host# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:80
user@host# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:81
user@host# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:83
```

   - On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@host# set interface all allowed-mac 00:05:85:3A:82:80
user@host# set interface all allowed-mac 00:05:85:3A:82:81
user@host# set interface all allowed-mac 00:05:85:3A:82:83
```

4. Enter **commit** from configuration mode.

Verification

Verifying That MAC Limiting Is Working Correctly on the Services Gateway

Purpose

Verify that MAC limiting is working on the services gateway.

Copyright © 2016, Juniper Networks, Inc.

Action Display the learned MAC addresses. The following sample output shows the results when two packets were sent from hosts on ge-0/0/1 and five packets requests were sent from hosts on ge-0/0/2, with both interfaces set to a MAC limit of 4 with the action drop:

```
user@host> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
VLAN MAC address Type Age Interfaces
employee-vlan * Flood - ge-0/0/2.0
employee-vlan 00:05:85:3A:82:77 Learn 0 ge-0/0/1.0
employee-vlan 00:05:85:3A:82:79 Learn 0 ge-0/0/1.0
employee-vlan 00:05:85:3A:82:80 Learn 0 ge-0/0/2.0
employee-vlan 00:05:85:3A:82:81 Learn 0 ge-0/0/2.0
employee-vlan 00:05:85:3A:82:83 Learn 0 ge-0/0/2.0
employee-vlan 00:05:85:3A:82:85 Learn 0 ge-0/0/2.0
```

Meaning The sample output shows that with a MAC limit of 4 for each interface, the packet for a fifth MAC address on ge-0/0/2 was dropped because it exceeded the MAC limit. The address was not learned, and thus an asterisk (*) rather than an address appears in the MAC address column in the first line of the sample output.

Related Documentation

- Understanding MAC Limiting on page 2999
- Ethernet Ports Switching Overview on page 2984
- Verifying Switching Mode Configuration on page 2986

Copyright © 2016, Juniper Networks, Inc.

PART 43

Configuration Statements and Operational Commands

- Configuration Statements on page 3005
- Operational Commands on page 3095

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 140

Configuration Statements

- [edit security forwarding-options] Hierarchy Level on page 3006
- Access Configuration Statement Hierarchy on page 3007
- Class-of-Service Configuration Statement Hierarchy on page 3015
- Interfaces Configuration Statement Hierarchy on page 3019
- VLANS Configuration Statement Hierarchy on page 3035
- authentication-order (Access Profile) on page 3040
- code-points (CoS) on page 3041
- destination-address (Security Policies) on page 3042
- domain-type (VLANs) on page 3042
- encapsulation (Interfaces) on page 3043
- ethernet-switching on page 3044
- family inet (Interfaces) on page 3045
- family inet6 on page 3048
- flow (Security Flow) on page 3051
- forwarding-classes (CoS) on page 3053
- global-mac-table-aging-time (Protocols) on page 3054
- global-mac-limit (Protocols) on page 3055
- global-mode (Protocols) on page 3056
- global-no-mac-learning (Protocols) on page 3056
- host-inbound-traffic on page 3057
- inet6 (Security Forwarding Options) on page 3058
- interfaces (CoS) on page 3059
- interfaces (Security Zones) on page 3060
- interface (Switching Options) on page 3061
- interface (VLANs) on page 3062
- l2-learning (Protocols) on page 3063
- loss-priority (CoS Loss Priority) on page 3064
- match (Security Policies) on page 3065

Copyright © 2016, Juniper Networks, Inc.

- native-vlan-id (Interfaces) on page 3066
- peer-selection-service on page 3067
- pgcp-service on page 3068
- policy (Security Policies) on page 3069
- port (Access RADIUS) on page 3071
- profile (Access) on page 3072
- radius-server (Access) on page 3075
- redundancy-group (Interfaces) on page 3076
- secure-wire on page 3076
- security-zone on page 3077
- shaping-rate (CoS Interfaces) on page 3079
- source-address (Access RADIUS) on page 3080
- source-address (Security Policies) on page 3081
- static-mac (VLANs) on page 3082
- switch-options (VLANs) on page 3083
- system-services (Security Zones Interfaces) on page 3084
- unframed | no-unframed (Interfaces) on page 3085
- vlans on page 3086
- vlan-id (VLAN) on page 3092
- vlan members (VLANs) on page 3093
- vlan-tagging (Interfaces) on page 3094

[edit security forwarding-options] Hierarchy Level

```
security {
  forwarding-options {
    family {
      inet6 {
        mode (drop | flow-based | packet-based);
      }
      iso {
        mode packet-based;
      }
      mpls {
        mode packet-based;
      }
    }
  }
  mirror-filter filter-name {
    destination-port port-number;
    destination-prefix destination-prefix;
    interface-in interface-name;
    interface-out interface-name;
    output {
      destination-mac mac-address;
```


Copyright © 2016, Juniper Networks, Inc.

```
        interface interface-name;
      }
      protocol protocol;
      source-port port-number;
      source-prefix source-prefix;
    }
    secure-wire secure-wire-name interface [interface-name-1 interface-name-2];
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Access Configuration Statement Hierarchy

Use the statements in the **access** configuration hierarchy to configure access to the device and authentication methods, including address assignment and address pool, user and firewall authentication, a group profile, LDAP options and LDAP server configuration, an access profile, RADIUS options and RADIUS server configuration, and SecurID server configuration.

```
access {
  address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-name;
    pool pool-name {
      family {
        inet {
          dhcp-attributes {
            boot-file boot-file-name;
            boot-server boot-server-name;
            domain-name domain-name;
            grace-period seconds;
            maximum-lease-time (seconds | infinite);
            name-server ipv4-address;
            netbios-node-type (b-node | h-node | m-node | p-node);
            next-server next-server-name;
            option dhcp-option-identifier-code {
              array {
                byte [8-bit-value];
                flag [ false| off |on |true];
                integer [32-bit-numeric-values];
                ip-address [ip-address];
                short [signed-16-bit-numeric-value];
                string [character string value];
                unsigned-integer [unsigned-32-bit-numeric-value];
                unsigned-short [16-bit-numeric-value];
              }
              byte 8-bit-value;
              flag (false | off | on | true);
              integer 32-bit-numeric-values;
              ip-address ip-address;
              short signed-16-bit-numeric-value;
```

Copyright © 2016, Juniper Networks, Inc.

```
            string character string value;
            unsigned-integer unsigned-32-bit-numeric-value;
            unsigned-short 16-bit-numeric-value;
          }
          option-match {
            option-82 {
              circuit-id match-value {
                range range-name;
              }
              remote-id match-value;
                range range-name;
              }
            }
          }
          propagate-ppp-settings [interface-name];
          propagate-settings interface-name;
          router ipv4-address;
          server-identifier ip-address;
          sip-server {
            ip-address ipv4-address;
            name sip-server-name;
          }
          tftp-server server-name;
          wins-server ipv4-address;
        }
        host hostname {
          hardware-address mac-address;
          ip-address reserved-address;
        }
        network network address;
        range range-name {
          high upper-limit;
          low lower-limit;
        }
        xauth-attributes {
          primary-dns ip-address;
          primary-wins ip-address;
          secondary-dns ip-address;
          secondary-wins ip-address;
        }
      }
      inet6 {
        dhcp-attributes {
          dns-server ipv6-address;
          grace-period seconds;
          maximum-lease-time (seconds | infinite);
          option dhcp-option-identifier-code {
            array {
              byte [8-bit-value];
              flag [ false| off |on |true];
              integer [32-bit-numeric-values];
              ip-address [ip-address];
              short [signed-16-bit-numeric-value];
              string [character string value];
              unsigned-integer [unsigned-32-bit-numeric-value];
              unsigned-short [16-bit-numeric-value];
```


Copyright © 2016, Juniper Networks, Inc.

```
                }
                byte 8-bit-value;
                flag (false | off | on | true);
                integer 32-bit-numeric-values;
                ip-address ip-address;
                short signed-16-bit-numeric-value;
                string character string value;
                unsigned-integer unsigned-32-bit-numeric-value;
                unsigned-short 16-bit-numeric-value;
              }
              propagate-ppp-settings [interface-name];
              sip-server-address ipv6-address;
              sip-server-domain-name domain-name;
            }
            prefix ipv6-network-prefix;
            range range-name {
              high upper-limit;
              low lower-limit;
              prefix-length delegated-prefix-length;
            }
          }
        link pool-name;
      }
    }
    address-pool pool-name {
      (address address-or-address-prefix ) {
      address-range {
        high upper-limit;
        low lower-limit;
        mask network-mask;
      }
      primary-dns name;
      primary-wins name;
      secondary-dns name;
      secondary-wins name;
    }
    address-protection;
    domain {
      delimiter delimiter;
      map domain-map-name {
        aaa-logical-system logical-system-name;
        aaa-routing-instance routing-instance-name;
        access-profile access-profile-name;
        address-pool address-pool-name;
        dynamic-profile dynamic-profile-name;
        padn destination-address; {
          mask destination-mask;
          metric metric-value
        }
        strip-domain;
        target-logical-system logical-system-name;
        target-routing-instance target-routing-instance;
      }
      parse-direction (left-to-right | right-to-left);
    }
    firewall-authentication {
```

Copyright © 2016, Juniper Networks, Inc.

```
    pass-through {
      default-profile profile-name;
      ftp {
        banner {
          fail string;
          login string;
          success string;
        }
      }
      http {
        banner {
          fail string;
          login string;
          success string;
        }
      telnet {
        banner {
          fail string;
          login string;
          success string;
        }
    }
    traceoptions {
      file {
        filename;
        files number;
        flag flag;
        match regular-expression;
        no-remote-trace;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
    }
    web-authentication {
      banner {
        success string;
      }
      default-profile profile-name;
    }
  }
  group-profile profile-name {
    ppp {
      cell-overhead;
      encapsulated-overhead encapsulated-overhead-value;
      framed-pool address-pool-name;
      idle-timeout seconds;
      interface-id interface-identifier;
      keepalive seconds;
      ppp-options {
        chap;
        pap;
      }
      primary-dns name;
      primary-wins name;
      secondary-dns name;
      secondary-wins name;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  gx-plus {
    global {
      max-outstanding-requests max-outstanding-requests;
    }
    partition partition-name {
      destination-host gx-plus-destination-host;
      destination-realm gx-plus-destination-realm;
      diameter-instance gx-plus-diameter-instance;
    }
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter filter-name;
    }
  }
  ldap-server hostname-or-address; {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  ppp-options {
    compliance {
      rfc(2486 | [rfc-number]);
    }
  }
  profile profile-name  {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      coa-immediate-update;
      duplication;
      immediate-update;
      order [accounting-method];
      statistics (time | volume-time);
      update-interval minutes;
    }
    accounting-order [accounting-method];
    address-assignment pool pool-name;
    authentication-order [ldap | none | password | radius | securid];
    authorization-order [jsrc];
    client client-name  {
      chap-secret  chap-secret;
      client-group [ group-names ];
```

```
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
  radius {
    accounting-server [server];
    attributes {
      exclude {
        acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
        acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
        acc-loop-cir-id [access-request | accounting-start | accounting-stop];
        accounting-authentic [accounting-off | accounting-on | accounting-start |
          accounting-stop];
        accounting-delay-time [accounting-off | accounting-on | accounting-start |
          accounting-stop];
        accounting-session-id [access-request];
        accounting-terminate-cause [accounting-off];
        act-data-rate-dn [access-request | accounting-start | accounting-stop];
        act-data-rate-up [access-request | accounting-start | accounting-stop];
        act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
        act-interlv-delay-up [access-request | accounting-start | accounting-stop];
        att-data-rate-dn [access-request | accounting-start | accounting-stop];
        att-data-rate-up [access-request | accounting-start | accounting-stop];
        called-station-id [access-request | accounting-start | accounting-stop];
        calling-station-id [access-request | accounting-start | accounting-stop];
        class [access-request | accounting-start | accounting-stop];
```


Copyright © 2016, Juniper Networks, Inc.

```
    delegated-ipv6-prefix [accounting-start | accounting-stop];
    dhcp-gi-address [access-request | accounting-start | accounting-stop];
    dhcp-mac-address [access-request | accounting-start | accounting-stop];
    dhcp-options [access-request | accounting-start | accounting-stop];
    downstream-calculated-qos-rate [access-request | accounting-start |
      accounting-stop];
    dsl-forum-attributes [access-request | accounting-start | accounting-stop];
    dsl-line-state [access-request | accounting-start | accounting-stop];
    dsl-type [access-request | accounting-start | accounting-stop];
    dynamic-iflset-name [accounting-start | accounting-stop];
    event-time-stamp [accounting-off | accounting-on | accounting-start |
      accounting-stop];
    framed-interface-id [access-request | accounting-start | accounting-stop];
    framed-ip-address [access-request | accounting-start | accounting-stop];
    framed-ip-netmask [access-request | accounting-start | accounting-stop];
    framed-ip-route [access-request | accounting-start | accounting-stop];
    framed-ipv6-pool [accounting-start | accounting-stop];
    framed-ipv6-prefix [accounting-start | accounting-stop];
    framed-ipv6-route [accounting-start | accounting-stop];
    framed-pool [accounting-start | accounting-stop];
    input-filter [accounting-start | accounting-stop];
    input-gigapackets [accounting-stop];
    input-gigawords [accounting-stop];
    input-ipv6-gigawords [accounting-stop];
    input-ipv6-octets [accounting-stop];
    input-ipv6-packets [accounting-stop];
    interface-description [access-request | accounting-start | accounting-stop];
    l2c-downstream-data [access-request | accounting-start | accounting-stop];
    l2c-upstream-data [access-request | accounting-start | accounting-stop];
    max-data-rate-dn [access-request | accounting-start | accounting-stop];
    max-data-rate-up [access-request | accounting-start | accounting-stop];
    max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
    max-interlv-delay-up [access-request | accounting-start | accounting-stop];
    min-data-rate-dn [access-request | accounting-start | accounting-stop];
    min-data-rate-up [access-request | accounting-start | accounting-stop];
    min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
    min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
    nas-identifier [access-request | accounting-start | accounting-stop];
    nas-port [access-request | accounting-off | accounting-on | accounting-start |
      accounting-stop];
    nas-port-id [access-request | accounting-start | accounting-stop];
    nas-port-type [access-request | accounting-start | accounting-stop];
    output-filter [accounting-start | accounting-stop];
    output-gigapackets [accounting-stop];
    output-gigawords [accounting-stop];
    output-ipv6-gigawords [accounting-stop];
    output-ipv6-octets [accounting-stop];
    output-ipv6-packets [accounting-stop];
    upstream-calculated-qos-rate [access-request | accounting-start |
      accounting-stop];
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    input-filter;
    logical-system-routing-instance;
```

Copyright © 2016, Juniper Networks, Inc.

```
            output-filter;
          }
        }
      authentication-server [server];
      radius-options {
        request-rate number;
        revert-interval seconds;
      }
      radius-server server-address {
        accounting-port port-number
        max-outstanding-requests number-of--outstanding-requests;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
      }
      service {
        accounting-order {
          activation-protocol;
          radius;
        }
      }
      session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
      }
    }
    radius-options {
      request-rate number;
      revert-interval seconds;
    }
    radius-server server-address {
      accounting-port port-number;
      max-outstanding-requests number-of-max-outstanding-requests;
      port port-number;
      retry attempts;
      routing-instance routing-instance-name;
      secret password;
      source-address source-address;
      timeout seconds;
    }
    securid-server server-name {
      configuration-file filepath;
    }
    terminate-code {
      aaa {
        deny {
          authentication-denied {
            radius acct-terminate-cause-value;
          }
          no-resources {
            radius acct-terminate-cause-value;
          }
```


Copyright © 2016, Juniper Networks, Inc.

```
            server-request-timeout {
              radius acct-terminate-cause-value;
            }
          }
          shutdown {
            administrative-reset {
              radius acct-terminate-cause-value;
            }
            remote-reset {
              radius acct-terminate-cause-value;
            }
          }
        }
        dhcp {
          client-request {
            radius acct-terminate-cause-value;
          }
          lost-carrier {
            radius acct-terminate-cause-value;
          }
          nak {
            radius acct-terminate-cause-value;
          }
          nas-logout {
            radius acct-terminate-cause-value;
          }
          no-offers {
            radius acct-terminate-cause-value;
          }
        }
      }
    }
```

Related Documentation

- *Understanding User Authentication Methods*
- Layer 2 Transparent and Switching Overview on page 2893
- Understanding User Authentication for Security Devices on page 5173

Class-of-Service Configuration Statement Hierarchy

Use the statements in the **class-of-service** configuration hierarchy to configure class-of-services (CoS) features.

```
class-of-service {
  adaptive-shapers adaptive-shaper-name {
    trigger becn {
      shaping-rate (absolute-rate | percent percent);
    }
  }
  application-traffic-control {
    rate-limiters rate-limiter-name {
      bandwidth-limit kbps;
      burst-size-limit bytes;
    }
```

```
    rule-sets rule-set-name {
      rule rule-name {
          match {
            application [application-name];
            application-any;
            application-group [application-group-name];
            application-known;
            application-unknown;
          }
          then {
            dscp-code-point dscp-value;
            forwarding-class class-name;
            log;
            loss-priority (high | low |medium-high | medium-low);
            rate-limit {
              loss-priority-high;
              client-to-server rate-limiter;
              server-to-client rate-limiter;
            }
          }
        }
      }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier name
      {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-points [alias-or-bit-string ];
        }
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 |exp |ieee-802.1 |ieee-802.1ad |inet-precedence) alias-name{
      dscp-bits;
      }
  }
  drop-profiles profile-name {
    fill-level percent {
      drop-probability number;
    }
    interpolate {
```


Copyright © 2016, Juniper Networks, Inc.

```
      drop-probability [number];
      fill-level [percent];
    }
  }
  forwarding-classes {
    class class-name {
      priority (high | low);
      queue-num number;
      spu-priority (high | low);
    }
    queue queue-number {
      class-name {
        priority (high | low);
      }
    }
  }
  forwarding-policy {
    class class-name {
      classification-override {
        forwarding-class class-name;
      }
    }
    next-hop-map next-hop-map-name {
      forwarding-class class-name {
        discard;
        lsp-next-hop [lsp-regular-expression];
        next-hop [next-hop-identifier];
        non-lsp-next-hop;
      }
    }
  }
  fragmentation-maps fragmentation-map-name {
    forwarding-class forwarding-class-name {
      drop-timeout milliseconds;
        (fragment-threshold bytes |no-fragmentation) ;
        multilink-class number;
      }
    }
  host-outbound-traffic {
    dscp-code-point static-dscp-code-point;
    forwarding-class class-name;
    tcp {
      raise-internet-control-priority;
    }
  }
  interfaces interface-name {
    input-traffic-control-profile profile-name;
      output-traffic-control-profile profile-name;
      output-traffic-control-profile-remaining profile-name;
      scheduler-map scheduler-map;
      shaping-rate bps;
      unit  logical-unit-number  {
        adaptive-shaper adaptive-shaper-name;
        classifiers {
          (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
        forwarding-class class-name;
        input-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        loss-priority-maps {
          frame-relay-de {
            (lpmap-name | default);
          }
        }
        output-traffic-control-profile {
          profile-name;
          shared-instance shared-instance-name;
        }
        rewrite-rules {
          (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad
            |inet-precedence)
        }
        scheduler-map scheduler-map-name;
          shaping-rate {
          rate;
        }
        vc-shared-scheduler;
        virtual-channel-group group-name;
      }
    }
  }
  loss-priority-maps {
    frame-relay-de loss-priority-map-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-points [bit-string];
      }
    }
  }
  rewrite-rules {
    (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      rewrite-rule-name {
      forwarding-class forwarding-class-name {
        loss-priority (high | low | medium-high | medium-low) {
          code-point alias-or-bit-string;
        }
        import (default | rewrite-rule-name);
      }
    }
  }
  scheduler-maps scheduler-map-name {
    forwarding-class class-name {
      scheduler scheduler-name;
      }
    }
  schedulers scheduler-name {
    buffer-size {
      exact;
      (percent percent | remainder percent  | temporal microseconds) ;
    }
    drop-profile-map {
```


Copyright © 2016, Juniper Networks, Inc.

```
        loss-priority (any | high | low | medium-high | medium-low);
        protocol any ;
        drop-profile profile;
      }
      priority (high | low | medium-high | medium-low | strict-high);
      shaping-rate (absolute-rate | percent percent);
      transmit-rate <exact> (percent percent | rate bits | remainder percent);
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
    traffic-control-profiles profile-name {
      delay-buffer-rate ( absolute-rate | cps cells-per-second | percent percent);
        guaranteed-rate (absolute-rate | percent percent);
        overhead-accounting (bytes bytes | cell-mode | frame-mode);
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    tri-color;
    virtual-channel-groups virtual-channel-group-name {
      virtual-channel-name {
        default;
        scheduler-map scheduler-map-name;
        shaping-rate (absolute-rate | percent percent);
      }
    }
    virtual-channels virtual-channel-name;
  }
```

Related Documentation

- SSL Proxy Overview on page 531
- Understanding Interfaces on page 2441

Interfaces Configuration Statement Hierarchy

Use the statements in the **interfaces** configuration hierarchy to configure interfaces on the device.

NOTE: For a gigabit ethernet interface, the **gigether-options** and **ether-options** are identical, but only the **gigether-options** are documented.

```
interfaces {
  interface-name {
    accounting-profile name;
```

```
clocking (external | internal);
dce;
description text;
disable;
e1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  invert-data data;
  loopback (local | remote);
  start-end-flag (shared | filler);
  timeslots time-slot-range;
}
e3-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  compatibility-mode {
    digital-link {
      subrate value;
    }
    kentrox {
      subrate value;
    }
    larscom;
  }
  fcs (16 | 32);
  framing (g.751 | g.832);
  idle-cycle-flag value;
  invert-data;
  loopback (local | remote);
  (no-payload-scrambler | payload-scrambler);
  (no-unframed | -unframed);
  start-end-flag (filler | shared);
}
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc |
  ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc |
  extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls
  | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
fastether-options {
  802.3ad interface-name {
    (backup | primary);
    lacp {
      port-priority port-number;
    }
  }
  (auto-negotiation | no-auto-negotiation);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
}
flexible-vlan-tagging;
gigether-options {
    802.3ad interface-name {
        (backup | primary);
        lacp {
            port-priority port-number;
        }
    }
    (auto-negotiation <remote-fault> (local-interface-offline | local-interface-online)
        | no-auto-negotiation);
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth [number];
        }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
}
gratuitous-arp-reply;
hierarchical-scheduler {
    maximum-hierarchy-levels 2;
}
hold-time {
    down milliseconds;
    up milliseconds;
}
keepalives {
    down-count number;
    interval number;
    up-count number;
}
link-mode (full-duplex | half-duplex);
lmi {
    lmi-type (ansi | c-lmi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte number;
    t392dce number;
}
logical-tunnel-options {
    per-unit-mac-disable;
}
mac mac-address;
mtu bytes;
native-vlan-idvlan-id;
```

Copyright © 2016, Juniper Networks, Inc.

```
no-gratuitous-arp-request;
no-keepalives;
optics-options {
  alarm {
    low-light-alarm (link-down | syslog);
  }
  warning {
    low-light-warning (link-down | syslog);
  }
  wavelength wavelength-options;
}
otn-options {
  bytes {
    transmit-payload-type number];
  }
  fec (efec | gfec | none);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
  trigger {
    oc-lof {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-lom {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-los {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-wavelength-lock {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    odu-ais {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    odu-bdi {
```


Copyright © 2016, Juniper Networks, Inc.

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-lck {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-oci {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-sd {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-bbe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-es {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-ses {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-uas {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-ttim {
```

Copyright © 2016, Juniper Networks, Inc.

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  opu-ptim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-ais {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-exe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-iae {
```


Copyright © 2016, Juniper Networks, Inc.

```
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-sd {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-tca-bbe {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-tca-es {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-tca-ses {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-tca-uas {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
        otu-ttim {
            hold-time {
                down milliseconds;
                up milliseconds;
                }
            ignore;
        }
    }
    tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi |
        otu-dapi |otu-expected-receive-dapi | otu-expected-receive-sapi |otu-sapi);
}
passive-monitor-mode;
(per-unit-scheduler | no-per-unit-schedule);
port-mirror-instance;
ppp-options {
```

```
    chap {
      access-profile name;;
      default-chap-secret secret;
      local-name name;
      no-rfc2486;
      passive;
    }
    compression {
      acfc;
      pfc;
    }
    dynamic-profile (dynamic-profile | junos-default-profile);
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
      access-profile name;
      default-password password;
      local-name name;
      local-password password;
      no-rfc2486;
      passive;
    }
  }
  promiscuous-mode;
  receive-bucket {
    overflow {
      discard;
      tag;
    }
    rate number;
    threshold number;
  }
  redundant-pseudo-interface-options {
    redundancy-group number;
  }
  satop-options {
    excessive-packet-loss-rate {
      sample-period milliseconds;
      threshold percentage;
    }
    idle-pattern number;
    (jitter-buffer-auto-adjust | jitter-buffer-latency milliseconds | jitter-buffer-packets
      number;
    payload-size number;
  }
  speed (100m |10m | 1g);
  stacked-vlan-tagging;
  switch-options {
    switch-port port-number {
      (auto-negotiation | no-auto-negotiation);
      cascade-port;
      link-mode (full-duplex | half-duplex);
```


Copyright © 2016, Juniper Networks, Inc.

```
      speed (100m |10m | 1g);
      vlan-id number;
    }
  }
  t1-options {
    alarm-compliance {
      accunet-t1-5-service;
    }
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flags (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
  }
  t3-options {
    bert-algorithm algorithm ;
    bert-error-rate rate ;
    bert-period seconds ;
    (cbit-parity | no-cbit-parity);
    compatibility-mode {
      adtran {
        subrate value;
      }
      digital-link {
        subrate value;
      }
      kentrox {
        subrate value;
      }
      larscom;
        subrate value;
      }
      verilink;
        subrate value;
      }
    }
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag (flags | ones);
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (no-payload-scrambler | payload-scrambler);
    (no-unframed | unframed);
    start-end-flag value (filler | shared);
  }
  traceoptions {
```

Copyright © 2016, Juniper Networks, Inc.

```
    flag (all | event | ipc | media);
}
transmit-bucket {
    overflow {
        discard;
    }
    rate number;
    threshold number;
}
(traps | no-traps);
unit unit-number {
    accept-source-mac {
        mac-address mac-address;
    }
    accounting-profile name;
    arp-resp (restricted | unrestricted);
    backup-options {
        interface interface-name;
    }
    bandwidth bandwidth;
    description text;
    disable;
    encapsulation (dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge |
        vlan-ccc | vlan-vpls |vlan-tcc);
    family {
        ccc {
            filter {
                group number;
                input filter-name;
                input-list [filter-name];
                output filter-name;
                output-list [filter-name];
            }
            policer {
                input input-policer-name;
                output output-policer-name;
            }
        }
        ethernet-switching {
            bridge-domain-type (svlan| bvlan);
            filter {
                group number;
                input filter-name;
                input-list [filter-name];
                output filter-name;
                output-list [filter-name];
            }
            interface-mode (access | trunk);
            native-vlan-id native-vlan-id;
            policer {
                input input-policer-name;
                output outputpolicer-name;
            }
            port-mode (access | tagged-access | trunk);
            reflective-relay;
            vlan-id vlan-id;
```


Copyright © 2016, Juniper Networks, Inc.

```
    vlan members [vlan-id];
    vlan-rewrite {
      translate {
        from-vlan-id;
        to-vlan-id ;
      }
    }
  }
  inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
    address (source–address/prefix) {
      arp destination-address {
        (mac mac-address | multicast-mac multicast-mac-address);
        publish publish-address;
      }
      broadcast address;
      preferred;
      primary;
      vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key-value;
        authentication-type (md5 | simple);
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds
        (preempt <hold-timeseconds> | no-preempt );
        priority value;
        track {
          interface interface-name {
            bandwidth-threshold bandwidth;
            priority-cost value;
          }
          priority-hold-time seconds;
          route route-address{
            routing-instance routing-instance;
            priority-cost value;
          }
        }
        virtual-address [address];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group value;
          active-interface interface-name;
        }
      }
      web-authentication {
        http;
        https;
        redirect-to-https;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  dhcp {
    client-identifier {
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      user-id (ascii string| hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp arp-name;
    input input-name;
    output output-name;
  }
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
    simple-filter;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    targeted-broadcast {
      (forward-and-send-to-re |forward-only);
    }
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
  }
  inet6 {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        ouput;
      }
    }
    address source–address/prefix {
      eui-64;
      ndp address {
        (mac mac-address | multicast-mac multicast-mac-address);
        publish;
      }
      preferred;
      primary;
      vrrp-inet6-group group_id {
        (accept-data | no-accept-data);
        advertisements-threshold number;
        authentication-key value;
        authentication-type (md5 | simple);
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        (preempt <hold-time seconds>| no-preempt );
        priority value;
        track {
          interface interface-name {
            bandwidth-threshold value;
            priority-cost value;
          }
          priority-hold-time seconds;
          route route-address{
            routing-instance routing-instance;
          }
        }
        virtual-inet6-address [address];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group value;
          active-interface interface-name;
        }
      }
      web-authentication {
        http;
        https;
        redirect-to-https;
      }
```

```
    }
    (dad-disable | no-dad-disable);
    dhcpv6-client {
      client-ia-type (ia-na | ia-pd);
      client-identifier duid-type (duid-ll | duid-llt | vendor);
      client-type (autoconfig | stateful);
      rapid-commit;
      req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server |
        sip-domain | sip-server |time-zone | vendor-spec);
      retransmission-attempt number;
      update-router-advertisement {
        interface interface-name;
      }
      update-server;
    }
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
    }
    mtu value;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
      input input-name;
      output output-name;
    }
    rpf-check {
      fail-filter filter-name;
      mode {
        loose;
      }
    }
    sampling {
      input;
      output;
    }
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
  }
  iso {
    address source-address;
    mtu value;
  }
  mlfr-end-to-end {
    bundle bundle-name;
  }
  mlfr-uni-nni {
    bundle bundle-name;
  }
  mlppp {
    bundle bundle-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    mpls {
      filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
      }
      mtu mtu-value;
      policer {
        input input-name;
        output output-name;
      }
    }
    tcc {
      policer {
        input input-name;
        output output-name;
      }
      proxy {
        inet-address inet-address;
      }
      remote {
        inet-address inet-address;
        mac-address mac-address;
      }
    }
    vpls {
      filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
      }
      policer {
        input input-name;
        output output-name;
      }
    }
  }
  input-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number ;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
  interface-shared-with {
    psd-name;
  }
  native-inner-vlan-id value;
  (no-traps | traps);
  output-vlan-map {
    inner-tag-protocol-id tpid;
```

```
            inner-vlan-id number;
            (pop | push | swap);
            tag-protocol-id tpid;
            vlan-id number;
          }
          ppp-options {
            chap {
              access-profile name;
              default-chap-secret name;
              local-name name;
              no-rfc2486;
              passive;
            }
            dynamic-profile profile-name;
            lcp-max-conf-req number;
            lcp-restart-timer milliseconds;
            loopback-clear-timer seconds;
            ncp-max-conf-req number;
            ncp-restart-timer milliseconds;
            no-termination-request;
            pap {
              access-profile name;
              default-password password;
              local-name name;
              local-password password;
              no-rfc2486;
              passive;
            }
          }
          proxy-arp (restricted | unrestricted);
          radio-router {
            bandwidth number;
            credit {
              interval number;
            }
            data-rate number;
            latency number;
            quality number;
            resource number;
            threshold number;
          }
          swap-by-poppush;
          traps;
          vlan-id vlan-id;
          vlan-id-range vlan-id-range;
          vlan members [vlan-id];
          vlan-id-range vlan-id1-vlan-id2;
          vlan-tags {
            (inner vlan-id | inner-range vlan-id1-vlan-id2);
            inner-list [vlan-id];
            outer vlan-id;
          }
        }
        vlan-tagging;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Interfaces on page 2441

VLANS Configuration Statement Hierarchy

Use the statements in the **vlans** configuration hierarchy to configure a bridging domain that includes a set of logical ports that share the same flooding or broadcast characteristics.

```
vlans vlans-name {
  description  text;
  domain-type bridge;
  forwarding-options {
    dhcp-relay {
      active-server-group active-server-group-name;
      dhcpv6 {
        active-server-group active-server-group-name;
        authentication {
          password password;
          username-include {
            circuit-type;
            client-id;
            delimiter delimiter;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
          }
        }
        dynamic-profile (dynamic-profile-name | junos-default-profile) {
          aggregate-clients (merge | replace);
          use-primary ( primary-profile-name | junos-default-profile);
        }
        group group-name {
          active-server-group active-server-group-name;
          authentication {
            password password;
            username-include {
              circuit-type;
              client-id;
              delimiter delimiter;
              domain-name domain-name;
              interface-name;
              logical-system-name;
              relay-agent-interface-id;
              relay-agent-remote-id;
              relay-agent-subscriber-id;
              routing-instance-name;
              user-prefix user-prefix;
            }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
dynamic-profile (dynamic-profile-name | junos-default-profile) {
  aggregate-clients (merge | replace) {
    use-primary (primary-profile-name | junos-default-profile);
  }
}
interface interface-name {
  dynamic-profile (dynamic-profile-name| junos-default-profile) {
    aggregate-clients (merge | replace) {
      use-primary (primary-profile-name | junos-default-profile);
    }
  }
  exclude;
  overrides {
    (allow-snooped-clients | no-allow-snooped-clients);
    interface-client-limit number;
    no-bind-on-request;
    send-release-on-delete;
  }
  service-profile profile-name;
  trace;
  upto upto-interface-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      detection-time {
        threshold milliseconds;
      }
      holddown-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | singlehop);
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (0 | 1 | automatic);
    }
  }
}
overrides {
  (allow-snooped-clients | no-allow-snooped-clients);
  interface-client-limit number;
  no-bind-on-request;
  send-release-on-delete;
}
relay-agent-interface-id {
  prefix {
    host-name;
    logical-system-name;
    routing-instance-name;
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
        use-interface-description (device | logical):
      }
      service-profile dynamic-profile-name;
    }
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          detection-time {
            threshold milliseconds;
          }
          holddown-interval milliseconds;
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          session-mode (automatic | multihop | singlehop);
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          version (0 | 1 | automatic);
        }
      }
    }
    overrides {
      (allow-snooped-clients | no-allow-snooped-clients);
      interface-client-limit number;
      no-bind-on-request;
      send-release-on-delete;
    }
    relay-agent-interface-id {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical):
    }
    server-group {
      server-group-name {
        server-ip-address;
      }
    }
    service-profile service-profile-name;
  }
  group group-name {
    active-server-group server-group-name;
    interface interface-name {
      exclude;
      upto interface-name;
    }
    relay-option-60 {
      vendor-option {
        default-local-server-group local-server-group-name;
        default-relay-server-group server-group-name;
```

```
            drop;
            equals {
              ascii ascii-name {
                drop;
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
              hexadecimal hexadecimal {
                drop;
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
            }
            starts-with {
              ascii ascii-name {
                drop;
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
              hexadecimalhexadecimal {
                drop;
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
            }
          }
        }
        relay-option-82 {
          circuit-id {
            prefix {
              host-name;
              logical-system-name;
              routing-instance-name;
            }
            use-interface-description (device | logical);
          }
        }
      }
      relay-option-60 {
        vendor-option {
          default-local-server-group local-server-group-name;
          default-relay-server-group server-group-name;
          drop;
          equals {
            ascii ascii-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group server-group-name;
            }
            hexadecimal hexadecimal {
              drop;
              local-server-group local-server-group-name;
              relay-server-group server-group-name;
            }
          }
          starts-with {
```


Copyright © 2016, Juniper Networks, Inc.

```
                ascii ascii-name {
                  drop;
                  local-server-group local-server-group-name;
                  relay-server-group server-group-name;
                }
                hexadecimal hexadecimal {
                  drop;
                  local-server-group local-server-group-name;
                  relay-server-group server-group-name;
                }
              }
            }
          }
          relay-option-82 {
            circuit-id {
              prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
              }
              use-interface-description (device | logical);
            }
          }
          server-group server-group-name {
            ip-address;
          }
        }
      }
      filter {
        input input-filter-name;
      }
      flood {
        input input-filter-name;
      }
    }
    interface interface-name;
    isolated-vlan isolated-vlan-name;
    l3-interface l3-interface-name;
    mcae-mac-flush
    mcae-mac-synchronize
    private-vlan (community | isolated)
    service-id service-id;
    switch-options {
      interface interface-name {
        encapsulation-type;
        ignore-encapsulation-mismatch;
        pseudowire-status-tlv;
        static-mac mac-address {
          vlan-id vlan-id;
        }
      }
      mac-table-aging-time seconds;
      mac-table-size {
        number;
        packet-action drop;
      }
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    vlan-id (all | none | vlan-id);
    vlan members [vlan-id];
    vxlan {
      decapsulate-accept-inner-vlan;
      encapsulate-inner-vlan;
      multicast-group ;
      ovsdb-managed ;
      unreachable-vtep-aging-timer;
      vni vni-number;
    }
  }
```

Related Documentation

- Layer 2 Transparent and Switching Overview on page 2893

authentication-order (Access Profile)

Syntax authentication-order [ldap | none | password | radius | securid];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement modified in Junos OS Release 9.1.

Description Set the order in which the Junos OS tries different authentication methods when verifying that a client can access the devices. For each login attempt, the software tries the authentication methods in order, from first to last.

Options

- **ldap**—Verify the client using LDAP.
- **none**—Specify no authentication performed.
- **password**—Verify the client using the information configured at the **[edit access profile *profile-name* client *client-name*]** hierarchy level.
- **radius**—Verify the client using RADIUS authentication services.
- **securid**—Verify the client using SecurID authentication services.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Access Configuration Statement Hierarchy on page 3007

Copyright © 2016, Juniper Networks, Inc.

code-points (CoS)

Syntax code-points [*aliases*] [*bit-patterns*];

Hierarchy Level [edit class-of-service classifiers (dscp | ieee-802.1) *classifier-name* forwarding-class *class-name* loss-priority *level*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure one or more code-point aliases or bit sets to apply to a forwarding class.

NOTE: OCX Series switches do not support MPLS, so they do not support EXP code points or code point aliases.

Options ***aliases***—Name of the alias or aliases.

bit-patterns—Value of the code-point bits, in decimal form.

Required Privilege Level interfaces—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441
- Example: Configuring BA Classifiers on Transparent Mode Devices on page 2941

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Policies)

Syntax

```
destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options

address—IP address (**any**, **any-ipv4**, **any-ipv6**), IP address set, or address book entry, or wildcard address (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

domain-type (VLANs)

Syntax

```
domain-type bridge;
```

Hierarchy Level

[edit vlans *vlans-name*]

Release Information

Statement modified in Junos OS Release 9.5.

Description

Define the type of domain for a Layer 2 VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Layer 2 Bridging and Transparent Mode Overview on page 2897

Copyright © 2016, Juniper Networks, Inc.

encapsulation (Interfaces)

Syntax encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | frame-relay-port-ccc | vlan-ccc | vlan-vpls);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify logical link layer encapsulation.

Options

- **cisco-hdlc**—For normal mode (when the device is using only one B-channel). Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points
- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Physical Encapsulation on an Interface on page 2623

Copyright © 2016, Juniper Networks, Inc.

ethernet-switching

Syntax

```
ethernet-switching {
  block-non-ip-all;
  bpdu-vlan-flooding;
  bypass-non-ip-unicast;
  no-packet-flooding {
    no-trace-route;
  }
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 9.5.

Description Changes default Layer 2 forwarding behavior.

Options

- **block-non-ip-all**—Block all Layer 2 non-IP and non-ARP traffic, including multicast and broadcast traffic.
- **bpdu-vlan-flooding**—Set 802.1D bridge protocol data unit (BPDU) flooding based on VLAN.
- **bypass-non-ip-unicast**—Allow all Layer 2 non-IP traffic to pass through the device.
- **no-packet-flooding**—Stop IP flooding and send ARP or ICMP requests to discover the destination MAC address for a unicast packet.
  - **no-trace-route**—Do not send ICMP requests to discover the destination MAC address for a unicast packet. Only ARP requests are sent. This option only allows the device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnetwork as the ingress IP address.

NOTE: The **block-non-ip-all** and **bypass-non-ip-unicast** options cannot be configured at the same time.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

family inet (Interfaces)

Syntax

```
inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address (source–address/prefix) {
    arp destination-address {
      (mac mac-address | multicast-mac multicast-mac-address);
      publish publish-address;
    }
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key-value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds
      (preempt <hold-time seconds> | no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
          priority-cost value;
        }
      }
      virtual-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  dhcp {
    client-identifier {
```

```
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      user-id (ascii string| hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp arp-name;
    input input-name;
    output output-name;
  }
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
    simple-filter;
  }
  targeted-broadcast {
    (forward-and-send-to-re |forward-only);
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
}
```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement introduced in a prior release of Junos OS.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.

NOTE: You use family inet to assign an IPv4 address. You use family inet6 to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

family inet6

Syntax

```
inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      ouput;
    }
  }
  address source–address/prefix {
    eui-64;
    ndp address {
      (mac mac-address | multicast-mac multicast-mac-address);
      publish;
    }
    preferred;
    primary;
    vrrp-inet6-group group_id {
      (accept-data | no-accept-data);
      advertisements-threshold number;
      authentication-key value;
      authentication-type (md5 | simple);
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      (preempt <hold-time seconds>| no-preempt );
      priority value;
      track {
        interface interface-name {
          bandwidth-threshold value;
          priority-cost value;
        }
        priority-hold-time seconds;
        route route-address{
          routing-instance routing-instance;
        }
      }
      virtual-inet6-address [address];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  (dad-disable | no-dad-disable);
  dhcpv6-client {
    client-ia-type (ia-na | ia-pd);
    client-identifier duid-type (duid-ll | duid-llt | vendor);
```


Copyright © 2016, Juniper Networks, Inc.

```
    client-type (autoconfig | stateful);
    rapid-commit;
    req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-domain
      | sip-server |time-zone | vendor-spec);
    retransmission-attempt number;
    update-router-advertisement {
      interface interface-name;
    }
    update-server;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  nd6-stale-time seconds;
  no-neighbor-learn;
  policer {
    input input-name;
    output output-name;
  }
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
  sampling {
    input;
    output;
  }
  unnumbered-address {
    interface-name;
    preferred-source-address preferred-source-address;
  }
}
```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement supported in Junos 10.2 for SRX Series devices.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.

NOTE: You use family inet6 to assign an IPv6 address. You use family inet to assign an IPv4 address. An interface can be configured with both an IPv4 and IPv6 address.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

flow (Security Flow)

Syntax

```
flow {
    aging {
        early-ageout seconds;
        high-watermark percent;
        low-watermark percent;
    }
    allow-dns-reply;
    ethernet-switching {
        block-non-ip-all;
        bpdu-vlan-flooding;
        bypass-non-ip-unicast;
        no-packet-flooding {
            no-trace-route;
        }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
    load distribution {
        session-affinity ipsec;
    }
    pending-sess-queue-length (high | moderate | normal);
    route-change-timeout seconds;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    tcp-mss {
        all-tcp mss value;
        gre-in {
            mss value;
        }
        gre-out {
            mss value;
        }
        ipsec-vpn {
            mss value;
        }
    }
    tcp-session {
        fin-invalidate-session;
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
            (session-ageout | session-timeout seconds);
        }
    }
    traceoptions {
        file {
            filename;
            files number;
```

```
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
      packet-filter filter-name {
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
      }
      rate-limit messages-per-second;
    }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.5.

Description Determine how the device manages packet flow. The device can regulate packet flow in the following ways:

- Enable or disable DNS replies when there is no matching DNS request.
- Set the initial session-timeout values.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Understanding Flow in Logical Systems for SRX Series Devices on page 3211

Copyright © 2016, Juniper Networks, Inc.

forwarding-classes (CoS)

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- *class-name*—Display the forwarding class name assigned to the internal queue number.

  NOTE: This option is supported only on high-end SRX Series devices, including the SRX1500, SRX5400, SRX5600, and SRX5800.

  NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **policing-priority**—Layer 2 policing. One forwarding class can be configured as **premium** and others are configured as **normal**.
- **priority**—Fabric priority value:
  - **high**—Forwarding class's fabric queuing has high priority.
  - **low**—Forwarding class's fabric queuing has low priority.
- *queue-number*—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, either **high** or **low**.

  NOTE: The spu-priority option is only supported on SRX1500, and SRX5000 line devices.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring AppQoS on page 589

global-mac-table-aging-time (Protocols)

Syntax
global-mac-table-aging-time *seconds*;

Hierarchy Level
[edit protocols l2-learning]

Release Information
Statement modified in Release 9.5 of Junos OS.

Description
Configure the timeout interval for entries in the MAC table.

Default
300 seconds

Options
seconds—Time elapsed before MAC table entries are timed out and entries are deleted from the table.
Range: 10 through 64000 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- l2-learning (Protocols) on page 3063

Copyright © 2016, Juniper Networks, Inc.

global-mac-limit (Protocols)

Syntax

```
global-mac-limit limit {
  packet-action drop;
}
```

Hierarchy Level [edit protocols l2-learning]

Release Information Statement modified in Release 9.5 of Junos OS.

Description Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router.

Default 131071 MAC addresses

Options ***limit***—Number of MAC addresses that can be learned on the device.

Range: 20 through 131071 addresses

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Layer 2 Bridging and Transparent Mode for Security Devices*
- *Junos OS Layer 2 Bridging and Switching Library for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

global-mode (Protocols)

Syntax global-mode (switching | transparent-bridge) ;

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 15.1X49-D40.

Description Specify the global mode for the SRX Series device as Layer 2 transparent bridge mode or switching mode. After changing the mode, you must reboot the device for the configuration to take effect.

NOTE: Switching mode is currently not supported on all SRX Series devices.

Default transparent-bridge

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- l2-learning (Protocols) on page 3063
- Layer 2 Transparent and Switching Overview on page 2893

global-no-mac-learning (Protocols)

Syntax global-no-mac-learning;

Hierarchy Level [edit protocols l2-learning]

Release Information Statement modified in Release 9.5 of Junos OS.

Description Disable MAC learning for the entire device.

Default MAC learning is enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Layer 2 Bridging and Switching Library for Security Devices*
- *Layer 2 Bridging and Transparent Mode for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

host-inbound-traffic

Syntax

```
host-inbound-traffic {
  protocols protocol-name {
    except;
  }
  system-services service-name {
    except;
  }
}
```

Hierarchy Level

[edit security zones functional-zone management],
[edit security zones functional-zone management interfaces *interface-name*],
[edit security zones security-zone *zone-name*],
[edit security zones security-zone *zone-name* interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Control the type of traffic that can reach the device from interfaces bound to the zone.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding How to Control Inbound Traffic Based on Traffic Types on page 1055
- Understanding How to Control Inbound Traffic Based on Protocols on page 1058

Copyright © 2016, Juniper Networks, Inc.

inet6 (Security Forwarding Options)

Syntax

```
inet6 {
  mode (drop | flow-based | packet-based);
}
```

Hierarchy Level [edit security forwarding-options family]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable packet-based or flow-based processing of IPv6 traffic. By default, the device drops IPv6 traffic.

NOTE: Packet-based processing is not supported on the following SRX Series devices: SRX1500, SRX5600, and SRX5800.

Options The **mode** statement is described separately.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- family inet6 on page 2705

Copyright © 2016, Juniper Networks, Inc.

interfaces (CoS)

Syntax

```
interfaces interface-name {
  input-traffic-control-profile profile-name;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
    scheduler-map scheduler-map;
    shaping-rate bps;
    unit  logical-unit-number  {
      adaptive-shaper adaptive-shaper-name;
      classifiers {
        (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence)
      }
      forwarding-class class-name;
      input-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      loss-priority-maps {
        frame-relay-de {
          (lpmap-name | default);
        }
      }
      output-traffic-control-profile {
        profile-name;
        shared-instance shared-instance-name;
      }
      rewrite-rules {
        (dscp |dscp-ipv6 |exp |frame-relay-de |ieee-802.1 |ieee-802.1ad |inet-precedence)
      }
      scheduler-map scheduler-map-name;
        shaping-rate {
        rate;
      }
      vc-shared-scheduler;
      virtual-channel-group group-name;
    }
  }
}
```

Hierarchy Level [edit class-of-service interface *interface-name* unit *number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Associate the class-of-service configuration elements with an interface.

Options interface *interface-name* unit *number*—The user-specified interface name and unit number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

interfaces (Security Zones)

Syntax

```
interfaces interface-name {
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
}
```

Hierarchy Level [edit security zones functional-zone management],
[edit security zones security-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the set of interfaces that are part of the zone.

Options ***interface-name*** —Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Security Zones on page 1050

Copyright © 2016, Juniper Networks, Inc.

interface (Switching Options)

Syntax

```
interface interface-name {
  encapsulation-type;
  ignore-encapsulation-mismatch;
  pseudowire-status-tlv;
  static-mac mac-address {
    vlan-id vlan-id;
  }
}
```

Hierarchy Level [edit vlans *vlans-name* switch-options]

Release Information Statement modified in Junos OS Release 9.5.

Description Specify the logical interfaces to include in the bridge domain (VLAN).

Options

- ***interface-name***—Name of a logical interface.
- ***encapsulation-type***—Encapsulation type for VPN.
- ***ignore-encapsulation-mismatch***—Allow different encapsulation types on local and remote devices.
- ***pseudowire-status-tlv***—Send pseudowire status.
- ***mac-address***—Static MAC address assigned to the logical interface.
- ***vlan-id***—VLAN identifier.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding VLANs on page 2900

Copyright © 2016, Juniper Networks, Inc.

interface (VLANs)

Syntax interface *interface-name*;

Hierarchy Level [edit vlans*vlans-name*]

Release Information Statement modified in Junos OS Release 9.5.

Description Specify an interface to include the VLAN.

Options ***interface-name***—Name of the integrated routing and bridging (IRB) interface to include in the VLAN. The format of the interface name is **irb.*x***, where ***x*** is the unit number of the interface you configured at the **[edit interfaces irb]** hierarchy level.

NOTE: You can specify only one IRB interface for each VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding VLANs on page 2900

 Copyright © 2016, Juniper Networks, Inc.

l2-learning (Protocols)

Syntax

```
l2-learning {
  global-mac-limit limit {
  packet-action-drop
  }
  global-mac-table-aging-time seconds;
  global-mode (switching | transparent-bridge) ;
  global-no-mac-learning;
}
```

Hierarchy Level [edit protocols]

Release Information Statement modified in Junos OS Release 9.5. Support for global mode added in Junos OS Release 15.1X49-D40.

Description Configure Layer 2 address learning and forwarding properties globally.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- global-mac-table-aging-time (Protocols) on page 3054
- global-mac-limit (Protocols) on page 3055
- global-no-mac-learning (Protocols) on page 3056
- global-mode (Protocols) on page 3056

Copyright © 2016, Juniper Networks, Inc.

loss-priority (CoS Loss Priority)

Syntax loss-priority *level* code-points[*values*];

Hierarchy Level [edit class-of-service loss-priority-maps frame-relay-de *map-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Map CoS values to a loss priority.

Options ***level*** can be one of the following:

- **high**—Packet has high loss priority.
- **medium-high**—Packet has medium-high loss priority.
- **medium-low**—Packet has medium-low loss priority.
- **low**—Packet has low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

match (Security Policies)

Syntax

```
match {
  application {
    [application];
    any;
  }
  destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
  }
  source-identity {
    [role-name];
    any;
    authenticated-user;
    unauthenticated-user;
    unknown-user;
  }
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated with **source-identity** option in Junos OS Release 12.1.

Description Configure security policy match criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

native-vlan-id (Interfaces)

Syntax native-vlan-id*vlan-id*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.

Options ***vlan-id***—Configure a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

peer-selection-service

Syntax

```
peer-selection-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable the peer selection service process.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the peer selection service process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Building Blocks Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

pgcp-service

Syntax

```
pgcp-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the Packet Gateway Control Protocol (PGCP) that is required for the border gateway function (BGF) feature.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the Packet Gateway Control Protocol (PGCP) process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Building Blocks Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

policy (Security Policies)

Syntax

```
policy policy-name {
  description description;
  match {
    application {
      [application];
      any;
    }
    destination-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
```

Copyright © 2016, Juniper Networks, Inc.

```
            ssl-proxy {
                profile-name profile-name;
            }
            uac-policy {
                captive-portal captive-portal;
            }
            utm-policy policy-name;
        }
        destination-address {
            drop-translated;
            drop-untranslated;
        }
        firewall-authentication {
            pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                web-redirect;
            }
            user-firewall {
                access-profile profile-name;
                domain domain-name
                ssl-termination-profile profile-name;
            }
            web-authentication {
                client-match user-or-group-name;
            }
        }
        services-offload;
        tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
        }
        tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
        }
    }
    reject;
  }
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. The **services-offload** option added in Junos OS Release 11.4. Statement updated with the **source-identity** option and the **description** option added in Junos OS Release 12.1. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Define a security policy.

 Copyright © 2016, Juniper Networks, Inc.

Options ***policy-name***—Name of the security policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320
- *Building Blocks Feature Guide for Security Devices*

port (Access RADIUS)

Syntax port *port-number*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure the port number on which to contact the RADIUS server.

Options ***port-number***—Port number on which to contact the RADIUS server.

Default: 1812 (as specified in RFC 2865)

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation

- profile (Access) on page 2734

Copyright © 2016, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order [accounting-method];
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order [accounting-method];
  address-assignment pool pool-name;
  authentication-order [ldap | none | password | radius | securid];
  authorization-order [jsrc];
  client client-name {
    chap-secret chap-secret;
    client-group [ group-names ];
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
```


Copyright © 2016, Juniper Networks, Inc.

```
radius {
  accounting-server [server];
  attributes {
    exclude {
      acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
      acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
      acc-loop-cir-id [access-request | accounting-start | accounting-stop];
      accounting-authentic [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-delay-time [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-session-id [access-request];
      accounting-terminate-cause [accounting-off];
      act-data-rate-dn [access-request | accounting-start | accounting-stop];
      act-data-rate-up [access-request | accounting-start | accounting-stop];
      act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      act-interlv-delay-up [access-request | accounting-start | accounting-stop];
      att-data-rate-dn [access-request | accounting-start | accounting-stop];
      att-data-rate-up [access-request | accounting-start | accounting-stop];
      called-station-id [access-request | accounting-start | accounting-stop];
      calling-station-id [access-request | accounting-start | accounting-stop];
      class [access-request | accounting-start | accounting-stop];
      delegated-ipv6-prefix [accounting-start | accounting-stop];
      dhcp-gi-address [access-request | accounting-start | accounting-stop];
      dhcp-mac-address [access-request | accounting-start | accounting-stop];
      dhcp-options [access-request | accounting-start | accounting-stop];
      downstream-calculated-qos-rate [access-request | accounting-start |
        accounting-stop];
      dsl-forum-attributes [access-request | accounting-start | accounting-stop];
      dsl-line-state [access-request | accounting-start | accounting-stop];
      dsl-type [access-request | accounting-start | accounting-stop];
      dynamic-iflset-name [accounting-start | accounting-stop];
      event-time-stamp [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      framed-interface-id [access-request | accounting-start | accounting-stop];
      framed-ip-address [access-request | accounting-start | accounting-stop];
      framed-ip-netmask [access-request | accounting-start | accounting-stop];
      framed-ip-route [access-request | accounting-start | accounting-stop];
      framed-ipv6-pool [accounting-start | accounting-stop];
      framed-ipv6-prefix [accounting-start | accounting-stop];
      framed-ipv6-route [accounting-start | accounting-stop];
      framed-pool [accounting-start | accounting-stop];
      input-filter [accounting-start | accounting-stop];
      input-gigapackets [accounting-stop];
      input-gigawords [accounting-stop];
      input-ipv6-gigawords [accounting-stop];
      input-ipv6-octets [accounting-stop];
      input-ipv6-packets [accounting-stop];
      interface-description [access-request | accounting-start | accounting-stop];
      l2c-downstream-data [access-request | accounting-start | accounting-stop];
      l2c-upstream-data [access-request | accounting-start | accounting-stop];
      max-data-rate-dn [access-request | accounting-start | accounting-stop];
      max-data-rate-up [access-request | accounting-start | accounting-stop];
      max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      max-interlv-delay-up [access-request | accounting-start | accounting-stop];
      min-data-rate-dn [access-request | accounting-start | accounting-stop];
```

Copyright © 2016, Juniper Networks, Inc.

```
            min-data-rate-up [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
            nas-identifier [access-request | accounting-start | accounting-stop];
            nas-port [access-request | accounting-off | accounting-on | accounting-start |
              accounting-stop];
            nas-port-id [access-request | accounting-start | accounting-stop];
            nas-port-type [access-request | accounting-start | accounting-stop];
            output-filter [accounting-start | accounting-stop];
            output-gigapackets [accounting-stop];
            output-gigawords [accounting-stop];
            output-ipv6-gigawords [accounting-stop];
            output-ipv6-octets [accounting-stop];
            output-ipv6-packets [accounting-stop];
            upstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];
          }
          ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
          }
        }
      authentication-server [server];
      radius-options {
        request-rate number;
        revert-interval seconds;
      }
      radius-server server-address {
        accounting-port port-number
        max-outstanding-requests number-of--outstanding-requests;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
      }
      service {
        accounting-order {
          activation-protocol;
          radius;
        }
      }
      session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
      }
    }
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Copyright © 2016, Juniper Networks, Inc.

Description Create a profile containing a set of attributes that define device management access.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441
- Understanding User Authentication for Security Devices on page 5173
- Layer 2 Transparent and Switching Overview on page 2893

radius-server (Access)

Syntax

```
radius-server server-address {
  port port-number;
  retry attempts;
  routing-instance routing-instance-name;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

Hierarchy Level
[edit access],
[edit access profile *profile-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure RADIUS for Layer 2 Tunneling Protocol (L2TP) or Point-to-Point Protocol (PPP) authentication.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options ***server-address***—Address of the RADIUS authentication server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

redundancy-group (Interfaces)

Syntax redundancy-group *number*;

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the redundancy group that a redundant Ethernet interface belongs to.

Options ***number*** —Number of the redundancy group that the redundant interface belongs to. Failover properties of the interface are inherited from the redundancy group.

Range: 1 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441

secure-wire

Syntax secure-wire *secure-wire-name* interface [*interface-name-1 interface-name-2*];

Hierarchy Level [edit security forwarding-options]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure mapping of interfaces through which traffic is forwarded unchanged.

Options **secure-wire**—Specify a name for the secure wire interface mapping.

interface—Specify a pair of peer logical interfaces that constitutes the secure wire mapping.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- *Building Blocks Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

security-zone

Syntax

```
security-zone zone-name {
  address-book {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
  }
  application-tracking;
  description text;
  host-inbound-traffic {
    protocols protocol-name {
      except;
    }
    system-services service-name {
      except;
    }
  }
  interfaces interface-name {
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
  }
  screen screen-name;
  tcp-rst;
}
```

Hierarchy Level [edit security zones]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description Define a security zone, which allows you to divide the network into different segments and apply different security options to each segment.

Copyright © 2016, Juniper Networks, Inc.

Options

zone-name —Name of the security zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [edit security zones] Hierarchy Level on page 324
- Security Zones and Interfaces Overview on page 1049
- Example: Configuring Application Firewall Rule Sets Within a Security Policy on page 560

Copyright © 2016, Juniper Networks, Inc.

shaping-rate (CoS Interfaces)

Syntax shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

Logical and physical interface traffic shaping is mutually exclusive. This means you can include the **shaping-rate** statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level or the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, but not both.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

Default If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces interface *interface-name*]** hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options *rate*—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: For logical interfaces, 1000 through 32,000,000,000 bps.

For physical interfaces, 1000 through 160,000,000,000 bps.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Large Delay Buffer on a Channelized T1 Interface on page 1493
- Understanding Ethernet Interfaces on page 2541

Copyright © 2016, Juniper Networks, Inc.

source-address (Access RADIUS)

Syntax source-address *source-address*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options *source-address*—Valid IP address configured on one of the device interfaces.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

source-address (Security Policies)

Syntax

```
source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations (in addition to the existing support of active/passive chassis cluster configurations) added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1.

Description

Define the matching criteria. You can specify one or more IP addresses, address sets, or wildcard addresses. You can specify wildcards **any**, **any-ipv4**, or **any-ipv6**.

Options

address—IP addresses, address sets, or wildcard addresses (represented as A.B.C.D/wildcard-mask). You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

static-mac (VLANs)

Syntax

```
static-mac mac-address {
    vlan-id vlan-id;
}
```

Hierarchy Level [edit vlans*vlan--name* switch-options interface *interface-name*]

Release Information Statement modified in Junos OS Release 9.5.

Description Configure a static MAC address for a logical interface in a VLAN.

Options

- ***mac-address***—MAC address
- ***vlan-id***—VLAN identifier

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding VLANs on page 2900

Copyright © 2016, Juniper Networks, Inc.

switch-options (VLANs)

Syntax

```
switch-options {
  interface interface-name {
    encapsulation-type;
    ignore-encapsulation-mismatch;
    pseudowire-status-tlv;
    static-mac mac-address {
      vlan-id vlan-id;
    }
  }
  mac-table-aging-time seconds;
  mac-table-size {
    number;
    packet-action drop;
  }
}
```

Hierarchy Level [edit vlans *vlans-name*]

Release Information Statement modified in Junos OS Release 9.5.

Description Configure Layer 2 learning and forwarding properties for a VLAN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Layer 2 Transparent and Switching Overview on page 2893

Copyright © 2016, Juniper Networks, Inc.

system-services (Security Zones Interfaces)

Syntax

```
system-services service-name {
  except;
}
```

Hierarchy Level [edit security zones security-zone *zone-name* interfaces *interface-name* host-inbound-traffic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the types of traffic that can reach the device on a particular interface.

Options

- *service-name* —Service for which traffic is allowed. The following services are supported:
  - **all**—Enable all possible system services available on the Routing Engine (RE).
  - **any-service**—Enable services on entire port range.
  - **bootp**—Enable traffic destined to BOOTP and DHCP relay agents.
  - **dhcp**—Enable incoming DHCP requests.
  - **dhcpv6**—Enable incoming DHCP requests for IPv6.
  - **dns**—Enable incoming DNS services.
  - **finger**—Enable incoming finger traffic.
  - **ftp**—Enable incoming FTP traffic.
  - **http**—Enable incoming J-Web or clear-text Web authentication traffic.
  - **https**—Enable incoming J-Web or Web authentication traffic over Secure Sockets Layer (SSL).
  - **ident-reset**—Enable the access that has been blocked by an unacknowledged identification request.
  - **ike**—Enable Internet Key Exchange traffic.
  - **netconf SSH**—Enable incoming NetScreen Security Manager (NSM) traffic over SSH.
  - **ntp**—Enable incoming Network Time Protocol (NTP) traffic.
  - **ping**—Allow the device to respond to ICMP echo requests.
  - **r2cp**—Enable incoming Radio Router Control Protocol traffic.
  - **reverse-ssh**—Reverse SSH traffic.
  - **reverse-telnet**—Reverse Telnet traffic.
  - **rlogin**—Enable incoming **rlogin** (remote login) traffic.
  - **rpm**—Enable incoming real-time performance monitoring (RPM) traffic.
  - **rsh**—Enable incoming Remote Shell (**rsh**) traffic.
  - **snmp**—Enable incoming SNMP traffic (UDP port 161).

Copyright © 2016, Juniper Networks, Inc.

- **snmp-trap**—Enable incoming SNMP traps (UDP port 162).
- **ssh**—Enable incoming SSH traffic.
- **telnet**—Enable incoming Telnet traffic.
- **tftp**—Enable TFTP services.
- **traceroute**—Enable incoming traceroute traffic (UDP port 33434).
- **xnm-clear-text**—Enable incoming Junos XML protocol traffic for all specified interfaces.
- **xnm-ssl**— Enable incoming Junos XML protocol-over-SSL traffic for all specified interfaces.

- **except**—(Optional) except can only be used if all has been defined.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

unframed | no-unframed (Interfaces)

Syntax
(unframed | no-unframed);

Hierarchy Level
[edit interfaces *interface-name* t3-options]

Release Information
Statement introduced in Junos OS Release 11.1.

Description
Enable or disable framing for the T3 interface on a 1-Port Clear Channel DS3/E3 GPIM on an SRX Series device. By default, unframed mode is enabled. Select no-unframed to enable framing. Select unframed to return to the default mode.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring a T3 Interface*

Copyright © 2016, Juniper Networks, Inc.

vlans

Syntax

```
vlans vlans-name {
  description  text;
  domain-type bridge;
  forwarding-options {
    dhcp-relay {
      active-server-group active-server-group-name;
      dhcpv6 {
        active-server-group active-server-group-name;
        authentication {
          password password;
          username-include {
            circuit-type;
            client-id;
            delimiter delimiter;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
          }
        }
        dynamic-profile (dynamic-profile-name | junos-default-profile) {
          aggregate-clients (merge | replace);
          use-primary ( primary-profile-name | junos-default-profile);
        }
        group group-name {
          active-server-group active-server-group-name;
          authentication {
            password password;
            username-include {
              circuit-type;
              client-id;
              delimiter delimiter;
              domain-name domain-name;
              interface-name;
              logical-system-name;
              relay-agent-interface-id;
              relay-agent-remote-id;
              relay-agent-subscriber-id;
              routing-instance-name;
              user-prefix user-prefix;
            }
          }
          dynamic-profile (dynamic-profile-name | junos-default-profile) {
            aggregate-clients (merge | replace) {
              use-primary (primary-profile-name | junos-default-profile);
            }
          }
          interface interface-name {
```


Copyright © 2016, Juniper Networks, Inc.

```
    dynamic-profile (dynamic-profile-name| junos-default-profile) {
      aggregate-clients (merge | replace) {
        use-primary (primary-profile-name | junos-default-profile);
      }
    }
    exclude;
    overrides {
      (allow-snooped-clients | no-allow-snooped-clients);
      interface-client-limit number;
      no-bind-on-request;
      send-release-on-delete;
    }
    service-profile profile-name;
    trace;
    upto upto-interface-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        detection-time {
          threshold milliseconds;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        session-mode (automatic | multihop | singlehop);
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (0 | 1 | automatic);
      }
    }
  }
  overrides {
    (allow-snooped-clients | no-allow-snooped-clients);
    interface-client-limit number;
    no-bind-on-request;
    send-release-on-delete;
  }
  relay-agent-interface-id {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical):
  }
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
```

```
            bfd {
              detection-time {
                threshold milliseconds;
              }
              holddown-interval milliseconds;
              minimum-interval milliseconds;
              minimum-receive-interval milliseconds;
              multiplier number;
              no-adaptation;
              session-mode (automatic | multihop | singlehop);
              transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
              }
              version (0 | 1 | automatic);
            }
          }
        }
        overrides {
          (allow-snooped-clients | no-allow-snooped-clients);
          interface-client-limit number;
          no-bind-on-request;
          send-release-on-delete;
        }
        relay-agent-interface-id {
          prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
          }
          use-interface-description (device | logical):
        }
        server-group {
          server-group-name {
            server-ip-address;
          }
        }
        service-profile service-profile-name;
      }
      group group-name {
        active-server-group server-group-name;
        interface interface-name {
          exclude;
          upto interface-name;
        }
        relay-option-60 {
          vendor-option {
            default-local-server-group local-server-group-name;
            default-relay-server-group server-group-name;
            drop;
            equals {
              ascii ascii-name {
                drop;
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
```


Copyright © 2016, Juniper Networks, Inc.

```
            hexadecimal hexadecimal {
              drop;
              local-server-group local-server-group-name;
              relay-server-group server-group-name;
            }
          }
          starts-with {
            ascii ascii-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group server-group-name;
            }
            hexadecimalhexadecimal {
              drop;
              local-server-group local-server-group-name;
              relay-server-group server-group-name;
            }
          }
        }
      }
      relay-option-82 {
        circuit-id {
          prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
          }
          use-interface-description (device | logical);
        }
      }
    }
    relay-option-60 {
      vendor-option {
        default-local-server-group local-server-group-name;
        default-relay-server-group server-group-name;
        drop;
        equals {
          ascii ascii-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group server-group-name;
          }
          hexadecimal hexadecimal {
            drop;
            local-server-group local-server-group-name;
            relay-server-group server-group-name;
          }
        }
        starts-with {
          ascii ascii-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group server-group-name;
          }
          hexadecimal hexadecimal {
            drop;
```

```
                local-server-group local-server-group-name;
                relay-server-group server-group-name;
              }
            }
          }
        }
        relay-option-82 {
          circuit-id {
            prefix {
              host-name;
              logical-system-name;
              routing-instance-name;
            }
            use-interface-description (device | logical);
          }
        }
        server-group server-group-name {
          ip-address;
        }
      }
      filter {
        input input-filter-name;
      }
      flood {
        input input-filter-name;
      }
    }
    interface interface-name;
    isolated-vlan isolated-vlan-name;
    l3-interface l3-interface-name;
    mcae-mac-flush
    mcae-mac-synchronize
    private-vlan (community | isolated)
    service-id service-id;
    switch-options {
      interface interface-name {
        encapsulation-type;
        ignore-encapsulation-mismatch;
        pseudowire-status-tlv;
        static-mac mac-address {
          vlan-id vlan-id;
        }
      }
      mac-table-aging-time seconds;
      mac-table-size {
        number;
        packet-action drop;
      }
    }
    vlan-id (all | none | vlan-id);
    vlan members [vlan-id];
    vxlan {
      decapsulate-accept-inner-vlan;
      encapsulate-inner-vlan;
      multicast-group ;
      ovsdb-managed ;
```


Copyright © 2016, Juniper Networks, Inc.

```
        unreachable-vtep-aging-timer;
        vni vni-number;
    }
}
```

Hierarchy Level [edit]

Release Information Statement modified in Junos OS Release 9.5.

Description Configure a domain that includes a set of logical ports that share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Options *vlans-name*—Name of the VLAN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Layer 2 Transparent and Switching Overview on page 2893

Copyright © 2016, Juniper Networks, Inc.

vlan-id (VLAN)

Syntax vlan-id (all | none | *number*);

Hierarchy Level [edit vlans *vlan-name*],
[edit logical-systems *logical-system-name* vlans*vlan-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vlans*vlan-name*],
[edit routing-instances *routing-instance-name* vlans *vlan--name*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for Layer 2 trunk ports added in Junos OS Release 9.2.
Support for SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify a VLAN identifier (VID) to include in the packets sent to and from the VLAN, or a VPLS routing instance.

NOTE: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

Options *number*—A valid VLAN identifier. If you configure multiple VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.

NOTE: If you specify a VLAN identifier, you cannot also use the all option. They are mutually exclusive.

all—Specify that the VLAN spans all the VLAN identifiers configured on the member logical interfaces.

NOTE: You cannot specify the all option if you include a routing interface in the VLAN.

none—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Multichassis link aggregation (MC-LAG) does not support the none option with the vlan-id statement with VLANs.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring VLANs on page 2901
- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 3360
- vlans on page 3086

vlan members (VLANs)

Syntax
vlan members [*vlan-id*];

Hierarchy Level
[edit vlans *vlan-name*]

Release Information
Statement modified in Junos OS Release 9.5.

Description
Specify multiple VLAN identifiers to create a VLAN for each VLAN identifier.

Options
vlan-id—A list of valid VLAN identifiers. A VLAN is created for each VLAN identifier in the list.

NOTE: If you specify a VLAN identifier list, you cannot configure an IRB interface in the bridge VLAN.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring VLANs on page 2901

Copyright © 2016, Juniper Networks, Inc.

vlan-tagging (Interfaces)

Syntax vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level [edit interfaces *interface*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.

Options **native-vlan-id**—Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

NOTE: The native-vlan-id can be configured only when either flexible-vlan-tagging mode or interface-mode trunk is configured.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Tagging on page 2485

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 141

Operational Commands

- clear oam ethernet connectivity-fault-management path-database
- clear oam ethernet connectivity-fault-management statistics
- clear security flow ip-action
- clear security flow session family
- show ethernet-switching mac-learning-log (View)
- show ethernet-switching table (View)
- show igmp-snooping route (View)
- show igmp-snooping vlans (View)
- show interfaces (SRX Series)
- show oam ethernet connectivity-fault-management adjacencies
- show oam ethernet connectivity-fault-management forwarding-state
- show oam ethernet connectivity-fault-management interfaces
- show oam ethernet connectivity-fault-management mep-database
- show oam ethernet connectivity-fault-management mep-statistics
- show oam ethernet connectivity-fault-management mip
- show oam ethernet connectivity-fault-management path-database
- show oam ethernet connectivity-fault-management routes
- show oam ethernet link-fault-management
- show security flow gate family
- show security flow ip-action
- show security flow session family
- show security flow statistics
- show security flow status
- show security forward-options secure-wire
- show security policies
- show security zones
- show vlans

Copyright © 2016, Juniper Networks, Inc.

clear oam ethernet connectivity-fault-management path-database

Syntax clear oam ethernet connectivity-fault-management path-database maintenance-domain *md-name* maintenance-association *ma-name* host <*mac-addr*>

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Clear the relevant path information from the database for the specified remote host.

Options **host**—(Optional) MAC address of remote host in xx:xx:xx:xx:xx:xx format.

maintenance-association —Name of the maintenance association.

maintenance-domain —Name of the maintenance domain.

Required Privilege Level clear

Related Documentation
- show oam ethernet connectivity-fault-management path-database on page 3156

List of Sample Output clear oam ethernet connectivity-fault- management path-database on page 3096

Sample Output

clear oam ethernet connectivity-fault- management path-database

```
user@host> clear oam ethernet connectivity-fault-management path-database
maintenance-domain private maintenance-association private-ma
Path database entries cleared for the remote-host
```


Copyright © 2016, Juniper Networks, Inc.

clear oam ethernet connectivity-fault-management statistics

Syntax
clear oam ethernet connectivity-fault-management statistics
interface
level

Release Information
Statement introduced in Junos OS Release 12.1X44-D10.

Description
Clear connectivity-fault-management statistics.

Options
Interface—Clear the statistics on an interface.

Level—The maintenance-domain level (0 through 7).

Required Privilege Level
View

Related Documentation
- show oam ethernet connectivity-fault-management mep-statistics on page 3152

List of Sample Output
clear oam ethernet connectivity-fault- management statistics on page 3097

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault- management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow ip-action

Syntax **clear security flow ip-action** [*filter*]

Release Information Command introduced in Junos OS Release 10.4. Logical systems option introduced in Junos OS Release 11.2.

Description Clear IP-action entries, based on filtered options, for IP sessions running on the device.

Options *filter*—Filter the display based on the specified criteria.

The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

all | [*filter*]—All active sessions on the device.

destination-port ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535.

destination-prefix ***destination-prefix***—Destination IP prefix or address.

family (inet | inet6) [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

logical-system ***logical-system-name*** **| all** [*filter*]—Specified logical system or all logical systems.

protocol ***protocol-name*** **|** ***protocol-number*** [*filter*]—Protocol name or number and filtered options.

- **ah** or **51**
- **egp** or **8**
- **esp** or **50**
- **gre** or **47**
- **icmp** or **1**
- **icmp6** or **58**
- **igmp** or **2**
- **ipip** or **4**
- **ospf** or **89**
- **pim** or **103**
- **rsvp** or **46**
- **sctp** or **132**
- **tcp** or **6**
- **udp** or **17**

Copyright © 2016, Juniper Networks, Inc.

root-logical-system [*filter*]—Default logical system information and filtered options.

source-port ***source-port***—Source port number of the traffic. Range is 1 through 65,535.

source-prefix ***source-prefix***—Source IP prefix or address of the traffic.

Required Privilege Level clear

Related Documentation
- show security flow ip-action on page 2051

List of Sample Output
clear security flow ip-action all on page 3099
clear security flow ip-action destination-prefix on page 3099
clear security flow ip-action family inet on page 3099
clear security flow ip-action protocol udp on page 3099

Output Fields When you enter this command, the system responds with the status of your request.

Sample Output

clear security flow ip-action all

```
user@host>clear security flow ip-action all
1008 ip-action entries cleared
```

clear security flow ip-action destination-prefix

```
user@host>clear security flow ip-action destination-prefix 192.0.2.5/24
87 ip-action entries cleared
```

clear security flow ip-action family inet

```
user@host>clear security flow ip-action family inet
2479 ip-action entries cleared
```

clear security flow ip-action protocol udp

```
user@host>clear security flow ip-action protocol udp
270 ip-action entries cleared
```

Copyright © 2016, Juniper Networks, Inc.

clear security flow session family

Syntax clear security flow session family (inet | inet6)

Release Information Command introduced in Junos OS Release 10.2.

Description Clear sessions that match the specified protocol family.

Options
- **inet**—Clear IPv4 sessions.
- **inet6**—Clear IPv6 sessions.

Required Privilege Level clear

Related Documentation
- show security flow session family on page 2103

List of Sample Output
clear security flow session family inet on page 3100
clear security flow session family inet6 on page 3100

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security flow session family inet

```
user@host> clear security flow session family inet
1 active sessions cleared
```

clear security flow session family inet6

```
user@host> clear security flow session family inet6
1 active sessions cleared
```


Copyright © 2016, Juniper Networks, Inc.

show ethernet-switching mac-learning-log (View)

Syntax **show ethernet-switching mac-learning-log**

Release Information Command introduced in Junos OS Release 9.5.

Description Displays the event log of learned MAC addresses.

Required Privilege Level view

Related Documentation
- Understanding MAC Limiting on page 2999

Output Fields Table 279 lists the output fields for the show ethernet-switching mac-learning-log command. Output fields are listed in the approximate order in which they appear.

Table 310: show interfaces Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
VLAN-IDX	VLAN index. An internal value assigned by Junos OS for each VLAN.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log

```
user@host> show ethernet-switching mac-learning-log
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
```

Copyright © 2016, Juniper Networks, Inc.

```
vlan_idx 15 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:05:00:00:05 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:30:48:90:54:89 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5e:00:01:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5e:00:01:08 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5e:00:01:09 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:19:e2:50:ac:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5e:00:01:04 was learned
[output truncated]
```


Copyright © 2016, Juniper Networks, Inc.

show ethernet-switching table (View)

Syntax **show ethernet-switching table (brief |detail |extensive) interface** *interface-name*

Release Information Command introduced in Junos OS Release 9.5.

Description Displays the Ethernet switching table.

Options
- **none**—(Optional) Display brief information about the Ethernet switching table.
- **brief** | **detail** | **extensive**—(Optional) Display the specified level of output.
- **interface-name**—(Optional) Display the Ethernet switching table for a specific interface.

Required Privilege Level view

Related Documentation
- Port Security Overview on page 2999
- Understanding MAC Limiting on page 2999

Output Fields Table 280 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 311: show ethernet-switching table Output Fields

Field Name	Field Description
VLAN	The name of a VLAN.
MAC address	The MAC address associated with the VLAN.
Type	The type of MAC address. Values are: • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.

Sample Output

show ethernet-switching table

```
user@host> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
```

Copyright © 2016, Juniper Networks, Inc.

```
F2 * Flood - All-members
F2 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
[output truncated]
```

Sample Output

show ethernet-switching table brief

```
user@host> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
F2 * Flood - All-members
F2 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
```


Copyright © 2016, Juniper Networks, Inc.

```
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
[output truncated]
```

Sample Output

show ethernet-switching table detail

```
user@host> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
```

Copyright © 2016, Juniper Networks, Inc.

```
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show ethernet-switching table interface ge-0/0/1

```
user@host> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN     MAC address         Type     Age Interfaces
V1       *                   Flood    - All-members
V1       00:00:05:00:00:05   Learn    0 ge-0/0/1.0
```

show igmp-snooping route (View)

Syntax show igmp-snooping route (brief | detail | ethernet-switching | inet | vlan)

Release Information Command introduced in Junos OS Release 9.5.

Description Display IGMP snooping route information.

Options

- **none**—Display general parameters.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **ethernet-switching**—(Optional) Display Ethernet switching information.
- **inet**—(Optional) Display inet information.
- **vlan *vlan-id* |*vlan-name***—(Optional) Display route information for the specified VLAN.

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

Output Fields Table 281 lists the output fields for the **show igmp-snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 312: show igmp-snooping route Output Fields

Field Name	Field Description
VLAN	Name of the VLAN.
Group	Multicast group address.
Next-hop	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@host> show igmp-snooping route
VLAN       Group          Next-hop
v11         224.1.1.1, *   533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12         224.1.1.3, *   534
Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@host> show igmp-snooping route vlan v1
Table: 0
VLAN      Group              Next-hop
v1        224.1.1.1, *       1266
Interfaces: ge-0/0/0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
v1        224.1.1.3, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.5, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.7, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.9, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.11, *      1266
Interfaces: ge-0/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

show igmp-snooping vlans (View)

Syntax **show igmp-snooping vlans**
<**brief** | **detail** >
<vlan *vlan-id* | *vlan-name* >

Release Information Command introduced in Junos OS Release 9.5.

Description Display IGMP snooping VLAN information.

Options
- **none**—Display general parameters.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **vlan** *vlan-id* | *vlan-name* —(Optional) Display VLAN information for the specified VLAN.

Required Privilege Level view

Related Documentation
- show igmp-snooping route (View) on page 2813

Output Fields lists the output fields for the **show igmp-snooping vlans** command. Output fields are listed in the approximate order in which they appear.

Table 313: show igmp-snooping vlans

Field Name	Field Description
VLAN	Name of the VLAN.
Interfaces	Number of interfaces in the VLAN.
Groups	Number of groups in the VLAN.
MRouters	Number of multicast routers associated with the VLAN.
Receivers	Number of host receivers in the VLAN.
Tag	Numerical identifier of the VLAN.
vlan-interface	Internal VLAN interface identifier.
Membership timeout	Membership timeout value.
Querier timeout	Timeout value for interfaces dynamically marked as router interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface.
Interface	Name of the interface.
Reporters	Number of dynamic groups on an interface.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show igmp-snooping vlans

```
user@host> show igmp-snooping vlans
VLAN  Interfaces  Groups  MRouters  Receivers
default  0           0       0         0
v1       11          50      0         0
v10      1           0       0         0
v11      1           0       0         0
v180     3           0       1         0
v181     3           0       0         0
v182     3           0       0         0
```

Sample Output

show igmp-snooping vlans vlan v10

```
user@host> show igmp-snooping vlans vlan v10
VLAN    Interfaces   Groups  MRouters   Receivers
v10     1            0       0          0
```

Sample Output

show igmp-snooping vlans vlan v10 detail

```
user@host> sshow igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
Membership timeout: 260, Querier timeout: 255
Interface: ge-0/0/10.0, tagged, Groups: 0, Reporters: 0
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces (SRX Series)

Syntax

```
show interfaces {
  <brief | detail | extensive | terse>
  controller interface-name
  descriptions interface-name
  destination-class (all | destination-class-name logical-interface-name)
  diagnostics optics interface-name
  far-end-interval interface-fpc/pic/port
  filters interface-name
  flow-statistics interface-name
  interval interface-name
  load-balancing (detail | interface-name)
  mac-database mac-address mac-address
  mc-ae id identifier unit number revertive-info
  media interface-name
  policers interface-name
  queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics
  redundancy (detail | interface-name)
  routing brief detail summary interface-name
  routing-instance (all | instance-name)
  snmp-index snmp-index
  source-class (all | destination-class-name logical-interface-name)
  statistics interface-name
  switch-port switch-port number
  transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all |
    interface-name)
  zone interface-name
}
```

Release Information Command modified in Junos OS Release 9.5.

Description Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

On SRX Series appliance, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-** ***pim*****/0/*****port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim*/0/*port***—E1 interface.
  - **e3-*pim*/0/*port***—E3 interface.

Copyright © 2016, Juniper Networks, Inc.

  - **fe-*pim*/0/*port***—Fast Ethernet interface.
  - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
  - **se-*pim*/0/*port***—Serial interface.
  - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
  - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
  - **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

- **brief | detail | extensive | terse**—(Optional) Display the specified level of output.
- **controller**—(Optional) Show controller information.
- **descriptions**—(Optional) Display interface description strings.
- **destination-class**—(Optional) Show statistics for destination class.
- **diagnostics**—(Optional) Show interface diagnostics information.
- **far-end-interval**—(Optional) Show far end interval statistics.
- **filters**—(Optional) Show interface filters information.
- **flow-statistics**—(Optional) Show security flow counters and errors.
- **interval**—(Optional) Show interval statistics.
- **load-balancing**—(Optional) Show load-balancing status.
- **mac-database**—(Optional) Show media access control database information.
- **mc-ae**—(Optional) Show MC-AE configured interface information.
- **media**—(Optional) Display media information.
- **policers**—(Optional) Show interface policers information.
- **queue**—(Optional) Show queue statistics for this interface.
- **redundancy**—(Optional) Show redundancy status.
- **routing**—(Optional) Show routing status.
- **routing-instance**—(Optional) Name of routing instance.
- **snmp-index**—(Optional) SNMP index of interface.
- **source-class**—(Optional) Show statistics for source class.
- **statistics**—(Optional) Display statistics and detailed output.
- **switch-port**—(Optional) Front end port number (0..15).
- **transport**—(Optional) Show interface transport information.
- **zone**—(Optional) Interface's zone.

Required Privilege Level view

Related Documentation

- Understanding Interfaces on page 2441

List of Sample Output

show interfaces Gigabit Ethernet on page 3121
show interfaces brief (Gigabit Ethernet) on page 3122
show interfaces detail (Gigabit Ethernet) on page 3122
show interfaces extensive (Gigabit Ethernet) on page 3124
show interfaces terse on page 3127
show interfaces controller (Channelized E1 IQ with Logical E1) on page 3127
show interfaces controller (Channelized E1 IQ with Logical DS0) on page 3127
show interfaces descriptions on page 3128
show interfaces destination-class all on page 3128
show interfaces diagnostics optics on page 3128
show interfaces far-end-interval coc12-5/2/0 on page 3129
show interfaces far-end-interval coc1-5/2/1:1 on page 3129
show interfaces filters on page 3130
show interfaces flow-statistics (Gigabit Ethernet) on page 3130
show interfaces interval (Channelized OC12) on page 3131
show interfaces interval (E3) on page 3131
show interfaces interval (SONET/SDH) on page 3132
show interfaces load-balancing on page 3132
show interfaces load-balancing detail on page 3132
show interfaces mac-database (All MAC Addresses on a Port) on page 3133
show interfaces mac-database (All MAC Addresses on a Service) on page 3133
show interfaces mac-database mac-address on page 3134
show interfaces mc-ae on page 3134
show interfaces media (SONET/SDH) on page 3134
show interfaces policers on page 3135
show interfaces policers interface-name on page 3135
show interfaces queue on page 3135
show interfaces redundancy on page 3136
show interfaces redundancy (Aggregated Ethernet) on page 3136
show interfaces redundancy detail on page 3137
show interfaces routing brief on page 3137
show interfaces routing detail on page 3137
show interfaces routing-instance all on page 3138
show interfaces snmp-index on page 3138
show interfaces source-class all on page 3138
show interfaces statistics (Fast Ethernet) on page 3139
show interfaces switch-port on page 3139
show interfaces transport pm on page 3140
show security zones on page 3141

Output Fields

Table 194 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—There are two possible conditions regarding the number of oversized frames: ◦ Packet length exceeds 1518 octets, or ◦ Packet length exceeds MRU • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 314: show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
```

Copyright © 2016, Juniper Networks, Inc.

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

Sample Output

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
```


Copyright © 2016, Juniper Networks, Inc.

```
Disabled,
 Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
 Device flags   : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags     : None
 CoS queues     : 8 supported, 8 maximum usable queues
 Hold-times     : Up 0 ms, Down 0 ms
 Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
 Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
 Statistics last cleared: Never
 Traffic statistics:
  Input  bytes  :                   0                   0 bps
  Output bytes  :                   0                   0 bps
  Input  packets:                   0                   0 pps
  Output packets:                   0                   0 pps
 Egress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort                   0                   0                   0

   1 expedited-fo                  0                   0                   0

   2 assured-forw                  0                   0                   0

   3 network-cont                  0                   0                   0

 Queue number:         Mapped forwarding classes
   0                   best-effort
   1                   expedited-forwarding
   2                   assured-forwarding
   3                   network-control
 Active alarms  : LINK
 Active defects : LINK
 Interface transmit statistics: Disabled

 Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
   Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Transit statistics:
    Input  bytes  :                   0                   0 bps
    Output bytes  :                   0                   0 bps
    Input  packets:                   0                   0 pps
    Output packets:                   0                   0 pps
   Security: Zone: public
   Flow Statistics :
   Flow Input statistics :
     Self packets :                    0
     ICMP packets :                    0
     VPN packets :                     0
     Multicast packets :               0
     Bytes permitted by policy :       0
     Connections established :         0
```

```
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

Sample Output

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
```


Copyright © 2016, Juniper Networks, Inc.

```
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %           bps     %           usec
    0 best-effort            95     950000000    95              0      low
none
    3 network-control         5      50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
```

```
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Local statistics:
 Input  bytes  :                   0
 Output bytes  :                   0
 Input  packets:                   0
 Output packets:                   0
Transit statistics:
 Input  bytes  :                   0                   0 bps
 Output bytes  :                   0                   0 bps
 Input  packets:                   0                   0 pps
 Output packets:                   0                   0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                    0
  ICMP packets :                    0
  VPN packets :                     0
  Multicast packets :               0
  Bytes permitted by policy :       0
  Connections established :         0
Flow Output statistics:
  Multicast packets :               0
  Bytes permitted by policy :       0
Flow error statistics (Packets dropped due to):
  Address spoofing:                 0
  Authentication failed:            0
  Incoming NAT errors:              0
  Invalid zone received packet:     0
  Multiple user authentications:    0
  Multiple incoming NAT:            0
  No parent for a gate:             0
  No one interested in self packets: 0
  No minor session:                 0
  No more sessions:                 0
  No NAT gate:                      0
  No route present:                 0
  No SA for incoming SPI:           0
  No tunnel found:                  0
  No session for a gate:            0
  No zone or NULL zone binding      0
  Policy denied:                    0
  Security association not active:  0
  TCP sequence number out of window: 0
  Syn-attack protection:            0
  User authentication errors:       0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
    Generation: 150
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1              --> 0/0
                                                 10.0.0.16             --> 0/0
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

Sample Output

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                      Admin Link
ce1-1/2/6                                                       up    up
e1-1/2/6                                                        up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                      Admin Link
ce1-1/2/3                                                       up    up

ds-1/2/3:1                                                      up    up

ds-1/2/3:2                                                      up    up
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

Sample Output

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                          Packets                 Bytes
     Destination class            (packet-per-second)    (bits-per-second)
                        gold                        0                     0
                                                 (                  0) (                  0)
                      silver                        0                     0
                                                 (                  0) (                  0)
  Logical interface so-0/1/3.0
                                          Packets                 Bytes
     Destination class            (packet-per-second)    (bits-per-second)
                        gold                        0                     0
                                                 (                  0) (                  0)
                      silver                        0                     0
                                                 (                  0) (                  0)
```

Sample Output

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                         :  7.408 mA
    Laser output power                         :  0.3500 mW / -4.56 dBm
    Module temperature                         :  23 degrees C / 73 degrees F
    Module voltage                             :  3.3450 V
    Receiver signal average optical power      :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Laser output power high warning            :  Off
    Laser output power low warning             :  Off
```


Copyright © 2016, Juniper Networks, Inc.

```
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

Sample Output

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
```


Copyright © 2016, Juniper Networks, Inc.

```
    Bytes permitted by policy :        3478
    Connections established :          1
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        16994
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
```

Copyright © 2016, Juniper Networks, Inc.

```
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State           Last change  Member count
ams0       Up              1d 00:50     2
ams1       Up              00:00:59     2
```

show interfaces load-balancing detail

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1             56              0              0
  00:00:c0:01:01:02       7023810      323095260              0              0
  00:00:c0:01:01:03       7023810      323095260              0              0
  00:00:c0:01:01:04       7023810      323095260              0              0
  00:00:c0:01:01:05       7023810      323095260              0              0
  00:00:c0:01:01:06       7023810      323095260              0              0
  00:00:c0:01:01:07       7023810      323095260              0              0
  00:00:c0:01:01:08       7023809      323095214              0              0
  00:00:c0:01:01:09       7023809      323095214              0              0
  00:00:c0:01:01:0a       7023809      323095214              0              0
  00:00:c0:01:01:0b       7023809      323095214              0              0
  00:00:c8:01:01:02      31016568     1426762128       38040381     1749857526
```

Copyright © 2016, Juniper Networks, Inc.

```
  00:00:c8:01:01:03      31016568     1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499     1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573     1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567     1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567     1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580      433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496     1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498     1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
```


Copyright © 2016, Juniper Networks, Inc.

```
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues      : 8 supported
  Last flapped    : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  SONET alarms    : None
  SONET defects   : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

Copyright © 2016, Juniper Networks, Inc.

```
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
  Queue Buffer Usage:
    Reserved buffer      :            118750000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :              6250000 bytes
  Queue-depth bytes      :
    Current              :                    0
..
..
```

Sample Output

show interfaces redundancy

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present               sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

show interfaces redundancy (Aggregated Ethernet)

```
user@host> show interfaces redundancy
Interface  State           Last change  Primary      Secondary    Current status
rlsq0      On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up

ae0
ae1
```


Copyright © 2016, Juniper Networks, Inc.

```
ae2
ae3
ae4
```

show interfaces redundancy detail

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

Sample Output

show interfaces routing brief

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
```

Copyright © 2016, Juniper Networks, Inc.

```
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

Sample Output

show interfaces routing-instance all

```
user@host> show interfaces terse routing-instance all
Interface   Admin   Link   Proto    Local             Remote Instance
at-0/0/1    up      up     inet     10.0.0.1/24
ge-0/0/0.0  up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0  up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0  up      up     inet     10.0.0.1/32
```

Sample Output

show interfaces snmp-index

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

Sample Output

show interfaces source-class all

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                                Packets                Bytes
      Source class                  (packet-per-second)    (bits-per-second)
                             gold              1928095            161959980
                                  (                889) (            597762)
                           bronze                    0                    0
```

 Copyright © 2016, Juniper Networks, Inc.

```
                                    (                  0) (                  0)
                             silver                    0                     0
                                    (                  0) (                  0)
  Logical interface so-0/1/3.0
                                                    Packets                 Bytes
      Source class                   (packet-per-second)    (bits-per-second)
                               gold                    0                     0
                                    (                  0) (                  0)
                             bronze                    0                     0
                                    (                  0) (                  0)
                             silver               116113               9753492
                                    (                939) (             631616)
```

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                                    Packets                 Bytes
      Destination class              (packet-per-second)    (bits-per-second)
                              silver1                    0                     0
                                    (                  0) (                  0)
                              silver2                    0                     0
                                    (                  0) (                  0)
                              silver3                    0                     0
                                    (                  0) (                  0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

Sample Output

show interfaces switch-port

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                          Receive         Transmit
          Total bytes                       28437086         21792250
```

Copyright © 2016, Juniper Networks, Inc.

```
          Total packets                      409145          88008
          Unicast packets                      9987          83817
          Multicast packets                  145002              0
          Broadcast packets                  254156           4191
          Multiple collisions                    23             10
          FIFO/CRC/Align errors                   0              0
          MAC pause frames                        0              0
          Oversized frames                        0
          Runt frames                             0
          Jabber frames                           0
          Fragment frames                         0
          Discarded frames                        0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

Sample Output

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                   COUNT          THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr         2008544300       0                NA              NA
  FEC-UncorrectedWords     0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
```


Copyright © 2016, Juniper Networks, Inc.

```
    PM                  MIN       MAX       AVG       THRESHOLD       TCA-ENABLED
     TCA-RAISED
  BER                   3.6e-5    5.8e-5    3.6e-5     10.0e-3        No
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                           CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                                 (MIN)
(MAX)    (MIN) (MAX)       (MIN) (MAX)
  Lane chromatic dispersion         0         0         0         0         0
 0        NA    NA         NA     NA
  Lane differential group delay     0         0         0         0         0
 0        NA    NA         NA     NA
    q Value                            120       120       120       120       0
   0        NA    NA         NA     NA
  SNR                               28        28        29        28        0
 0        NA    NA         NA     NA
  Tx output power(0.01dBm)          -5000     -5000     -5000     -5000     -300
 -100     No    No         No     No
  Rx input power(0.01dBm)           -3642     -3665     -3626     -3637     -1800
 -500     No    No         No     No
  Module temperature(Celsius)       46        46        46        46        -5
 75       No    No         No     No
  Tx laser bias current(0.1mA)      0         0         0         0         0
 0        NA    NA         NA     NA
  Rx laser bias current(0.1mA)      1270      1270      1270      1270      0
 0        NA    NA         NA     NA
  Carrier frequency offset(MHz)     -186      -186      -186      -186      -5000
 5000     No    No         No     No
```

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
```

Copyright © 2016, Juniper Networks, Inc.

```
Interfaces bound: 1
Interfaces:
  ge-0/0/2.0
```


Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management adjacencies

Syntax show oam ethernet connectivity-fault-management adjacencies <interface-name>

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display connectivity-fault-management adjacencies.

Options **interface-name**—Dispalys the name of the interface.

Required Privilege Level view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management adjacencies on page 3143

Output Fields Table 315 lists the output fields for the **show oam ethernet connectivity-fault-management adjacencies** command. Output fields are listed in the approximate order in which they appear

Table 315: show oam ethernet connectivity-fault-management adjacencies Output Fields

Field Name	Field Description
Mep-id	Maintenance association end point (MEP) identifier.
Interface	Interface identifier.
State	Indicates if the connectivity check protocol is up.
Timer to Expire	Indicates the expiration time.

Sample Output

show oam ethernet connectivity-fault- management adjacencies

```
user@host> show oam ethernet connectivity-fault-management adjacencies
Mep-id     Interface        State        Timer to Expire
         101     fe-0/0/4.0         ok                 29
```

show oam ethernet connectivity-fault-management forwarding-state

Syntax show oam ethernet connectivity-fault-management forwarding-state <interface>

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display the Ethernet OAM forwarding state for received packets.

Options **<interface>**—Displays the Ethernet OAM state for an interface.

Required Privilege Level view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management forwarding-state on page 3144

Output Fields Table 316 lists the output fields for the **show oam ethernet connectivity-fault-management forwarding-state** command. Output fields are listed in the approximate order in which they appear.

Table 316: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description
Interface name	Interface identifier.
Level	Maintenance domain level.
Direction	MEP direction configured.
Filter action	Filter action for messages at the level.
Nexthop type	Next-hop type.
Nexthop index	Next-hop index number.

Sample Output

show oam ethernet connectivity-fault- management forwarding-state

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-0/0/1.0 vlan:100
Instance name: INSTANCE_0 bd_vlan_100
  Maintenance domain forwarding state:

  Level    Direction      Filter action      Nexthop      Nexthop
                                             type         index
  0                       Drop               Discard
```


Copyright © 2016, Juniper Networks, Inc.

```
1                   Drop            Discard
2                   Drop            Discard
3                   Drop            Discard
4                   Drop            Discard
5                   Drop            Discard
6                   Drop            Discard
7         down      Receive         Receive
```

Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management interfaces

Syntax show oam ethernet connectivity-fault-management interfaces <interface name>

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display Ethernet OAM information for the specified interface.

Options **<interface name>**—Displays connectivity fault management information for the specified interface.

Required Privilege Level view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management interfaces on page 3146

Output Fields Table 317 lists the output fields for the **show oam ethernet connectivity-fault-management interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 317: show oam ethernet connectivity-fault-management interfaces Output Fields

Field Name	Field Description
Interfaces	Interface identifier.
Link	The local link status is Up, down, or oam-down.
Status	The status is active or inactive.
Level	Maintenance domain level configured.
MEP Identifier	Maintenance association end point (MEP) identifier.
Neighbors	Number of MEP neighbors.

Sample Output

show oam ethernet connectivity-fault- management interfaces

```
user@host> show oam ethernet connectivity-fault-management interfaces
Interfaces      Link      Status       Level       MEP         Neighbours
                                                          Identifier


        ge-0/0/1.0      Up        Active       7           1000            0
```

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management mep-database

Syntax show oam ethernet connectivity-fault-management mep-database

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Displays Ethernet OAM maintenance endpoint database information.

Options **<local-mep>**—Identifier for local maintenance endpoint (1 through 8191).

maintenance-association —Name of the maintenance association.

maintenance-domain —Name of the maintenance domain.

remote-mep —Identifier for remote maintenance endpoint (1 through 8191).

Required Privilege Level View

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management mep-database on page 3150

Output Fields Table 318 lists the output fields for the **show oam ethernet connectivity-fault-management mep-database** command. Output fields are listed in the approximate order in which they appear.

Table 318: show oam ethernet connectivity-fault-management mep-database Output Fields

Field Name	Field Description
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity check status.
Interval	Continuity check message interval.
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.

Copyright © 2016, Juniper Networks, Inc.

Table 318: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
MAC address	MAC address configured for the MEP.
Auto-discovery	Indicates whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and Link Trace messages (LTMs) transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.
Remote MEP not receiving CCM	Indicates that the remote MEP is not receiving CCMs.
Erroneous CCM received	Indicates that erroneous CCMs have been received.
Cross-connect CCM received	Indicates that cross-connect CCMs have been received.
RDI sent by some MEP	Indicates that the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Link Trace messages (LTMs) transmitted.
LTMs received	LTMs received.
LTRs sent	Link Trace Replies (LTRs) transmitted.
LTRs received	LTRs received.
Sequence number of next LTM request	Sequence number of the next Link Trace LTM request to be transmitted.

Copyright © 2016, Juniper Networks, Inc.

Table 318: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
1DMs sent	If the MEP is an initiator for a one-way ETH-DM session, then this is the number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays 0.
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session, then this is the number of valid 1DM frames received. For all other cases, this field displays 0.
Invalid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session, then this is the number of invalid 1DM frames received. For all other cases, this field displays 0.
DMMs sent	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays 0.
DMRs sent	If the MEP is a responder for a ETH-DM session, then this is the number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays 0.
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of valid DMRs received. For all other cases, this field displays 0.
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of invalid DMRs received. For all other cases, this field displays 0.

Sample Output

show oam ethernet connectivity-fault- management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain Customer1
Maintenance domain name: Customer1, Format: string, Level: 7
  Maintenance association name: Track_vlan_100, Format: string
  Continuity-check status: enabled, Interval: 1s
  MEP identifier: 1000, Direction: down, MAC address: 80:71:1f:ad:53:81
  Auto-discovery: disabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                 : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                         : no
```


Copyright © 2016, Juniper Networks, Inc.

```
Statistics:
  CCMs sent                                  : 170114
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 1
  LTRs sent                                  : 1
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 0
  Valid 1DMs received                        : 0
  Invalid 1DMs received                      : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                        : 0
  Invalid DMRs received                      : 0
```

Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management mep-statistics

Syntax

```
show oam ethernet connectivity-fault-management mep-statistics
count
local-mep
maintenance-association
maintenance-domain
remote-mep
```

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display Ethernet OAM maintenance endpoint statistics.

NOTE: The delay measurement (DM) statistics are not valid for SRX Series devices, which supports only the IEEE 802.1ag standard.

Options

count —Number of statistics per maintenance endpoint (1 through 100).

local-mep —Identifier for local maintenance endpoint (1 through 8191).

maintenance-association—Name of maintenance association.

maintenance-domain—Name of maintenance domain.

remote-mep —Identifier for remote maintenance endpoint (1 through 8191).

Required Privilege Level view

Related Documentation

- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management mep-statistics on page 3154

Output Fields Table 319 lists the output fields for the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 319: show oam ethernet connectivity-fault-management mep-statistics Output Fields

Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) identifier.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.

Copyright © 2016, Juniper Networks, Inc.

Table 319: show oam ethernet connectivity-fault-management mep-statistics Output Fields *(continued)*

Field Name	Field Description
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Link Trace messages (LTMs) transmitted.
LTMs received	Link Trace messages received.
LTRs sent	Link Trace responses (LTRs) transmitted.
LTRs received	Link Trace responses received.
Sequence number of next LTM request	Sequence number of the next Link Tracemessage request to be transmitted.
1DMs sent	If the MEP is an initiator in a one-way ETH-DM session, then this is the number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays 0.
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session, then this is the number of valid 1DM frames received. For all other cases, this field displays 0.
Invalid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session, then this is the number of invalid 1DM frames received. For all other cases, this field displays 0.
DMMs sent	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays 0.
DMRs sent	If the MEP is a responder for a ETH-DM session, then this is the number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays 0.
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of valid DMRs received. For all other cases, this field displays 0.

Copyright © 2016, Juniper Networks, Inc.

Table 319: show oam ethernet connectivity-fault-management mep-statistics Output Fields *(continued)*

Field Name	Field Description
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session, then this is the number of invalid DMRs received. For all other cases, this field displays 0.

Sample Output

show oam ethernet connectivity-fault- management mep-statistics

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain private maintenance-association private-ma remote-mep 100
MEP identifier: 101, MAC address: 80:71:1f:ad:53:81
  CCMs sent                                       : 83
  CCMs received out of sequence                   : 0
  LBMs sent                                       : 0
  Valid in-order LBRs received                    : 0
  Valid out-of-order LBRs received                : 0
  LBRs received with corrupted data               : 0
  LBRs sent                                       : 0
  LTMs sent                                       : 0
  LTMs received                                   : 0
  LTRs sent                                       : 0
  LTRs received                                   : 0
  Sequence number of next LTM request             : 0
  1DMs sent                                       : 0
  Valid 1DMs received                             : 0
  Invalid 1DMs received                           : 0
  DMMs sent                                       : 0
  DMRs sent                                       : 0
  Valid DMRs received                             : 0
  Invalid DMRs received                           : 0
```


Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management mip

Syntax show oam ethernet connectivity-fault-management mip
interface-name
vlan

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display MIP information.

Options **interface-name** —Displays information of the specified logical interface.

vlan—Displays information about the specified VLAN (1 through 4094).

Required Privilege Level view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management mip on page 3155

Output Fields Table 320 lists the output fields for the **show oam ethernet connectivity-fault-management mip** command. Output fields are listed in the approximate order in which they appear.

Table 320: show oam ethernet connectivity-fault-management mip Output Fields

Field Name	Field Description
Default Maintenance-domain	The default maintenance domain name.
Interface	Interface identifier.
Level	Maintenance domain level configured.

Sample Output

show oam ethernet connectivity-fault- management mip

```
user@host> show oam ethernet connectivity-fault-management mip vlan 100
default maintenance-domain mhf    : default

  Interface       Level
  ge-0/0/1.0          5
  fe-0/0/4.0          5
```

show oam ethernet connectivity-fault-management path-database

Syntax

show oam ethernet connectivity-fault-management path-database
<host>
maintenance-association
maintenance-domain

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Description

Display the Link Trace path database for a remote host.

Options

<host>—MAC address of the remote host in xx:xx:xx:xx:xx:xx format.

maintenance-association —Name of the maintenance association.

maintenance-domain —Name of the maintenance domain.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output

show oam ethernet connectivity-fault- management path-database on page 3157

Output Fields

Table 321 lists the output fields for the **show oam ethernet connectivity-fault-management path-database** command. Output fields are listed in the approximate order in which they appear.

Table 321: show oam ethernet connectivity-fault-management path-database Output Fields

Field Name	Field Description
Interface	Interface Identifier.
Maintenance Domain	Maintenance domain name.
Maintenance Association	Maintenance association name.
Level	Maintenance domain level configured for the maintenance domain.
Hop	Sequential hop count of the Link Trace path.
TTL	Number of hops remaining in the Link Trace message (LTM). The time to live (TTL) is decremented at each hop.
Source MAC Address	MAC address of the 802.1ag maintenance association intermediate point (MIP) that is forwarding the LTM.
Next-hop MAC Address	MAC address of the 802.1ag node that is the next hop in the LTM path.

Copyright © 2016, Juniper Networks, Inc.

Table 321: show oam ethernet connectivity-fault-management path-database Output Fields *(continued)*

Field Name	Field Description
Transaction Identifier	Identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming Link Trace Reply (LTR) with a previously sent LTM.

Sample Output

show oam ethernet connectivity-fault- management path-database

```
user@host> show oam ethernet connectivity-fault-management path-database
Interface : fe-0/0/4
      Maintenance Domain: private, Level: 5
      Maintenance Association: private-ma, Local Mep: 100

    Hop   TTL    Source MAC address        Next-hop MAC address
    Transaction Identifier:0
    1     63     80:71:1f:ad:50:04         80:71:1f:ad:50:01
    2     62     80:71:1f:ad:53:81         00:00:00:00:00:00
    Transaction Identifier:1
    1     63     80:71:1f:ad:50:04         80:71:1f:ad:50:01
    2     62     80:71:1f:ad:53:81         00:00:00:00:00:00
    Transaction Identifier:2
    1     63     80:71:1f:ad:50:04         80:71:1f:ad:50:01
    2     62     80:71:1f:ad:53:81         00:00:00:00:00:00
    Transaction Identifier:3
    1     63     80:71:1f:ad:50:04         80:71:1f:ad:50:01
    2     62     80:71:1f:ad:53:81         00:00:00:00:00:00
```

Copyright © 2016, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management routes

Syntax show oam ethernet connectivity-fault-management routes

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display connectivity-fault-management bridge routes.

Required Privilege Level view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097

List of Sample Output show oam ethernet connectivity-fault- management routes on page 3158

Output Fields Table 322 lists the output fields for the **show oam ethernet connectivity-fault-management routes** command. Output fields are listed in the approximate order in which they appear.

Table 322: show oam ethernet connectivity-fault-management routes Output Fields

Field Name	Field Description
VLAN	The configured VLAN interface.
MAC	MAC address configured for the route.
Next-hop index	Software index of the next hop that is used to route the traffic for a given prefix.
Action	The next-hop action.

Sample Output

show oam ethernet connectivity-fault- management routes

```
user@host> show oam ethernet connectivity-fault-management routes
VLAN              MAC                Next-hop index  Action
vlan1             00:00:00:00:00:00          563             Flood
vlan1             01:80:c2:00:00:30                          Discard
vlan1             01:80:c2:00:00:31                          Discard
vlan1             01:80:c2:00:00:32                          Discard
vlan1             01:80:c2:00:00:33                          Discard
vlan1             01:80:c2:00:00:34                          Discard
vlan1             01:80:c2:00:00:35                          Receive
vlan1             01:80:c2:00:00:36          563             Flood
vlan1             01:80:c2:00:00:37          563             Flood
vlan1             01:80:c2:00:00:38                          Discard
vlan1             01:80:c2:00:00:39                          Discard
vlan1             01:80:c2:00:00:3a                          Discard
vlan1             01:80:c2:00:00:3b                          Discard
vlan1             01:80:c2:00:00:3c                          Discard
vlan1             01:80:c2:00:00:3d                          Receive
vlan1             01:80:c2:00:00:3e           563            Flood
```


Copyright © 2016, Juniper Networks, Inc.

```
vlan1            01:80:c2:00:00:3f          563           Flood
vlan2            00:00:00:00:00:00          563           Flood
```

show oam ethernet link-fault-management

Syntax
show oam ethernet link-fault-management
<brief | detail>
<interface-name>

Release Information
Statement for branch SRX Series devices introduced in Junos OS Release 9.5.

Description
Display Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options
brief | detail—(Optional) Display the specified level of output.

interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level
view

Related Documentation
- clear oam ethernet connectivity-fault-management path-database on page 3096
- clear oam ethernet connectivity-fault-management statistics on page 3097
- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 2603
- Example: Configuring Ethernet OAM Link Fault Management on page 2605

List of Sample Output
show oam ethernet link-fault-management brief on page 3164
show oam ethernet link-fault-management detail on page 3164

Output Fields
Table 290 lists the output fields for the **show oam ethernet link-fault-management** command. Output fields are listed in the approximate order in which they appear.

Table 323: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Status of the established link. • **Fail**—A link fault condition exists. • **Running**—A link fault condition does not exist.	All levels
Discovery state	State of the discovery mechanism: • **Passive Wait** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok**	All levels
Peer address	Address of the OAM peer.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 323: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. • **Remote-Stable**—Indicates remote OAM client acknowledgment of, and satisfaction with, local OAM state information. **False** indicates that remote DTE has either not seen or is unsatisfied with local state information. **True** indicates that remote DTE has seen and is satisfied with local state information. • **Local-Stable**—Indicates local OAM client acknowledgment of, and satisfaction with, remote OAM state information. **False** indicates that local DTE either has not seen or is unsatisfied with remote state information. **True** indicates that local DTE has seen and is satisfied with remote state information. • **Remote-State-Valid**—Indicates the OAM client has received remote state information found within local information TLVs (type, length, values) of received Information OAM PDUs. **False** indicates that the OAM client has not seen remote state information. **True** indicates that the OAM client has seen remote state information.	All levels
Remote loopback status	An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	Remote entity information. • **Remote MUX action**—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • **Remote parser action**—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to the higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • **Discovery mode**—Indicates whether discovery mode is active or inactive. • **Unidirectional mode**—Indicates the ability to operate a link in unidirectional mode for diagnostic purposes. • **Remote loopback mode**—Indicates whether remote loopback is supported or not supported. • **Link events**—Indicates whether interpreting link events is supported or not supported on the remote peer. • **Variable requests**—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.	All levels
OAM Receive Statistics		
Information	Number of information PDUs received.	**detail**
Event	Number of loopback control PDUs received.	**detail**
Variable request	Number of variable request PDUs received.	**detail**
Variable response	Number of variable response PDUs received.	**detail**
Loopback control	Number of loopback control PDUs received.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 323: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Organization specific	Number of vendor organization specific PDUs received.	**detail**
OAM Transmit Statistics		
Information	Number of information PDUs transmitted.	**detail**
Event	Number of event notification PDUs transmitted.	**detail**
Variable request	Number of variable request PDUs transmitted.	**detail**
Variable response	Number of variable response PDUs transmitted.	**detail**
Loopback control	Number of loopback control PDUs transmitted.	**detail**
Organization specific	Number of vendor organization specific PDUs transmitted.	**detail**
OAM Received Symbol Error Event information		
Events	Number of symbol error event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	**detail**
Threshold	Number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	Number of symbol errors in the period reported in the received event PDU.	**detail**
Total errors	Number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	**detail**
OAM Received Frame Error Event Information		
Events	Number of errored frame event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Duration of the window in terms of the number of 100 ms period intervals.	**detail**
Threshold	Number of detected errored frames required for the event to be generated.	**detail**
Errors in period	Number of detected errored frames in the period.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 323: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Total errors	Number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	**detail**
OAM Received Frame Period Error Event Information		
Events	Number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	Duration of the frame seconds window.	**detail**
Threshold	Number of frame seconds errors in the period.	**detail**
Errors in period	Number of frame seconds errors in the period.	**detail**
Total errors	Number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	**detail**
OAM Transmitted Symbol Error Event Information		
Events	Number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the transmitted PDU.	**detail**
Threshold	Number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	Number of symbol errors in the period reported in the transmitted event PDU.	**detail**
Total errors	Number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
OAM Transmitted Frame Error Event Information		
Events	Number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	Duration of the window in terms of the number of 100-ms period intervals.	**detail**
Threshold	Number of detected errored frames required for the event to be generated.	**detail**
Errors in period	Number of detected errored frames in the period.	**detail**
Total errors	Number of errored frames that have been detected after the OAM sublayer was reset.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
  Status: Running, Discovery state: Send Any
  Peer address: 00:90:69:72:2c:83
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  Remote loopback status: Disabled on local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-0/0/1
    Status: Running, Discovery state: Send Any
    Peer address: 00:90:69:0a:07:14
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 186365, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM transmit statistics:
      Information: 186347, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM received symbol error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: supported, Link events: supported
      Variable requests: unsupported
```


Copyright © 2016, Juniper Networks, Inc.

show security flow gate family

Syntax show security flow gate family (inet | inet6)

Release Information Command introduced in Junos OS Release 10.4.

Description Display filtered summary of information about existing gates, types of gates, and the maximum allowed number of gates.

Options
- inet—Displays IPv4 information.
- inet6—Displays IPv6 gate information.

Required Privilege Level view

Related Documentation
- show security flow gate on page 449

Output Fields Table 1 lists the output fields for the **show security flow gate family** command. Output fields are listed in the approximate order in which they appear.

Table 324: show security flow gate family Output Fields

Field Name	Field Description
Valid gates	Number of valid gates.
Pending gates	Number of pending gates.
Invalidated gates	Number of invalid gates.
Gates in other states	Number of gates in other states.
Total gates	Total number of gates.

Sample Output

```
user@host> show security flow gate family inet6
Hole: 2001:13::8-0-0->2001:12::8-33135-33135

  Translated: ::/0->::/0

  Protocol: tcp

  Application: FTP ALG/79

  Age: 24 seconds

  Flags: 0x8080

  Zone: zserver
```

Copyright © 2016, Juniper Networks, Inc.

```
  Reference count: 1

  Resource: 1-2-2


Valid gates: 1

Pending gates: 0

Invalidated gates: 0

Gates in other states: 0

Total gates: 1


user@host> show security flow gate family inet6 destination-prefix 2001:12::8 or source-prefix
Hole: 2001:13::8-0-0->2001:12::8-33135-33135

  Translated: ::/0->::/0

  Protocol: tcp

  Application: FTP ALG/79

  Age: 26 seconds

  Flags: 0x8080

  Zone: zserver

  Reference count: 1

  Resource: 1-2-2


Valid gates: 1

Pending gates: 0

Invalidated gates: 0

Gates in other states: 0

Total gates: 1
```


Copyright © 2016, Juniper Networks, Inc.

show security flow ip-action

Syntax **show security flow ip-action** [<filter>] [summary family (inet | inet6)]

Release Information Command introduced in Junos OS Release 10.1. Logical systems option added in Junos OS Release 11.2 . Summary option introduced in Junos OS Release 12.1.

Description Display the current IP-action settings, based on filtered options, for IP sessions running on the device.

Options

- *filter*—Filter the display based on the specified criteria.

  The following filters display those sessions that match the criteria specified by the filter. Refer to the sample output for filtered output examples.

  all | [*filter*]—All active sessions on the device.

  destination-port ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535.

  destination-prefix ***destination-prefix***—Destination IP prefix or address.

  family (inet | inet6) [*filter*]—IPv4 traffic or IPv6-NATPT traffic and filtered options.

  logical-system ***logical-system-name*** **| all** [*filter*]—Specified logical system or all logical systems.

  protocol ***protocol-name*** **|** ***protocol-number*** [*filter*]—Protocol name or number and filtered options.

  - **ah** or **51**
  - **egp** or **8**
  - **esp** or **50**
  - **gre** or **47**
  - **icmp** or **1**
  - **icmp6** or **58**
  - **igmp** or **2**
  - **ipip** or **4**
  - **ospf** or **89**
  - **pim** or **103**
  - **rsvp** or **46**
  - **sctp** or **132**
  - **tcp** or **6**
  - **udp** or **17**

  root-logical-system [*filter*]—Default logical system information and filtered options.

Copyright © 2016, Juniper Networks, Inc.

source-port *source-port*—Source port number of the traffic. Range is 1 through 65,535.

source-prefix *source-prefix*—Source IP prefix or address of the traffic.

- **summary** —Summary information about IP-action entries.

  family—Display summary of IP-action entries by family. This option is used to filter the output.

  - **inet**—Display summary of IPv4 entries.
  - **inet6**—Display summary of IPv6 entries.

Required Privilege Level view

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- clear security flow ip-action on page 1900
- clear security flow session destination-port on page 1905
- clear security flow ip-action on page 1900

List of Sample Output

show security flow ip-action on page 3169
show security flow ip-action destination-port on page 3170
show security flow ip-action destination-prefix on page 3171
show security flow ip-action family inet protocol on page 3171
show security flow ip-action family inet logical-system all on page 3172
show security flow ip-action source-prefix on page 3173
show security flow ip-action summary on page 3174
show security flow ip-action summary family inet on page 3174
show security flow ip-action summary family inet6 on page 3174

Output Fields Table 208 lists the output fields for the **show security flow ip-action** command. Output fields are listed in the approximate order in which they appear.

Table 325: show security flow ip-action Output Fields

Field Name	Field Description
Src-Addr	Source address of outbound IP traffic.
Src-Port	Source port number of outbound IP traffic.
Dst-Addr	Destination address of inbound IP traffic.
Dst-Port/Proto	Destination port number and protocol type of inbound IP traffic.
Timeout (sec)	Configured timeouts and time remaining for an IP session.
Zone	Security zone associated with an IP session.
Action	Configured action type, for example, block, close, and notify.

Copyright © 2016, Juniper Networks, Inc.

Table 325: show security flow ip-action Output Fields *(continued)*

Field Name	Field Description
State	The active mode and passive mode describe the states of the **ip-action** entry.
IPv4 action count	The total number of IPv4 entries.
IPv6 action count	The total number of IPv6 entries.

Sample Output

show security flow ip-action

```
user@host> show security flow ip-action
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
 203.0.113.1       *       203.0.113.4         21/tcp         293/300      *
      close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *       203.0.113.4          21/tcp         293/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *      203.0.113.4          21/tcp         293/300       *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *    203.0.113.4          21/tcp         292/300      *
   close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
 203.0.113.1         *    203.0.113.4          21/tcp         292/300      *
   close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
```

Copyright © 2016, Juniper Networks, Inc.

```
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action destination-port

```
user@host> show security flow ip-action destination-port 21

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
 close         Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4        21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1        *        203.0.113.4         21/tcp         274/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         273/300      *
  close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```


Copyright © 2016, Juniper Networks, Inc.

show security flow ip-action destination-prefix

```
user@host>  show security flow ip-action destination-prefix 203.0.113.4/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
192.0.2.3           *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
192.0.2.3           *        203.0.113.4         21/tcp         245/300      *
 close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
  close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         21/tcp         245/300      *
 close        Active
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet protocol

```
user@host>  show security flow ip-action family inet protocoludp
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
        close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC2
```

Copyright © 2016, Juniper Networks, Inc.

```
Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
     close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
     close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1        *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        203.0.113.4         69/udp         287/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action family inet logical-system all

```
user@host> show security flow ip-action family inet logical-system all

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
        close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
 close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *       203.0.113.4         69/udp         267/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
  close        Active            root-logical-system
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
```


Copyright © 2016, Juniper Networks, Inc.

```
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         267/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         266/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State             Logical-System
203.0.113.1         *        203.0.113.4         69/udp         266/300      *
  close        Passive           root-logical-system
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action source-prefix

```
user@host> show security flow ip-action source-prefix 192.0.2.3/8

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
  close        Passive
 IPv4 action count: 1 on FPC0.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
   close        Passive
 IPv4 action count: 1 on FPC0.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4        69/udp         244/300      *
    close        Passive
 IPv4 action count: 1 on FPC0.PIC3

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
     close        Active
 IPv4 action count: 1 on FPC1.PIC0

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
    close        Passive
 IPv4 action count: 1 on FPC1.PIC1

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
     close        Passiveo
 IPv4 action count: 1 on FPC1.PIC2

Src-Addr        Src-Port Dst-Addr        Dst-Port/Proto Timeout(sec) Zone
Action       State
203.0.113.1         *        192.0.2.4         69/udp         244/300      *
```

Copyright © 2016, Juniper Networks, Inc.

```
    close        Passive
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary

```
user@host> show security flow ip-action summary

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
 IPv6 action count: 0 on FPC0.PIC1
 IPv6 action count: 0 on FPC0.PIC2
 IPv6 action count: 0 on FPC0.PIC3
 IPv6 action count: 0 on FPC1.PIC0
 IPv6 action count: 0 on FPC1.PIC1
 IPv6 action count: 0 on FPC1.PIC2
 IPv6 action count: 0 on FPC1.PIC3
 IPv6 action count: Active mode 0 on all PICs
```

show security flow ip-action summary family inet

```
user@host> show security flow ip-action summary inet

 IPv4 action count: 1 on FPC0.PIC1
 IPv4 action count: 1 on FPC0.PIC2
 IPv4 action count: 1 on FPC0.PIC3
 IPv4 action count: 1 on FPC1.PIC0
 IPv4 action count: 1 on FPC1.PIC1
 IPv4 action count: 1 on FPC1.PIC2
 IPv4 action count: 1 on FPC1.PIC3
 IPv4 action count: Active mode 1 on all PICs
```

show security flow ip-action summary family inet6

```
user@host> show security flow ip-action summary family inet6

 IPv6 action count: 1 on FPC0.PIC1
 IPv6 action count: 1 on FPC0.PIC2
 IPv6 action count: 1 on FPC0.PIC3
 IPv6 action count: 1 on FPC1.PIC0
 IPv6 action count: 1 on FPC1.PIC1
 IPv6 action count: 1 on FPC1.PIC2
 IPv6 action count: 1 on FPC1.PIC3
 IPv6 action count: Active mode 1 on all PICs
```


Copyright © 2016, Juniper Networks, Inc.

show security flow session family

Syntax show security flow session family (inet | inet6)
[brief | extensive | summary]

Release Information Command introduced in Junos OS Release 10.2.

Description Display filtered summary of information about existing sessions, including types of sessions, active and failed sessions, and the maximum allowed number of sessions.

Options
- **inet**—Display details summary of IPv4 sessions.
- **inet6**—Display details summary of IPv6 sessions.
- brief | extensive | summary–Display the specified level of output.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session family on page 1907

List of Sample Output
show security flow session family inet on page 3176
show security flow session family inet brief on page 3177
show security flow session family inet extensive on page 3177
show security flow session family inet summary on page 3179

Output Fields Table 219 lists the output fields for the **show security flow session family** command. Output fields are listed in the approximate order in which they appear.

Table 326: show security flow session family Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.

Copyright © 2016, Juniper Networks, Inc.

Table 326: show security flow session family Output Fields *(continued)*

Field Name	Field Description
Flag	Internal flag depicting the state of the session, used for debugging purposes.
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Application	Name of the application.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Number of maximum sessions.

Sample Output

show security flow session family inet

```
root> show security flow session family inet
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000107, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 200.0.0.10/3 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000202
  Out: 60.0.0.1/19001 --> 200.0.0.10/3;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000202
Total sessions: 1

Flow Sessions on FPC10 PIC3:
```


Copyright © 2016, Juniper Networks, Inc.

```
Session ID: 430000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000110
  Out: 60.0.0.1/19001 --> 200.0.0.10/2;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000110

Session ID: 430000117, Policy name: default-policy-00/2, Timeout: 4, Valid
  In: 200.0.0.10/4 --> 60.0.0.1/19001;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000111
  Out: 60.0.0.1/19001 --> 200.0.0.10/4;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000111
Total sessions: 2
```

show security flow session family inet brief

```
root> show security flow session family inet brief
Flow Sessions on FPC10 PIC1:
Total sessions: 0

Flow Sessions on FPC10 PIC2:

Session ID: 420000115, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/1 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000206
  Out: 60.0.0.1/19769 --> 200.0.0.10/1;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000206

Session ID: 420000117, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/2 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 420000207
  Out: 60.0.0.1/19769 --> 200.0.0.10/2;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 420000207
Total sessions: 2

Flow Sessions on FPC10 PIC3:

Session ID: 430000119, Policy name: default-policy-00/2, Timeout: 2, Valid
  In: 200.0.0.10/3 --> 60.0.0.1/19769;icmp, If: ge-7/1/0.0, Pkts: 1, Bytes: 84,
CP Session ID: 430000112
  Out: 60.0.0.1/19769 --> 200.0.0.10/3;icmp, If: .local..0, Pkts: 1, Bytes: 84,
CP Session ID: 430000112
Total sessions: 1
```

show security flow session family inet extensive

```
root> show security flow session family inet extensive
Flow Sessions on FPC10 PIC1:

Session ID: 410000111, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76455, Duration: 0
   In: 200.0.0.10/4 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
```

```
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
   Out: 60.0.0.1/20537 --> 200.0.0.10/4;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 410000242
Total sessions: 1

Flow Sessions on FPC10 PIC2:

Session ID: 420000123, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 2
Session State: Valid
Start time: 76454, Duration: 2
   In: 200.0.0.10/3 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
   Out: 60.0.0.1/20537 --> 200.0.0.10/3;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 420000210
Total sessions: 1

Flow Sessions on FPC10 PIC3:

Session ID: 430000131, Status: Normal
Flags: 0x80400040/0x0/0x2800023
Policy name: default-policy-00/2
Source NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 4, Current timeout: 4
Session State: Valid
Start time: 76421, Duration: 1
   In: 200.0.0.10/5 --> 60.0.0.1/20537;icmp,
    Interface: ge-7/1/0.0,
    Session token: 0x6, Flag: 0xc0000021
    Route: 0xa0010, Gateway: 200.0.0.10, Tunnel: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
   Out: 60.0.0.1/20537 --> 200.0.0.10/5;icmp,
    Interface: .local..0,
    Session token: 0x2, Flag: 0x40000030
    Route: 0xfffb0006, Gateway: 60.0.0.1, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 84
    CP Session ID: 430000118
Total sessions: 1
```

show security flow session family inet summary

```
root> show security flow session family inet summary
Flow Sessions on FPC10 PIC1:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC2:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4

Flow Sessions on FPC10 PIC3:

Valid sessions: 2
Pending sessions: 0
Invalidated sessions: 2
Sessions in other states: 0
Total sessions: 4
```

Copyright © 2016, Juniper Networks, Inc.

show security flow statistics

Syntax **show security flow statistics**

Release Information Command introduced in Junos OS Release 10.2.
Starting with Junos OS Release 15.1X49-D10, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based datapath packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).
The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Description Display flow-related system statistics.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow statistics on page 3180
show security flow statistics (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3) on page 3181

Output Fields Table 233 lists the output fields for the **show security flow statistics** command. Output fields are listed in the approximate order in which they appear.

Table 327: show security flow statistics Output Fields

Field Name	Field Description
Current sessions	Number of current sessions.
Packets forwarded	Number of packets forwarded.
Packets dropped	Number of Packets dropped.
Fragment packets	Number of fragment packets.

Sample Output

show security flow statistics

```
root> show security flow statistics
Flow Statistics of FPC4 PIC1:
    Current sessions: 63
    Packets forwarded: 3001
    Packets dropped: 1281
    Fragment packets: 0

  Flow Statistics of FPC5 PIC0:
    Current sessions: 22
    Packets forwarded: 859
```


Copyright © 2016, Juniper Networks, Inc.

```
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC5 PIC1:
    Current sessions: 22
    Packets forwarded: 858
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics Summary:
    System total valid sessions: 107
    Packets forwarded: 4718
    Packets dropped: 1281
    Fragment packets: 0
```

show security flow statistics (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
root> show security flow statistics
Flow Statistics of FPC0 PIC1:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC0 PIC2:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics of FPC0 PIC3:
    Current sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0

  Flow Statistics Summary:
    System total valid sessions: 0
    Packets forwarded: 0
    Packets dropped: 0
    Fragment packets: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security flow status

Syntax **show security flow status**

Release Information Command introduced in Junos OS Release 10.2; session distribution mode option added in Junos OS Release 12.1X44-D10; enhanced route scaling mode option added in Junos OS Release 12.1X45-D10. GTP-U distribution option added in Junos OS Release 15.1X49-D40.
Starting with Junos OS Release 15.1X49-D10, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based data path packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).
The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Description Display the flow processing modes and logging status.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow status on page 3183
show security flow status (IPsec Performance Acceleration) on page 3183
show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3) on page 3184

Output Fields Table 234 lists the output fields for the **show security flow status** command. Output fields are listed in the approximate order in which they appear.

Table 328: show security flow status Output Fields

Field Name	Field Description
Flow forwarding mode	Flow processing mode. • Inet forwarding mode • Inet6 forwarding mode • MPLS forwarding mode • ISO forwarding mode • Session distribution mode • Enhanced route scaling mode
Flow trace status	Flow logging status. • Flow tracing status • Flow tracing options

Copyright © 2016, Juniper Networks, Inc.

Table 328: show security flow status Output Fields *(continued)*

Field Name	Field Description
flow session distribution	SPU load distribution mode. • RR-based • Hash-based GTP-U distribution • Enabled
Flow packet ordering	packet-ordering mode. • Hardware • Software
Flow ipsec performance acceleration	IPsec VPN performance acceleration status.

Sample Output

show security flow status

```
root> show security flow status
Flow forwarding mode:
Inet forwarding mode: flow based
Inet6 forwarding mode: flow based
MPLS forwarding mode: drop
ISO forwarding mode: drop
Enhanced route scaling mode: Enabled (reboot needed to disable)
Flow trace status
Flow tracing status: on
Flow tracing options: all
Flow session distribution
Distribution mode: Hash-based
GTP-U distribution: Enabled
Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
```

show security flow status (IPsec Performance Acceleration)

```
root> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
Flow trace status
    Flow tracing status: off
Flow session distribution
    Distribution mode: RR-based
    GTP-U distribution: Enabled Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
Flow ipsec performance acceleration: on
```

Copyright © 2016, Juniper Networks, Inc.

show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
root> show security flow status
node0:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware

node1:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
```


Copyright © 2016, Juniper Networks, Inc.

show security forward-options secure-wire

Syntax show security forward-options secure-wire <*secure-wire-name*>

Release Information Command introduced in Junos OS Release 12.3X48-D10.

Description Display information about secure wire mappings.

Options
- none—Display information about all configured secure wire mappings.
- *secure-wire-name*—(Optional) Display information about the specified secure wire mapping.

Required Privilege Level view

Related Documentation
- Understanding Secure Wire on page 2955

List of Sample Output
show security forward-options secure-wire on page 3185
show security forward-options secure-wire pw1 on page 3186

Output Fields Table 329 lists the output fields for the **show security forward-options secure-wire** command. Output fields are listed in the approximate order in which they appear.

Table 329: show security forward-options secure-wire Output Fields

Field Name	Field Description
Secure wire	Name of the secure wire mapping.
Interface	One of the peer interfaces in the secure wire mapping.
Link	Operational status of the interface link.
Interface	The second peer interface in the secure wire mapping.
Link	Operational status of the interface link.

Sample Output

show security forward-options secure-wire

```
user@host> show security forward-options secure-wire
Secure wire                    Interface       Link   Interface       Link

pw1                            ge-11/1/0.0     up     ge-11/1/1.0     up
pw2                            ge-11/0/0.0     up     ge-11/0/1.0     up
pw3                            ge-11/1/2.0     down   ge-11/1/3.0     down
Total secure wires: 3
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security forward-options secure-wire pw1

```
user@host> show security forward-options secure-wire pw1
Secure wire                    Interface       Link   Interface       Link

pw1                            ge-11/1/0.0     up     ge-11/1/1.0     up
```


Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 3190
show security policies policy-name p1 detail on page 3191
show security policies (services-offload) on page 3192
show security policies detail on page 3192
show security policies detail (TCP Options) on page 3193
show security policies policy-name p1 (Negated Address) on page 3193
show security policies policy-name p1 detail (Negated Address) on page 3194
show security policies global on page 3194

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 330: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 330: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 330: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```


Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                 272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
      Initial direction:                9072                  272 bps
      Reply direction  :                9072                  272 bps
    Input  packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Output packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Session rate       :                 108                    3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

show security zones

Syntax
show security zones
<detail | terse>
< *zone-name* >

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones.

Options
- none—Display information about all zones.
- **detail | terse**—(Optional) Display the specified level of output.
- ***zone-name*** —(Optional) Display information about the specified zone.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones on page 3196
show security zones abc on page 3196
show security zones abc detail on page 3196
show security zones terse on page 3197

Output Fields
Table 32 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 331: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                    Type
my-internal             Security
my-external             Security
dmz                     Security
```

Copyright © 2016, Juniper Networks, Inc.

show vlans

Syntax
show vlans
<brief | detail | extensive>
<interface *interface-name*>
<logical-system (*logical-system* | *all*)>
<operational>

Release Information Command introduced in Junos OS Release 8.4.

Description Display VLAN information.

Options
none—Display information for all VLANs.

brief | detail | extensive—(Optional) Display the specified level of output.

interface *interface-name*— (Optional) Display information about a specific interface.

logical system—(Optional) Display name of the logical system or all.

operational—(Optional) Display information for the operational bridging instances.

Required Privilege Level view

List of Sample Output
show vlans on page 3198
show vlans brief on page 3198
show vlans detail on page 3199

Sample Output

show vlans

```
user@host> show vlans
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22
default-switch          vlan-333                 333
                                                              ge-0/0/3.0*
                                                              ge-0/0/4.0*
default-switch          default               1
default-switch          vlan100               100
                                                              ge-0/0/1.0*
```

show vlans brief

```
user@host> show vlans brief
Routing instance        VLAN name             Tag          Interfaces
default-switch          vlan-22                  22
default-switch          vlan-333                 333
                                                              ge-0/0/3.0*
                                                              ge-0/0/4.0*
default-switch          default               1
```


Copyright © 2016, Juniper Networks, Inc.

```
default-switch          vlan100               100
                                                                ge-0/0/1.0*
```

show vlans detail

```
user@host> show vlans detail
Routing instance: default-switch
  VLAN Name: vlan-22                            State: Active
Tag: 22
Internal index: 2, Generation Index: 1, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Number of interfaces: Tagged 0    , Untagged 0
Total MAC count: 0

Routing instance: default-switch
  VLAN Name: vlan-333                           State: Active
Tag: 333
Internal index: 3, Generation Index: 2, Origin: Static
MAC aging time: 300 seconds
VXLAN Enabled : No
Interfaces:
    ge-0/0/3.0*,tagged,trunk
    ge-0/0/4.0*,tagged,trunk
Number of interfaces: Tagged 2    , Untagged 0
Total MAC count: 0
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Logical Systems Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 44

Overview

- Introduction to Logical Systems on page 3205
- Understanding Master Logical Systems on page 3221
- Understanding User Logical Systems on page 3225

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 142

Introduction to Logical Systems

- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Understanding the Fundamentals and Constraints of Logical Systems on page 3208
- Understanding Licenses for Logical Systems on SRX Series Devices on page 3209
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210
- Understanding Flow in Logical Systems for SRX Series Devices on page 3211

Understanding Logical Systems for SRX Series Services Gateways

Logical systems for SRX Series devices enable you to partition a single device into secure contexts. Each logical system has its own discrete administrative domain, logical interfaces, routing instances, security firewall and other security features. By transforming an SRX Series device into a multitenant logical systems device, you can give various departments, organizations, customers, and partners–depending on your environment–private use of portions of its resources and a private view of the device. Using logical systems, you can share system and underlying physical machine resources among discrete user logical systems and the master logical system.

The top part of Figure 131 shows the three main configuration components of a logical system. The lower part of the figure shows a single device with a master logical system and discrete user logical systems.

Copyright © 2016, Juniper Networks, Inc.

Figure 131: Understanding Logical Systems

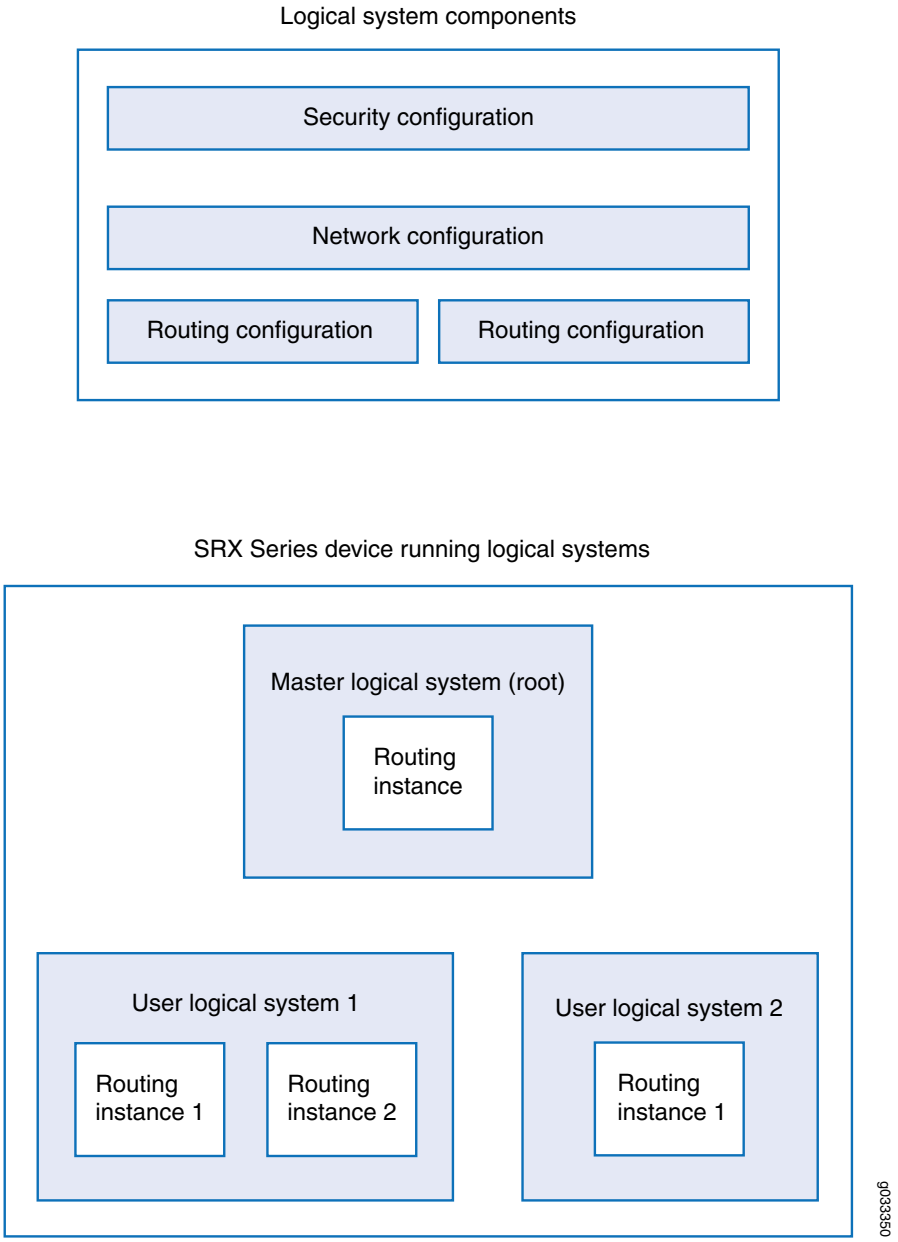

Logical systems on SRX Series devices offer many benefits, allowing you to:

- Curtail costs. Using logical systems, you can reduce the number of physical devices required for your company. Because you can consolidate services for various groups of users on a single device, you reduce both hardware costs and power expenditure.
- Create many logical systems on a single device and provision resources and services for them quickly. Because services are converged, it is easier for the master, or root, administrator to manage a single device configured for logical systems than it is to manage many discrete devices.

Copyright © 2016, Juniper Networks, Inc.

You can deploy an SRX Series device running logical systems in many environments, in particular, in the enterprise and in the data center.

- In the enterprise, you can create and provision logical systems for various departments and groups.

  You can configure logical systems to enable communication among groups sharing the device. When you create logical systems for various departments on the same device, users can communicate with one another without traffic leaving the device if you have configured an interconnect logical system to serve as an internal switch. For example, members of the product design group, the marketing department, and the accounting department sharing an SRX Series Services Gateway running logical systems can communicate with one another just as they could if separate devices were deployed for their departments. You can configure logical systems to interconnect through *logical tunnel (lt-0/0/0)* internal interfaces. The lt-0/0/0 interfaces on the interconnect logical system connect to an lt-0/0/0 interface that you configure for each logical system. The interconnect logical system switches traffic between logical systems. The SRX Series device running logical systems provides for high, fast interaction among all logical systems created on the device when an interconnect logical system is used.

  Logical systems on the same device can also communicate with one another directly through ports on the device, as if they were separate devices. Although this method allows for direct connections between logical systems, it consumes more resources—you must configure interfaces and an external switch—and therefore it is more costly.

- In the data center, as a service provider, you can deploy an SRX Series device running logical systems to offer your customers secure and private user logical systems and discrete use of the device's resources.

  For example, one corporation might require 10 user logical systems and another might require 20. Because logical systems are secure, private, and self-contained, data belonging to one logical system cannot be viewed by administrators or users of other logical systems. That is, employees of one corporation cannot view the logical systems of another corporation.

Logical systems include both master and user logical systems and their administrators. The roles and responsibilities of the master administrator and those of a user logical system administrator differ greatly. This differentiation of privileges and responsibilities is considered role-based administration and control.

NOTE: **To use the internal switch, which is optional, you must also configure an interconnect logical system. The interconnect logical system does not require an administrator.**

Related Documentation

- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Copyright © 2016, Juniper Networks, Inc.

Understanding the Fundamentals and Constraints of Logical Systems

This topic covers basic information about logical systems features and limitations.

- By default, logical systems delivers a master logical system, which exists at the root level. You can purchase licenses for logical systems that you intend to create with the total not exceeding 32.
- You can configure up to 32 security profiles, from 1 through 32, with ID 0 reserved for the internally configured default security profile. When the maximum number of security profiles is reached, if you want to add a new security profile, you must first delete one or more existing security profiles, commit the configuration, and then create the new security profile and commit it. You cannot add a new security profile and remove an existing one within a single configuration commit.

  If you want to add more than one new security profile, the same rule is true. You must first delete the equivalent number of of existing security profiles, commit the configuration, and then create the new security profiles and commit them.
- You can configure one or more master administrators to oversee administration of the device and the logical systems they configure.

  As master administrator for an SRX Series Services Gateway running logical systems, you have root control over the device, its resources, and the logical systems that you create. You allocate security, networking, and routing resources to user logical systems. You can configure one logical system to serve as an interconnect logical system virtual private LAN service (VPLS) switch. The interconnect logical system, which is not mandatory, does not require security resources. However, if you configure an interconnect logical system, you must bind a dummy security profile to it. The master administrator configures it and all lt-0/0/0 interfaces for it.
- A user logical system can have one or more administrators, referred to as user logical system administrators. The master administrator creates login accounts for these administrators and assigns them to a user logical system. Currently, the master administrator must configure all user logical system administrators. The first assigned user logical administrator cannot configure additional user logical system administrators for his logical system. As a user logical system administrator, you can configure the resources assigned to your user logical system, including logical interfaces assigned by the master administrator, routing instances and their routes, and security components. You can display configuration information only for your logical system.
- A logical system can include more than one routing instance based on available system resources.
- You cannot configure class of service on lt-0/0/0 interfaces.
- The trace and debug features are supported at the root level only.
- Commit rollback is supported at the root level only.
- Quality of service (QoS) classification across interconnected logical systems does not work.

Copyright © 2016, Juniper Networks, Inc.

- The master administrator can configure Application Layer Gateways (ALGs) at the root level. The configuration is inherited by all user logical systems. It cannot be configured discretely for user logical systems.
- The master administrator can configure IDP policies at the root level and then apply an IDP policy to a user logical system.
- Only the master administrator can create user accounts and login IDs for users for all logical systems. The master administrator creates these user accounts at the root level and assigns them to the appropriate user logical systems.
- The same name cannot be used in two separate logical systems. For example, if logical-system1 includes a user with Bob configured as the username, then other logical systems on the device cannot include a user with the username Bob.
- Configuration for users for all logical systems and all user logical systems administrators must be performed at the root level by the master administrator. A user logical system administrator cannot create other user logical system administrators or user accounts for their logical systems.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Understanding Licenses for Logical Systems on SRX Series Devices

This topic provides licensing information for SRX Series devices running logical systems. For general licensing information, such as how to install a license, see the *Installation and Upgrade Guide for Security Devices*.

By default, a device running logical systems delivers a master logical system at the root level. You can purchase licenses for other logical systems that you intend to create. If you intend to configure an interconnect logical system to use as a switch, it also requires a license.

Complications arise if the number of logical systems that you configure exceeds the number of licenses that you have purchased. The system will allow you to configure additional logical systems. However, when you attempt to commit their configurations, the system issues a warning message similar to the following: **Warning: 2 more license(s) are needed, logical system won't work without license!**. The message indicates the number of logical systems without licenses. We recommend that you do not configure more logical systems than the number of licenses you have purchased.

If you configure more logical systems than the number of licenses that you have purchased, the additional logical systems will not be activated until a license is available. The system will drop packets destined to them. They are inactive.

When a logical system is deleted, its license is freed up. That license is assigned to an inactive logical system, and the logical system is activated.

Copyright © 2016, Juniper Networks, Inc.

You can use the **show system license status logical-system all** command on the command-line interface (CLI) to determine which logical systems are active.

```
user@host>  show system license status logical-system all
logical system name      license status
root-logical-system      enabled
LSYS2              enabled
LSYS0              enabled
LSYS11             enabled
LSYS12             enabled
LSYS23             enabled
LSYS10             enabled
LSYS13             enabled
LSYS18             enabled
```

When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Understanding the Interconnect Logical System and Logical Tunnel Interfaces

This topic covers the interconnect logical system that serves as an internal virtual private LAN service (VPLS) switch connecting one logical system on the device to another. The topic also explains how logical tunnel (lt-0/0/0) interfaces are used to connect logical systems through the interconnect logical system.

A device running logical systems can use an internal VPLS switch to pass traffic without it leaving the device. The interconnect logical system switches traffic across logical systems that use it. Although a virtual switch is used typically, it is not mandatory. If you choose to use a virtual switch, you must configure the interconnect logical system. There can be only one interconnect logical system on a device.

For communication between logical systems on the device to occur, you must configure an lt-0/0/0 interface on each logical system that will use the internal switch, and you must associate it with its peer lt-0/0/0 interface on the interconnect logical system, effectively creating a logical tunnel between them. You define a peer relationship at each end of the tunnel when you configure the logical system's lt-0/0/0 interfaces.

Copyright © 2016, Juniper Networks, Inc.

You might want all logical systems on the device to be able to communicate with one another without using an external switch. Alternatively, you might want some logical systems to connect across the internal switch but not all of them.

The interconnect logical system does not require security resources assigned to it through a security profile. However, you must assign a dummy security profile containing no resources to the interconnect logical system. Otherwise you will not be able to successfully commit the configuration for it.

WARNING: **If you configure an lt-0/0/0 interface in any user logical system or the master logical system and you do not configure an interconnect logical system containing a peer lt-0/0/0 interface for it, the commit will fail.**

An SRX Series device running logical systems can be used in a chassis cluster. Each node has the same configuration, including the interconnect logical system.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related Documentation

- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 3342
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Understanding Logical Systems in the Context of Chassis Cluster on page 3369

Understanding Flow in Logical Systems for SRX Series Devices

This topic explains how packets are processed in flow sessions on SRX Series devices running logical systems. It describes how an SRX Series device running logical systems handles pass-through traffic in a single logical system and between logical systems. It also covers self-traffic as self-initiated traffic within a logical system and self-traffic terminated on another logical system. Before addressing logical systems, the topic provides basic information about the SRX Series architecture in with respect to packet processing and sessions. Finally, it addresses sessions and how to change session characteristics.

The concepts explained in this example rely on the topology shown in Figure 132.

Copyright © 2016, Juniper Networks, Inc.

Figure 132: Logical Systems, Their Virtual Routers, and Their Interfaces

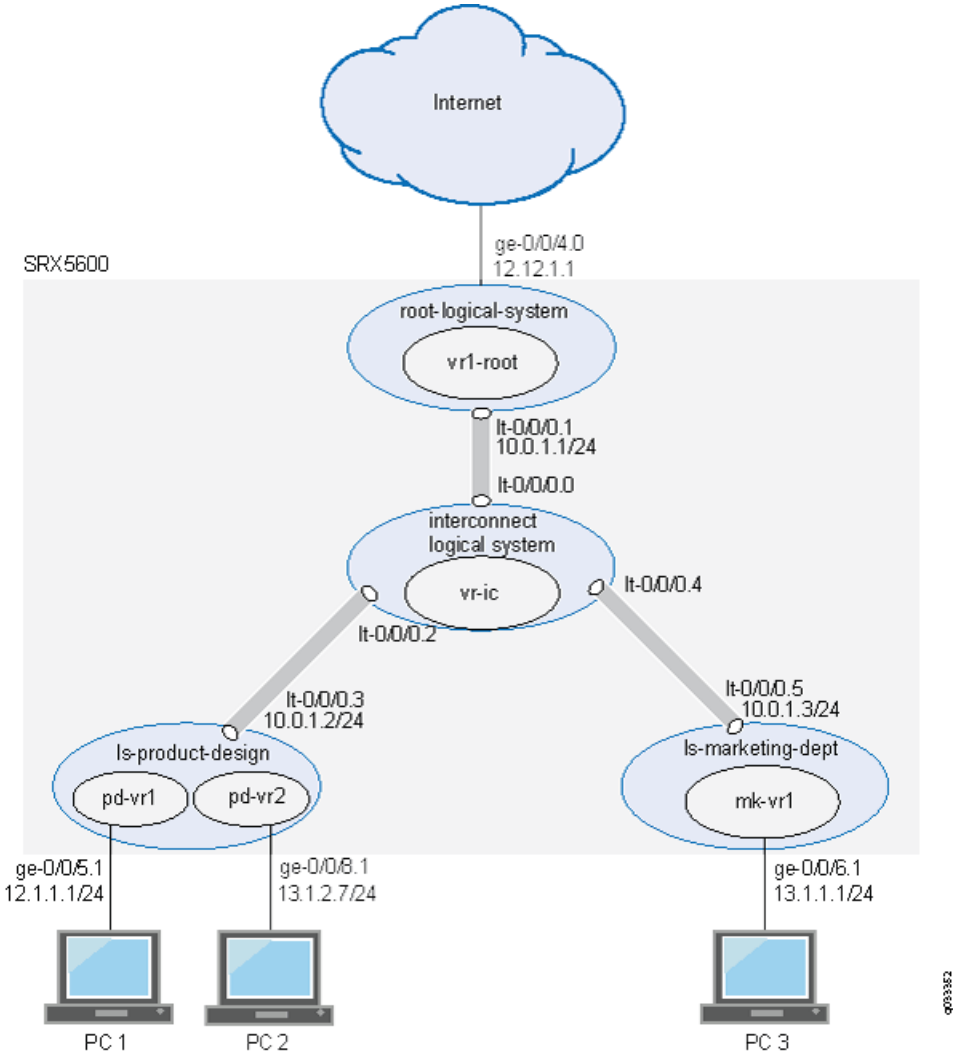

- Understanding Junos OS SRX Series Services Gateways Architecture on page 3213
- Session Creation for Devices Running Logical Systems on page 3214
- Understanding Flow on Logical Systems on page 3214
- Understanding Packet Classification on page 3214
- Handling Pass-Through Traffic for Logical Systems on page 3215
- Handling Self-Traffic on page 3216
- Understanding Session and Gate Limitation Control on page 3218
- Understanding Sessions on page 3218
- About Configuring Sessions on page 3218

Copyright © 2016, Juniper Networks, Inc.

Understanding Junos OS SRX Series Services Gateways Architecture

Junos OS is a distributed parallel processing high throughput and high performance system. The distributed parallel processing architecture of the services gateways includes multiple processors to manage sessions and run security and other services processing. This architecture provides greater flexibility and allows for high throughput and fast performance.

The SRX5000 line devices include I/O cards (IOC) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. One or more Services Processing Units (SPUs) run on an SPC.

These processing units have different responsibilities. All flow-based services for a packet are executed on a single SPU. Otherwise, however, the lines are not clearly divided in regard to the kinds of services that run on these processors. (For details on flow-based processing, see "Juniper Networks Devices Processing Overview" on page 1665.)

For example:

- An NPU processes packets discretely. It performs sanity checks and applies some screens that are configured for the interface, such as denial-of-service (DoS) screens, to the packet.
- An SPU manages the session for the packet flow and applies security features and other services to the packet. It also applies packet-based stateless firewall filters, classifiers, and traffic shapers to the packet.
- The system uses one processor as a central point to take care of arbitration and allocation of resources and distribute sessions in an intelligent way. The central point assigns an SPU to be used for a particular session when the first packet of its flow is processed.

These discrete, cooperating parts of the system, including the central point, each store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

This architecture allows the device to distribute processing of all sessions across multiple SPUs. It also allows an NPU to determine if a session exists for a packet, to check the packet, and to apply screens to it. How a packet is handled depends on whether it is the first packet of a flow.

Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that are established for the first packet of the packet stream when the flow session is established. Most packet processing occurs within a flow. For the distributed processing architecture of the services gateway, some packet-based processing, such as traffic shaping, occurs on the NPU. Some packet-based processing, such as application of classifiers to a packet, occurs on the SPU.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, Aplication Layer Gateway (ALG)s configured for it, if NAT should be

Copyright © 2016, Juniper Networks, Inc.

applied to translate the packet's source and/or destination IP address—are assessed for the first packet of a flow.

Session Creation for Devices Running Logical Systems

Session establishment for SRX Series devices running logical systems differs in minor ways from that of SRX series devices not running logical systems. Despite the complexities that logical systems introduce, traffic is handled in a manner similar to how it is handled on SRX Series devices not running logical systems. Flow-based packet processing, which is stateful, requires the creation of sessions. In considering flow based processing and session establishment for logical systems, it helps to think of each logical system on the device as a discrete device with respect to session establishment.

A session is created, based on routing and other classification information, to store information and allocate resources for a flow. Basically, a session is established when traffic enters a logical system interface, route lookup is performed to identify the next hop interface, and policy lookup is performed.

Optionally, logical systems enable you to configure an internal software switch. This virtual private LAN switch (VPLS) is implemented as an interconnect logical system. It enables both transit traffic and traffic terminated at a logical system to pass between logical systems. To enable traffic to pass between logical systems, logical tunnel (lt-0/0/0) interfaces across the interconnect logical system are used.

Communication between logical systems across the interconnect logical system requires establishment of two sessions: one for traffic that enters a logical system and exits its lt-0/0/0 interface, and one for traffic that enters the lt-0/0/0 interface of another logical system and either exits the device through one of its physical interface or is destined for it.

NOTE: Packet sequence occurs at the ingress and the egress interfaces. Packets traveling between logical systems might not be processed in the order in which they were received on the physical interface.

Understanding Flow on Logical Systems

To understand how traffic is handled for logical systems, it is helpful to consider each logical system as a discrete device.

NOTE: Traffic is processed for the master logical system in the same way as it is for user logical systems on the device.

Understanding Packet Classification

Packet classification is assessed the same way for SRX Series devices running with or without logical systems. Filters and class-of-service features are typically associated with an interface to influence which packets are allowed to transit the system and to

Copyright © 2016, Juniper Networks, Inc.

apply special actions to packets as needed. (Within a flow, some packet-based processing also takes place on an SPU.)

Packet classification is based on the incoming interface and performed at the ingress point. Traffic for a dedicated interface is classified to the logical system that contains that interface. Within the context of a flow, packet classification is based on both the physical interface and the logical interface.

Handling Pass-Through Traffic for Logical Systems

For SRX Series devices not running logical systems, pass-through traffic is traffic that enters and exits a device. You can think of pass-through traffic for logical systems similarly, but as having a larger dimension as a result of the nature of a multitenant device. For SRX Series devices running logical systems, pass-through traffic can exist within a logical system or between logical systems.

- Pass-Through Traffic Within a Logical System on page 3215
- Pass-Through Traffic Between Logical Systems on page 3215

Pass-Through Traffic Within a Logical System

For pass-through traffic within a logical system, traffic comes in on an interface belonging to one of the logical system's virtual routing instances, and it is sent to another of its virtual routing instances. To exit the device, the traffic is sent out an interface belonging to the second virtual routing instance. The traffic does not transit between logical systems but rather enters and exits the device in a single logical system. Pass-through traffic within a logical system is transmitted according to the routing tables in each of its routing instances.

Consider how pass-through traffic is handled within a logical system given the topology shown in Figure 132.

- When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
- Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1 with pd-vr2 identified as the next hop.
- A second route lookup is performed in pd-vr2 to identify the egress interface to use—in this case— ge-0/0/8.
- The packet is sent out ge-0/0/8 to the network.
- The security policy lookup is performed in ls-product-design, and one session is established.

Pass-Through Traffic Between Logical Systems

Pass-through traffic between logical systems is complicated by fact that each logical system has an ingress and an egress interface that the traffic must transit. It is as if traffic were coming into and going out from two devices.

Copyright © 2016, Juniper Networks, Inc.

Two sessions must be established for pass-through traffic between logical systems. (Note that policy lookup is performed in both logical systems).

- On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
- On the egress logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and its egress interface (a physical interface).

Consider how pass-through traffic is handled across logical systems in the topology shown in Figure 132.

- A session is established in the incoming logical system.
  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.
  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, which is the ingress interface in the ls-marketing-dept.
  - A session is established between ge-0/0/5 and lt-0/0/0.3.
- A session is established in the outgoing logical system.
  - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.
  - Packet processing continues in the ls-marketing-dept logical system.
  - To identify the egress interface, route lookup for the packet is performed in the mk-vr1 routing instances.
  - The outgoing interface is identified as ge-0/0/6, and the packet is transmitted from the interface to the network.

Handling Self-Traffic

Self-traffic is traffic that originates in a logical system on the device and is either sent out to the network from that logical system or is terminated on another logical system on the device.

Self-Initiated Traffic

Self-initiated traffic is generated from a source logical system context and forwarded directly to the network from the logical system interface.

Copyright © 2016, Juniper Networks, Inc.

The following process occurs:

- When a packet is generated in a logical system, a process for handling the traffic is started in the logical system.
- Route lookup is performed to identify the egress interface, and a session is established.
- The logical system performs a policy lookup and processes the traffic accordingly.
- If required, a management session is set up.

Consider how self-initiated traffic is handled across logical systems given the topology shown in Figure 132.

- A packet is generated in the ls-product-design logical system, and a process for handling the traffic is started in the logical system.
- Route lookup performed in pd-vr2 to identifies the egress interface as ge-0/0/8.
- A session is established.
- The packet is transmitted to the network from ge-0/0/8.

Traffic Terminated on a Logical System

When a packet enters the device on an interface belonging to a logical system and the packet is destined for another logical system on the device, the packet is forwarded between the logical systems in the same manner as is pass-through traffic. However, route lookup in the second logical system identifies the local egress interface as the packet destination. Consequently the packet is terminated on the second logical system as self-traffic.

- For terminated self-traffic, two policy lookups are performed, and two sessions are established.
  - On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).
  - On the destination logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and the local interface.

Consider how terminated self-traffic is handled across logical systems in the topology shown in Figure 132.

- A session is established in the incoming logical system.
  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.
  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.
  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, the ingress interface in the ls-marketing-dept.
  - A session is established between ge-0/0/5 and lt-0/0/0.3.

Copyright © 2016, Juniper Networks, Inc.

- A management session is established in the destination logical system.
  - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.
  - Packet processing continues in the ls-marketing-dept logical system.
  - Route lookup for the packet is performed in the mk-vr1 routing instance. The packet is terminated in the destination logical system as self-traffic.
  - A management session is established.

Understanding Session and Gate Limitation Control

The logical systems flow module provides session and gate limitation to ensure that these resources are shared fairly among the logical systems. Resources allocation and limitations for each logical system are specified in the security profile bound to the logical system.

- For session limiting, the system checks the first packet of a session against the maximum number of sessions configured for the logical system. If the maximum is reached, the device drops the packet and logs the event.
- For gate limiting, the device checks the first packet of a session against the maximum number of gates configured for the logical system. If the maximum number of gates for a logical system is reached, the device rejects the gate open request and logs the event.

Understanding Sessions

Sessions are created based on routing and other classification information to store information and allocate resources for a flow. You can change some characteristics of sessions, such as when a session is terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

About Configuring Sessions

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 1800 seconds. The default timeout for UDP is 60 seconds. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds. If no traffic uses the If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse.

You can affect the life of a session in the following ways:

- Age out sessions, based on how full the session table is.
- Set an explicit timeout for aging out TCP sessions.
- Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.
- You can configure sessions to accommodate other systems as follows:
  - Disable TCP packet security checks.

Copyright © 2016, Juniper Networks, Inc.

- Change the maximum segment size.

Related Documentation

- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 143

Understanding Master Logical Systems

- Understanding the Master Logical System and the Master Administrator Role on page 3221
- SRX Series Logical System Master Administrator Configuration Tasks Overview on page 3222

Understanding the Master Logical System and the Master Administrator Role

When, as a master administrator, you initialize an SRX Series device running logical systems, a master logical system is created at the root level. You can log in to the device as root and change the root password.

By default, all system resources are assigned to the master logical system, and the master administrator allocates them to the user logical systems.

As master administrator, you manage the device and all its logical systems. You also manage the master logical system and configure its assigned resources. There can be more than one master administrator managing a device running logical systems.

- The master administrator's role and main responsibilities include:
  - Creating user logical systems and configuring their administrators. You can create one or more user logical system administrators for each user logical system.
  - Creating login accounts for users for all logical systems and assigning them to the appropriate logical systems.
  - Configuring an interconnect logical system if you want to allow communication between logical systems on the device. The interconnect logical system acts as an internal switch. It does not require an administrator.

    To configure an interconnect logical system, you configure lt-0/0/0 interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for establishment of tunnels.
  - Configuring security profiles to provision portions of the system's security resources to user logical systems and the master logical system.

    Only the master administrator can create, change, and delete security profiles and bind them to logical systems.

Copyright © 2016, Juniper Networks, Inc.

NOTE: A user logical system administrator can configure interface, routing, and security resources allocated to his logical system.

- Creating logical interfaces to assign to user logical systems. (The user logical system administrator configures logical interfaces assigned to his logical system.)
- Viewing and managing user logical systems, as required, and deleting user logical systems. When a user logical system is deleted, its allocated reserved resources are released for use by other logical systems.
- Configuring IDP, AppTrack, application identification, and application firewall features. The master administrator can also use trace and debug at the root level, and he can perform commit rollbacks. The master administrator manages the master logical system and configures all the features that a user logical system administrator can configure for his or her own logical systems including routing instances, static routes, dynamic routing protocols, zones, security policies, screens, and firewall authentication.

Related Documentation

- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 3342

SRX Series Logical System Master Administrator Configuration Tasks Overview

This topic identities and describes the master administrator's tasks in the order in which they are performed.

An SRX Series device running logical systems is managed by a master administrator. The master administrator has the same capabilities as the root administrator of an SRX Series device not running logical systems. However, the master administrator's role and responsibilities extend beyond those of other SRX Series device administrators because an SRX Series device running logical systems is partitioned into discrete logical systems, each with its own resources, configuration, and management concerns. The master administrator is responsible for creating these user logical systems and provisioning them with resources.

For an overview of the master administrator's role and responsibilities, see "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

As the master administrator, you perform the following tasks to configure an SRX Series device running logical systems:

Copyright © 2016, Juniper Networks, Inc.

1. Configure a root password. Initially the master administrator logs in to the device as the root user without needing to specify a password. After you log in to the device, you must define a root password for later use.

   See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 3241 for configuration information.

2. Create user logical systems and their administrators and users. Optionally, create an interconnect logical system.

   For each user logical system that you want to configure on the device, you must create a logical system, define one or more administrators for it, and add users to it.

   The master administrator configures login accounts for user logical system administrators and users and associates them with the user logical system. A user logical system can have more than one administrator; the master administrator must define and add all user logical system administrators and add them to their user logical systems.

   The master administrator adds users to user logical systems on behalf of the user logical system administrator. For example, if you have created a user logical system for the product design department, you must create user accounts for the users who belong to that department and associate them with the user logical system. The user logical system administrator does not have the ability to do this. Rather, the user logical administrator tells you the user accounts that you must create and add for his logical system.

   If you intend to use an internal virtual private LAN service (VPLS) switch to allow logical systems to communicate with one another, you must create an interconnect logical system. An interconnect logical system does not require an administrator.

   - For configuration information, see "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242
   - For information on user logical system administrators, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 3227.
   - For information on the interconnect logical system, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 3210.

3. Configure one or more security profiles. Security profiles assign security resources to logical systems. You can assign a single security profile to more than one logical system if you intend to allocate the same kinds and amounts of resources to them.

   - For configuration information, see "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 3258.
   - For information on security profiles, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 3253.

4. Configure interfaces, routing instances, and static routes for logical systems, as appropriate.

   - If you plan to use an interconnect logical system, configure its logical tunnel interfaces and add them to its virtual routing instance.

Copyright © 2016, Juniper Networks, Inc.

- Configure interfaces for the master logical system. Optionally, create its logical tunnel interface to allow it to communicate with other logical systems on the device. Create a virtual routing instance for the master logical system and add its interfaces and static routes to it. Also configure logical interfaces for user logical systems with VLAN tagging.

  NOTE: **The master administrator tells the user logical system administrators which interfaces are assigned to their logical systems. It is the user logical system administrator's responsibility to configure their interfaces.**

- Optionally, configure logical tunnel interfaces for any user logical systems that you want to allow to communicate with one another using the internal VPLS switch.
- For configuration information, see "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 3342.
- For information about the interconnect logical system and logical tunnel (lt-0/0/0) interfaces, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 3210.

5. Enable CPU utilization control and configure the CPU control target and reserved CPU quotas for logical systems. See "Example: Configuring CPU Utilization (Master Administrators Only)" on page 3465.
6. Optionally, configure dynamic routing protocols for the master logical system. See "Example: Configuring OSPF Routing Protocol for the Master Logical System" on page 3350
7. Configure zones, security policies, and security features for the master logical system. See "Example: Configuring Security Features for the Master Logical System" on page 3271.
8. Configure IDP for the master logical system. See "Example: Configuring an IDP Policy for the Master Logical System" on page 3281.
9. Configure application firewall services on the master logical system. See "Understanding Logical System Application Firewall Services" on page 3287 and "Example: Configuring Application Firewall Services for a Master Logical System" on page 3288.
10. Configure a route-based VPN to secure traffic between a logical system and a remote site. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 3294.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 144

Understanding User Logical Systems

- User Logical System Configuration Overview on page 3225
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227
- Example: Configuring User Logical Systems on page 3228

User Logical System Configuration Overview

When the master administrator creates a user logical system, he assigns a user logical system administrator to manage it. A user logical system can have multiple user logical system administrators.

As a user logical system administrator, you can access and view resources in your user logical system but not those of other user logical systems or the master logical system. You can configure resources allocated to your user logical system, but you cannot modify the numbers of allocated resources.

The following procedure lists the tasks that the user logical system administrator performs to configure resources in the user logical system:

1. Log in to the user logical system with the login and password configured by the master administrator:

   a. Telnet or SSH to the management IP address configured on the device. Log in to the user logical system with the administrator login and password provided by the master administrator.

      You enter a UNIX shell in the user logical system configured by the master administrator.

   b. The presence of the > prompt indicates the CLI has started. The prompt is preceded by a string that contains your username, the hostname of the router, and the name of the user logical system. When the CLI starts, you are at the top level in operational mode. You enter configuration mode by entering the **configure** operational mode command. The CLI prompt changes from user@host: *logical-system* > to user@host: *logical-system*#.

      To exit the CLI and return to the UNIX shell, enter the **quit** command.

Copyright © 2016, Juniper Networks, Inc.

2. Configure the logical interfaces assigned to the user logical system by the master administrator. Configure one or more routing instances and the routing protocols and options within each instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 3360.

3. Configure security resources for the user logical system:

   a. Create zones for the user logical system and bind the logical interfaces to the zones. Address books can be created that are attached to zones for use in policies. See "Example: Configuring Zones for a User Logical System" on page 3301.

   b. Configure screen options at the zone level. See "Example: Configuring Screen Options for a User Logical System" on page 3304.

   c. Configure security policies between zones in the user logical system. See "Example: Configuring Security Policies in a User Logical System" on page 3308.

      Custom applications or application sets can be created for specific types of traffic. To create a custom application, use the **application** configuration statement at the [**edit applications**] hierarchy level. To create an application set, use the **application-set** configuration statement at the [**edit applications**] hierarchy level.

   d. Configure firewall authentication. The master administrator creates access profiles in the master logical system. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 3269.

      The user logical system administrator then configures a security policy that specifies firewall authentication for matching traffic and configures the type of authentication (pass-through or Web authentication), default access profile, and success banner. See "Example: Configuring Firewall Authentication for a User Logical System" on page 3313.

   e. Configure a route-based VPN tunnel to secure traffic between a user logical system and a remote site. The master administrator assigns a secure tunnel interface to the user logical system and configures IKE and IPsec SAs for the VPN tunnel. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 3294.

      The user logical system administrator then configures a route-based VPN tunnel. See "Example: Configuring a Route-Based VPN Tunnel in a User Logical System" on page 3335.

   f. Configure Network Address Translation (NAT). See "Example: Configuring Network Address Translation for a User Logical System" on page 3356.

   g. Enable IDP. The master administrator configures IDP policies at the root level and specifies an IDP policy in the security profile that is bound to a logical system. See "Example: Configuring an IDP Policy for a User Logical System" on page 3321.

      The user logical system administrator then enables IDP in a security policy. See "Example: Enabling IDP in a User Logical System Security Policy" on page 3323.

   h. Display or clear application system cache (ASC) entries. See "Understanding Logical System Application Identification Services" on page 3286.

 Copyright © 2016, Juniper Networks, Inc.

i. Configure application firewall services on a user logical system. See "Understanding Logical System Application Firewall Services" on page 3287 and "Example: Configuring Application Firewall Services for a User Logical System" on page 3327.

j. Configure the AppTrack application tracking tool. See "Example: Configuring AppTrack for a User Logical System" on page 3331.

Related Documentation

- Example: Configuring User Logical Systems on page 3228
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Understanding User Logical Systems and the User Logical System Administrator Role

Logical systems allow a master administrator to partition an SRX Series device into discrete contexts called user logical systems. User logical systems are self-contained, private contexts, separate both from one another and from the master logical system. A user logical system has its own security, networking, logical interfaces, routing configurations, and one or more user logical system administrators.

When the master administrator creates a user logical system, he assigns one or more user logical system administrators to manage it. A user logical system administrator has a view of the device that is limited to his logical system. Although a user logical system is managed by a user logical system administrator, the master administrator has a global view of the device and access to all user logical systems. If necessary, the master administrator can manage any user logical system on the device.

The role and responsibilities of a user logical system administrator differ from those of the master administrator. As a user logical system administrator, you can access, configure, and view the configuration for your user logical system resources, but not those of other user logical systems or the master logical system.

As a user logical system administrator, you can:

- Configure zones, address books, security policies, user lists, custom services, and so forth, for your user logical system environment, based on the resources allocated to it.

  For example, if the master administrator allocates 40 zones to your user logical system, you can configure and administer those zones, but you cannot change the allocated number.

- Configure routing instances and assign allotted interfaces to them. Create static routes and add them to your routing instances. Configure routing protocols.
- Configure, enable, and monitor application firewall policy on your user logical system.
- Configure AppTrack.

Copyright © 2016, Juniper Networks, Inc.

- View all assigned logical interfaces and configure their attributes. The attributes that you configure for logical interfaces for your user logical system cannot be seen by other user logical system administrators.
- Run operational commands for your user logical system.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 3342
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

Example: Configuring User Logical Systems

This example shows the configuration of interfaces, routing instances, zones, and security policies for user logical systems.

- Requirements on page 3228
- Overview on page 3228
- Configuration on page 3230
- Verification on page 3238

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Be sure you know which logical interfaces and optionally, which logical tunnel interface (and its IP address) are allocated to your user logical system by the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Overview

This example configures the ls-marketing-dept and ls-accounting-dept user logical systems shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example configures the parameters described in Table 332 and Table 333.

Copyright © 2016, Juniper Networks, Inc.

Table 332: ls-marketing-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/6.1	• IP address 13.1.1.1/24 • VLAN ID 800

Routing instance	mk-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/6.1 and lt-0/0/0.5
		• Static routes:
		  ◦ 12.1.1.0/24 next-hop 10.0.1.2
		  ◦ 14.1.1.0/24 next-hop 10.0.1.4
		  ◦ 12.12.1.0/24 next-hop 10.0.1.1

Zones	ls-marketing-trust	Bind to interface ge-0/0/6.1.
	ls-marketing-untrust	Bind to interface lt-0/0/0.5
Address books	marketing-internal	• Address marketers: 13.1.1.0/24 • Attach to zone ls-marketing-trust
	marketing-external	• Address design: 12.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, accounting • Attach to zone ls-marketing-untrust
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-marketing-trust • To zone: ls-marketing-untrust • Source address: marketers • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-marketing-untrust • To zone: ls-marketing-trust • Source address: otherlsys • Destination address: marketers • Application: any

Table 333: ls-accounting-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/7.1	• IP address 14.1.1.1/24 • VLAN ID 900

Copyright © 2016, Juniper Networks, Inc.

Table 333: ls-accounting-dept Logical System Configuration *(continued)*

Feature	Name	Configuration Parameters

Routing instance	acct-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/7.1 and lt-0/0/0.7
		• Static routes:
		  ◦ 12.1.1.0/24 next-hop 10.0.1.2
		  ◦ 13.1.1.0/24 next-hop 10.0.1.3
		  ◦ 12.12.1.0/24 next-hop 10.0.1.1

Zones	ls-accounting-trust	Bind to interface ge-0/0/7.1.
	ls-accounting-untrust	Bind to interface lt-0/0/0.7
Address books	accounting-internal	• Address accounting: 14.1.1.0/24 • Attach to zone ls-accounting-trust
	accounting-external	• Address design: 12.1.1.0/24 • Address marketing: 13.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, marketing • Attach to zone ls-accounting-untrust
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-accounting-trust • To zone: ls-accounting-untrust • Source address: accounting • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-accounting-untrust • To zone: ls-accounting-trust • Source address: otherlsys • Destination address: accounting • Application: any

Configuration

- Configuring the ls-marketing-dept User Logical System on page 3230
- Configuring the ls-accounting-dept User Logical System on page 3234

Configuring the ls-marketing-dept User Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/6 unit 1 family inet address 13.1.1.1/24
set interfaces ge-0/0/6 unit 1 vlan-id 800
set routing-instances mk-vr1 instance-type virtual-router
set routing-instances mk-vr1 interface ge-0/0/6.1
set routing-instances mk-vr1 interface lt-0/0/0.5
set routing-instances mk-vr1 routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances mk-vr1 routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
set routing-instances mk-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
set security zones security-zone ls-marketing-trust interfaces ge-0/0/6.1
set security zones security-zone ls-marketing-untrust interfaces lt-0/0/0.5
set security address-book marketing-external address design 12.1.1.0/24
set security address-book marketing-external address accounting 14.1.1.0/24
set security address-book marketing-external address others 12.12.1.0/24
set security address-book marketing-external address-set otherlsys address design
set security address-book marketing-external address-set otherlsys address accounting
set security address-book marketing-external attach zone ls-marketing-untrust
set security address-book marketing-internal address marketers 13.1.1.0/24
set security address-book marketing-internal attach zone ls-marketing-trust
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
  permit-all-to-otherlsys match source-address marketers
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
  permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
  permit-all-to-otherlsys match application any
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
  permit-all-to-otherlsys then permit
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
  permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
  permit-all-from-otherlsys match destination-address marketers
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
  permit-all-from-otherlsys match application any
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
  permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsmarketingadmin1@host:ls-marketing-dept> configure
   lsmarketingadmin1@host:ls-marketing-dept#
   ```

2. Configure the logical interface for a user logical system.

   ```
   [edit interfaces]
   lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 family inet address
     13.1.1.1/24
   lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 vlan-id 800
   ```

3. Configure the routing instance and assign interfaces.

   ```
   [edit routing-instances]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 instance-type virtual-router
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 interface ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 interface lt-0/0/0.5
```

4. Configure static routes.

```
[edit routing-instances]
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
  12.1.1.0/24 next-hop 10.0.1.2
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
  14.1.1.0/24 next-hop 10.0.1.4
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
  12.12.1.0/24 next-hop 10.0.1.1
```

5. Configure security zones and assign interfaces to each zone.

```
[edit security zones]
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-trust
  interfaces ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-untrust
  interfaces lt-0/0/0.5
```

6. Create address book entries.

```
[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal
  address marketers 13.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  address design 12.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  address accounting 14.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  address others 12.12.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  address-set otherlsys address design
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  address-set otherlsys address accounting
```

7. Attach address books to zones.

```
[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal
  attach zone ls-marketing-trust
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
  attach zone ls-marketing-untrust
```

8. Configure a security policy that permits traffic from the ls-marketing-trust zone to the ls-marketing-untrust zone.

```
[edit security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
  match source-address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
  match destination-address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
  match application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys then
  permit
```


Copyright © 2016, Juniper Networks, Inc.

9. Configure a security policy that permits traffic from the ls-marketing-untrust zone to the ls-marketing-trust zone.

```
[edit security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
  match source-address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
  match destination-address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
  match application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
  then permit
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsmarketingadmin1@host:ls-marketing-dept# show routing instances
  mk-vr1 {
    instance-type virtual-router;
    interface ge-0/0/6.1;
    interface lt-0/0/0.5;
    routing-options {
      static {
        route 12.1.1.0/24 next-hop 10.0.1.2;
        route 14.1.1.0/24 next-hop 10.0.1.4;
        route 12.12.1.0/24 next-hop 10.0.1.1;
      }
    }
  }
lsmarketingadmin1@host:ls-marketing-dept# show security
  address-book {
    marketing-external {
      address product-designers 12.1.1.0/24;
      address accounting 14.1.1.0/24;
      address others 12.12.1.0/24;
      address-set otherlsys {
        address product-designers;
        address accounting;
      }
      attach {
        zone ls-marketing-untrust;
      }
    }
    marketing-internal {
      address marketers 13.1.1.0/24;
      attach {
        zone ls-marketing-trust;
      }
    }
  }
    policies {
      from-zone ls-marketing-trust to-zone ls-marketing-untrust {
        policy permit-all-to-otherlsys {
          match {
```

Copyright © 2016, Juniper Networks, Inc.

```
            source-address marketers;
            destination-address otherlsys;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone ls-marketing-untrust to-zone ls-marketing-trust {
        policy permit-all-from-otherlsys {
          match {
            source-address otherlsys;
            destination-address marketers;
            application any;
          }
          then {
            permit;
          }
        }
      }
    }
    zones {
      security-zone ls-marketing-trust {
        interfaces {
          ge-0/0/6.1;
        }
      }
      security-zone ls-marketing-untrust {
        interfaces {
          lt-0/0/0.5;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the ls-accounting-dept User Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/7 unit 1 family inet address 14.1.1.1/24
set interfaces ge-0/0/7 unit 1 vlan-id 900
set routing-instances acct-vr1 instance-type virtual-router
set routing-instances acct-vr1 interface ge-0/0/7.1
set routing-instances acct-vr1 interface lt-0/0/0.7
set routing-instances acct-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
set routing-instances acct-vr1 routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances acct-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set security address-book accounting-internal address accounting 14.1.1.0/24
set security address-book accounting-internal attach zone ls-accounting-trust
set security address-book accounting-external address design 12.1.1.0/24
```


Copyright © 2016, Juniper Networks, Inc.

```
set security address-book accounting-external address marketing 13.1.1.0/24
set security address-book accounting-external address others 12.12.1.0/24
set security address-book accounting-external address-set otherlsys address design
set security address-book accounting-external address-set otherlsys address marketing
set security address-book accounting-external attach zone ls-accounting-untrust
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
  permit-all-to-otherlsys match source-address accounting
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
  permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
  permit-all-to-otherlsys match application any
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
  permit-all-to-otherlsys then permit
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
  permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
  permit-all-from-otherlsys match destination-address accounting
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
  permit-all-from-otherlsys match application any
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
  permit-all-from-otherlsys then permit
set security zones security-zone ls-accounting-trust interfaces ge-0/0/7.1
set security zones security-zone ls-accounting-untrust interfaces lt-0/0/0.7
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsaccountingadmin1@host:ls-accounting-dept> configure
   lsaccountingadmin1@host:ls-accounting-dept#
   ```

2. Configure the logical interface for a user logical system.

   ```
   [edit interfaces]
   lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 family inet
     address 14.1.1.1/24
   lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 vlan-id 900
   ```

3. Configure the routing instance and assign interfaces.

   ```
   [edit routing-instances]
   lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 instance-type
     virtual-router
   lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 interface ge-0/0/7.1
   lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 interface lt-0/0/0.7
   ```

4. Configure static routes.

   ```
   [edit routing-instances]
   lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static
     route 12.1.1.0/24 next-hop 10.0.1.2
   lsaccountingadmin1@host:ls-accounting-deptt# set acct-vr1 routing-options static
     route 13.1.1.0/24 next-hop 10.0.1.3
   ```

Copyright © 2016, Juniper Networks, Inc.

```
lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static
  route 12.12.1.0/24 next-hop 10.0.1.1
```

5. Configure security zones and assign interfaces to each zone.

```
[edit security zones]
lsaccountingadmin1@host:ls-accounting-dept# set security-zone ls-accounting-trust
  interfaces ge-0/0/7.1
lsaccountingadmin1@host:ls-accounting-dept# set security-zone
  ls-accounting-untrust interfaces lt-0/0/0.7
```

6. Create address book entries.

```
[edit security]
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal
  address accounting 14.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external address design 12.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external address marketing 13.1.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external address others 12.12.1.0/24
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external address-set otherlsys address design
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external address-set otherlsys address marketing
```

7. Attach address books to zones.

```
[edit security]
lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal
  attach zone ls-accounting-trust
lsaccountingadmin1@host:ls-accounting-dept# set address-book
  accounting-external attach zone ls-accounting-untrust
```

8. Configure a security policy that permits traffic from the ls-accounting-trust zone to the ls-accounting-untrust zone.

```
[edit security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust]
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
  match source-address accounting
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
  match destination-address otherlsys
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
  match application any
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
  then permit
```

9. Configure a security policy that permits traffic from the ls-accounting-untrust zone to the ls-accounting-trust zone.

```
[edit security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust]
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
  match source-address otherlsys
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
  match destination-address accounting
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
  match application any
```


Copyright © 2016, Juniper Networks, Inc.

```
lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
  then permit
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsaccountingadmin1@host:ls-accounting-dept# show routing-instances
acct-vr1 {
  instance-type virtual-router;
  interface ge-0/0/7.1;
  interface lt-0/0/0.7;
  routing-options {
    static {
      route 12.12.1.0/24 next-hop 10.0.1.1;
      route 12.1.1.0/24 next-hop 10.0.1.2;
      route 13.1.1.0/24 next-hop 10.0.1.3;
    }
  }
}
lsaccountingadmin1@host:ls-accounting-dept# show security
address-book {
  accounting-internal {
    address accounting 14.1.1.0/24;
    attach {
      zone ls-accounting-trust;
    }
  }
  accounting-external {
    address design 12.1.1.0/24;
    address marketing 13.1.1.0/24;
    address others 12.12.1.0/24;
    address-set otherlsys {
      address design;
      address marketing;
    }
    attach {
      zone ls-accounting-untrust;
    }
  }
}
  policies {
    from-zone ls-accounting-trust to-zone ls-accounting-untrust {
      policy permit-all-to-otherlsys {
        match {
          source-address accounting;
          destination-address otherlsys;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone ls-accounting-untrust to-zone ls-accounting-trust {
```

Copyright © 2016, Juniper Networks, Inc.

```
        policy permit-all-from-otherlsys {
          match {
            source-address otherlsys;
            destination-address accounting;
            application any;
          }
          then {
            permit;
          }
        }
      }
    }
    zones {
      security-zone ls-accounting-trust {
        interfaces {
          ge-0/0/7.1;
        }
      }
      security-zone ls-accounting-untrust {
        interfaces {
          lt-0/0/0.7;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Policy Configuration on page 3238

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation

- User Logical System Configuration Overview on page 3225
- Understanding Logical System Interfaces and Routing Instances on page 3341
- Understanding Logical System Zones on page 3299
- Understanding Logical System Security Policies on page 3306

Copyright © 2016, Juniper Networks, Inc.

PART 45

Getting Started for Master Administrators

- Configuring Device for Master Logical Systems on page 3241

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 145

Configuring Device for Master Logical Systems

- Example: Configuring a Root Password for the Device (Master Administrators Only) on page 3241
- Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only) on page 3242

Example: Configuring a Root Password for the Device (Master Administrators Only)

- Requirements on page 3241
- Overview on page 3241
- Configuration on page 3241

Requirements

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 to understand how this task fits into the overall configuration process.

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

The Junos OS software is installed on the router before it is delivered from the factory. When you power on your router, it is ready for you to configure. Initially you log in as *root* user without using a password.

After you log in, you can configure a password for the root user, or, in logical systems terms, the master administrator. The master administrator has root privileges over the device.

Configuration

- Configuring the Root Password on page 3242

Copyright © 2016, Juniper Networks, Inc.

Configuring the Root Password

Step-by-Step Procedure

- Configure a root password for the device.

  user@host# **set system root-authentication Talk22rt6**

Related Documentation

- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)

This example shows how to create user logical systems and assign administrators to them. It shows how to add users to a user logical system. And the example shows how to create an interconnect logical system, which is optional.

NOTE: Only the master administrator can create user login accounts for administrators and users. If a user logical system administrator wants to add users to his logical system, he must convey the information to the master administrator, who will add the users.

- Requirements on page 3242
- Overview on page 3242
- Configuration on page 3244
- Verification on page 3248

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 to understand how this task fits into the overall configuration process.

This example is for a company that includes product design, marketing, and accounting departments. The company wants to curtail hardware and energy costs, but not at the risk of exposing data across departments or to the Internet.

Each department has its own security requirements in regard both to other departments and to the Internet. To meet its requirements for cost control without forfeiting security, the company deploys the SRX5600 device. The master administrator configures three user logical systems giving each department a logical device that is private and fully secured.

This topic covers how to:

Copyright © 2016, Juniper Networks, Inc.

- Create user logical systems and an interconnect logical system that is used as an internal VPLS switch to allow traffic to pass from one logical system to another.
- Create administrators for user logical systems other than the interconnect logical system. A user logical system can have more than one administrator. The interconnect logical system does not require an administrator.
- Add users to a user logical system.

NOTE: This example shows how to configure only two users—lsdesignuser1 and lsdesignuser2. In reality, every user logical system will include many users that would require configurations similar to those shown in this example.

Figure 133 shows an SRX5600 device deployed and configured for logical systems. The configuration examples reflect this deployment.

Figure 133: SRX Series Device Configured for Logical Systems

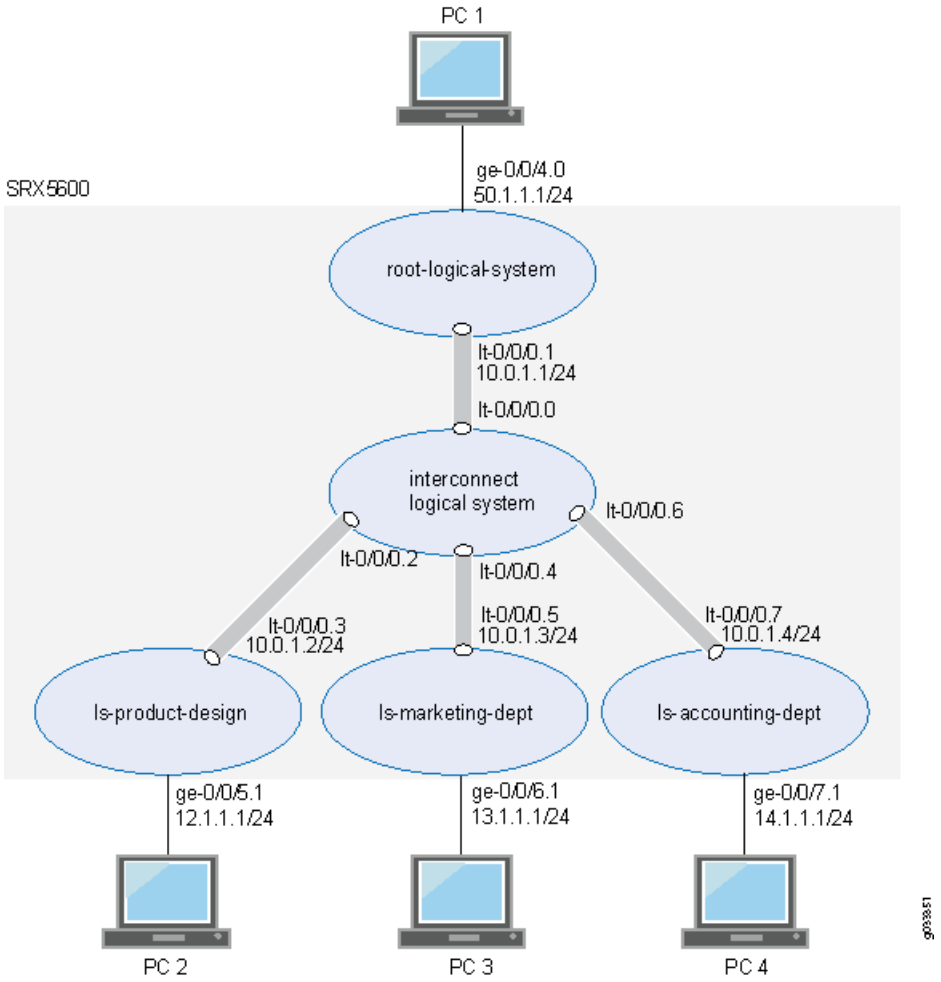

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System on page 3244

Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems ls-product-design
set system login class ls-design-admin logical-system ls-product-design
set system login class ls-design-admin permissions all
set system login user lsdesignadmin1 full-name lsdesignadmin1
set system login user lsdesignadmin1 class ls-design-admin
set system login user lsdesignadmin1 authentication encrypted-password "$ABC123"
set system login class ls-design-user logical-system ls-product-design
set system login class ls-design-user permissions view
set system login user lsdesignuser1 full-name lsdesignuser1
set system login user lsdesignuser1 class ls-design-user
set system login user lsdesignuser1 authentication encrypted-password "$ABC123"
set system login user lsdesignuser2 full-name lsdesignuser2
set system login user lsdesignuser2 class ls-design-user
set system login user lsdesignuser2 authentication encrypted-password "$ABC123"
set logical-systems ls-marketing-dept
set system login class ls-marketing-admin logical-system ls-marketing-dept
set system login class ls-marketing-admin permissions all
set system login user lsmarketingadmin1 class ls-marketing-admin
set system login user lsmarketingadmin1 full-name lsmarketingadmin1
set system login user lsmarketingadmin1 authentication encrypted-password "$ABC123"
set system login user lsmarketingadmin2 full-name lsmarketingadmin2
set system login user lsmarketingadmin2 class ls-marketing-admin
set system login user lsmarketingadmin2 authentication encrypted-password "$ABC123"
set logical-systems ls-accounting-dept
set system login class ls-accounting-admin logical-system ls-accounting-dept
set system login class ls-accounting-admin permissions all
set system login user lsaccountingadmin1 full-name lsaccountingadmin1
set system login user lsaccountingadmin1 class ls-accounting-admin
set system login user lsaccountingadmin1 authentication encrypted-password "$ABC123"
set logical-systems interconnect-logical-system
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Create the first user logical system and define its administrator.

   a. Create the user logical system.

      ```
      [edit]
      user@host# set logical-systems ls-product-design
      ```


Copyright © 2016, Juniper Networks, Inc.

b. Assign the user login class to the user logical system.

```
[edit system]
user@host# set login class ls-design-admin logical-system ls-product-design
```

c. Create the login class to give the user logical system administrator full permission over the user logical system.

```
[edit system]
user@host# set login class ls-design-admin permissions all
```

d. Assign a full name to the user logical system administrator.

```
[edit system]
user@host# set login user lsdesignadmin1 full-name lsdesignadmin1
```

e. Associate the login class with the user logical system administrator to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login user lsdesignadmin1 class ls-design-admin
```

f. Create a user login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsdesignadmin1 authentication plain-text-password
New password: Talk1234
Retype new password: Talk1234
```

2. Configure the first user for the logical system.

a. Configure the user login class and assign it to the user logical system.

```
[edit system]
user@host# set login class ls-design-user logical-system ls-product-design
```

b. To give the first user the ability to see the logical system's resources and settings but not change them, assign **view** as the permission to the login class.

```
[edit system]
user@host# set login class ls-design-user permissions view
```

c. Assign a full name to the logical system user.

```
[edit system]
user@host# set login user lsdesignuser1 full-name lsdesignuser1
```

d. Associate the login class with the user to allow the user to log in to the user logical system.

```
user@host# set login user lsdesignuser1 class ls-design-user
```

e. Create a user login password for the user.

```
[edit system]
user@host# set login user lsdesignuser1 authentication plain-text-password
New password: Talk4234
```

Copyright © 2016, Juniper Networks, Inc.

```
Retype new password: Talk4234
```

3. Create the second user for logical system ls-product-design.

   a. Assign a full name to the user.

   ```
   [edit system]
   user@host# set login user lsdesignuser2 full-name lsdesignuser2
   ```

   b. Associate the user with the login class to allow the user to log in to the user logical system.

   ```
   user@host# set login user lsdesignuser2 class ls-design-user
   ```

   c. Create a user login password.

   ```
   [edit system]
   user@host# set login user lsdesignuser2 authentication plain-text-password
   New password: Talk9234
   Retype new password: Talk9234
   ```

4. Create the second user logical system and define its administrator.

   a. Create the user logical system.

   ```
   [edit]
   user@host# set logical-systems ls-marketing-dept
   ```

   b. Configure the user login class and assign it to the user logical system.

   ```
   [edit system]
   user@host# set login class ls-marketing-admin logical-system ls-marketing-dept
   ```

   c. To give the user logical system administrator control over the user logical system, assign **all** as the permissions to the login class.

   ```
   [edit system]
   user@host# set login class ls-marketing-admin permissions all
   ```

   d. Assign a full name to the user logical system administrator.

   ```
   [edit system]
   user@host# set login user lsmarketingadmin1 full-name lsmarketingadmin1
   ```

   e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

   ```
   [edit system]
   user@host# set login user lsmarketingadmin1 class ls-marketing-admin
   ```

   f. Create a user login password for the user logical system administrator.

   ```
   [edit system]
   user@host# set login user lsmarketingadmin1 authentication plain-text-password
   New password: Talk2345
   Retype new password: Talk2345
   ```

5. Create a second user logical system administrator for the ls-marketing-dept logical system.

   a. Assign a full name to the user logical system administrator.

Copyright © 2016, Juniper Networks, Inc.

```
[edit system]
user@host# set login user lsmarketingadmin2 full-name lsmarketingadmin2
```

b. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login lsmarketingadmin2 class ls-marketing-admin
```

c. Create a user login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsmarketingadmin2 authentication plain-text-password
New password: Talk6345
Retype new password: Talk6345
```

6. Create the third user logical system and define its administrator.

a. Create the user logical system.

```
[edit]
user@host# set logical-systems ls-accounting-dept
```

b. Configure the user login class and assign it to the user logical system.

```
[edit system]
user@host# set login class ls-accounting-admin logical-system
  ls-accounting-dept
```

c. To give the user logical system administrator control over the user logical system, assign permissions to the login class.

```
[edit system]
user@host# set login class ls-accounting-admin permissions all
```

d. Assign a full name to the user logical system administrator.

```
[edit system]
user@host# set login user lsaccountingadmin1 full-name lsaccountingadmin1
```

e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

```
[edit system]
user@host# set login user lsaccountingadmin1 class ls-accounting-admin
```

f. Create a login password for the user logical system administrator.

```
[edit system]
user@host# set login user lsaccountingadmin1 authentication
  plain-text-password
New password: Talk5678
Retype new password: Talk5678
```

7. Configure an interconnect logical system to allow logical systems to pass traffic from one to another.

```
user@host# set logical-systems interconnect-logical-system
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command to verify that the logical systems were created. Also enter the **show system login class** command for each class that you defined.

To ensure that the logical systems administrators were created, enter the **show system login user** command.

If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems ?
interconnect-logical-system;
ls-accounting-dept;
ls-marketing-dept;
ls-product-design;

user@host# show system login class ls-design-admin
logical-system ls-product-design;
permissions all;

user@host# show system login class ls-design-user
logical-system ls-product-design
permissions view;

user@host show system login class ls-marketing-admin
logical-system ls-marketing-dept;
permissions all;

user@host show system login class ls-accounting-admin
logical-system ls-accounting-dept;
permissions all;

user@host show system login user ?
lsaccountingadmin1 lsaccountingadmin1
lsdesignadmin1 lsdesignadmin1
lsdesignuser2 lsdesignuser2
lsmarketingadmin1 lsmarketingadmin1
lsmarketingadmin2 lsmarketingadmin2
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying User Logical Systems and Login Configurations from the Master Logical System on page 3248
- Verifying User Logical Systems and Login Configurations Using Telnet on page 3249

Verifying User Logical Systems and Login Configurations from the Master Logical System

Purpose Verify that the user logical systems exist and that you, as the master administrator, can enter them from root. Return from a user logical system to the master logical system.

Action From operational mode, enter the following command:

```
root@host> set cli logical-system ls-product-design
```

 Copyright © 2016, Juniper Networks, Inc.

```
Logical system:ls-product-design
root@host:ls-product-design>

root@host:ls-product-design> clear cli logical-system
Cleared default logical system
root@host>

root@host> set cli logical-system ls-marketing-dept
Logical system:ls-marketing-dept
root@host:ls-marketing-dept>

root@host:ls-marketing-dept> clear cli logical-system
Cleared default logical system
root@host>

root@host> set cli logical-system ls-accounting-dept
Logical system:ls-accounting-dept
root@host:ls-accounting-dept>

root@host:ls-accounting-dept> clear cli logical-system
Cleared default logical system
root@host>
```

Verifying User Logical Systems and Login Configurations Using Telnet

Purpose Verify that the user logical systems you created exist and that the administrators' login IDs and passwords that you created are correct.

Action Use Telnet to log in to each user logical system as its user administrator would do.

1. Run Telnet specifying the IP address of your SRX Series device. For example:

   ```
   telnet 10.11.11.19
   ```

2. Enter the login ID and password for the administrator for one of the user logical systems that you created. After you log in, the prompt shows the administrator name. Notice how this result differs from the result produced when you log in to the user logical system from the master logical system at root. Repeat this procedure for all of your user logical systems.

   ```
   login: lsdesignadmin1
   Password: Talk1234
   lsdesignadmin1@host: ls-product-design>
   ```

Related Documentation

- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258
- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 3342

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 46

Configuring Security Features

- Configuring Master Logical System Security Profiles on page 3253
- Configuring Master Logical System Security Features on page 3267
- Configuring User Logical System Security Features on page 3299

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 146

Configuring Master Logical System Security Profiles

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

Understanding Logical System Security Profiles (Master Administrators Only)

Logical systems allow you to virtually divide a supported SRX Series device into multiple devices, isolating one from another, securing them from intrusion and attacks, and protecting them from faulty conditions outside their own contexts. To protect logical systems, security resources are configured in a manner similar to how they are configured for a discrete device. However, as the master administrator, you must allocate the kinds and amounts of security resources to logical systems. The logical system administrator allocates resources for his own logical system.

An SRX Series device running logical systems can be partitioned into user logical systems, an interconnect logical system, if desired, and the default master logical system. When the system is initialized, the master logical system is created at the root level. All system resources are assigned to it, effectively creating a default master logical system security profile. To distribute security resources across logical systems, the master administrator creates security profiles that specify the kinds and amounts of resources to be allocated to a logical system that the security profile is bound to. Only the master administrator can configure security profiles and bind them to logical systems. The user logical system administrator configures these resources for his or her logical system.

Logical systems are defined largely by the resources allocated to them, including security components, interfaces, routing instances, static routes, and dynamic routing protocols. When the master administrator configures a user logical system, he binds a security profile to it. Any attempt to commit a configuration for a user logical system without a security profile bound to it will fail.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Logical Systems Security Profiles on page 3254
- How the System Assesses Resources Assignment and Use Across Logical Systems on page 3254
- Cases: Assessments of Reserved Resources Assigned Through Security Profiles on page 3256

Logical Systems Security Profiles

As master administrator, you can configure a single security profile to assign resources to a specific logical system, use the same security profile for more than one logical system, or use a mix of both methods. You can configure up to 32 security profiles on an SRX Series device running logical systems. When you reach the limit, you must delete a security profile and commit the configuration change before you can create and commit another security profile. In many cases fewer security profiles are needed because you might bind a single security profile to more than one logical system.

Security profiles allow you to:

- Share the device's resources, including policies, zones, addresses and address books, flow sessions, and various forms of NAT, among all logical systems appropriately. You can dedicate various amounts of a resource to the logical systems and allow them to compete for use of the free resources.

  Security profiles protect against one logical system exhausting a resource that is required at the same time by other logical systems. Security profiles protect critical system resources and maintain a fair level of performance among user logical systems when the device is experiencing heavy traffic flow. They defend against one user logical system dominating the use of resources and depriving other user logical systems of them.

- Configure the device in a scalable way to allow for future creation of additional user logical systems.

You must delete a logical system's security profile before you delete that logical system.

How the System Assesses Resources Assignment and Use Across Logical Systems

To provision a logical system with security resources, you, as a master administrator, configure a security profile that specifies for each resource:

- A reserved quota that guarantees that the specified resource amount is always available to the logical system.

- A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems must compete for global resources.

Copyright © 2016, Juniper Networks, Inc.

If a reserved quota is not configured for a resource, the default value is 0. If a maximum allowed quota is not configured for a resource, the default value is the global system quota for the resource (global system quotas are platform-dependent). The master administrator must configure appropriate maximum allowed quota values in the security profiles so the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device. The master administrator must configure the appropriate maximum-allowed quota values in the security profiles so that the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device.

The system maintains a count of all allocated resources that are reserved, used, and made available again when a logical system is deleted. This count determines whether resources are available to use for new logical systems or to increase the amount of the resources allocated to existing logical systems through their security profiles.

When a user logical system is deleted, its reserved resource allocations are released for use by other logical systems.

Resources configured in security profiles are characterized as static modular resources or dynamic resources. For static resources, we recommend setting a maximum quota for a resource equal or close to the amount specified as its reserved quota, to allow for scalable configuration of logical systems. A high maximum quota for a resource might give a logical system greater flexibility through access to a larger amount of that resource, but it would constrain the amount available to allocate to a new user logical system.

The difference between reserved and maximum allowed amounts for a dynamic resource is not important because dynamic resources are aged out and do not deplete the pool available for assignment to other logical systems.

The following resources can be specified in a security profile:

- Security policies, including schedulers
- Security zones
- Addresses and address books for security policies
- Application firewall rule sets
- Application firewall rules
- Firewall authentication
- Flow sessions and gates
- NAT, including:
  - Cone NAT bindings
  - NAT destination rule
  - NAT destination pool
  - NAT IP address in source pool without Port Address Translation (PAT)

Copyright © 2016, Juniper Networks, Inc.

NOTE: IPv6 addresses in IPv6 source pools without PAT are not included in security profiles.

- NAT IP address in source pool with PAT
- NAT port overloading
- NAT source pool
- NAT source rule
- NAT static rule

NOTE: All resources except flow sessions are static.

You can modify a logical system security profile dynamically while the security profile is assigned to other logical systems. However, to ensure that the system resource quota is not exceeded, the system takes the following actions:

- If a static quota is changed, system daemons that maintain logical system counts for resources specified in security profiles revalidate the security profile. This check identifies the number of resources assigned across all logical systems to determine whether the allocated resources, including their increased amounts, are available.

  These quota checks are the same quota checks that the system performs when you add a new user logical system and bind a security profile to it. The are also performed when you bind a different security profile from the security profile that is presently assigned to it to an existing user logical system (or the master logical system).

- If a dynamic quota is changed, no check is performed, but the new quota is imposed on future resource usage.

Cases: Assessments of Reserved Resources Assigned Through Security Profiles

To understand how the system assesses allocation of reserved resources through security profiles, consider the following three cases that address allocation of one resource, zones. To keep the example simple, 10 zones are allocated in security-profile-1: 4 reserved zones and 6 maximum zones. This example assumes that the full maximum amount specified—six zones—is available for the user logical systems. The system maximum number of zones is 10.

These cases address configuration across logical systems. They test to see whether a configuration will succeed or fail when it is committed based on allocation of zones.

Table 334 shows the security profiles and their zone allocations.

Copyright © 2016, Juniper Networks, Inc.

Table 334: Security Profiles Used for Reserved Resource Assessments

Two Security Profiles Used in the Configuration Cases
security-profile-1 • zones reserved quota = 4 • zones maximum quota = 6 NOTE: Later the master administrator dynamically increases the reserved zone count specified in this profile.
master-logical-system-profile • zones maximum quota = 10 • no reserved quota

Table 335 shows three cases that illustrate how the system assesses reserved resources for zones across logical systems based on security profile configurations.

- The configuration for the first case succeeds because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 8, which is less than the system maximum resource quota.
- The configuration for the second case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.
- The configuration for the third case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.

Table 335: Reserved Resource Allocation Assessment Across Logical Systems

Reserved Resource Quota Checks Across Logical Systems
Example 1: Succeeds This configuration is within bounds: 4+4+0=8, maximum capacity =10. Security Profiles Used • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The master-logical-system-profile profile is used exclusively for the master logical system. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • master-logical-system = 0 reserved zones.

Copyright © 2016, Juniper Networks, Inc.

Table 335: Reserved Resource Allocation Assessment Across Logical Systems *(continued)*

Reserved Resource Quota Checks Across Logical Systems
Example 2: Fails This configuration is out of bounds: 4+4+4=12, maximum capacity =10. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • master-logical-system = 0 reserved zones. • new-user-logical-system = 4 reserved zones. Security Profiles • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The master-logical-system-profile is bound to the master logical system and used exclusively for it. • The master administrator configures a new user logical system called new-user-logical-system and binds security-profile-1 to it.
Example 3: Fails This configuration is out of bounds: 6+6=12, maximum capacity =10. The master administrator modifies the reserved zones quota in security-profile-1, increasing the count to 6. • user-logical-system-1 = 6 reserved zones. • user-logical-system-2 = 6 reserved zones. • master-logical-system = 0 reserved zones.

Related Documentation

- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258
- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Example: Configuring Logical Systems Security Profiles (Master Administrators Only)

This example shows how a master administrator configures three logical system security profiles to assign to user logical systems and the master logical system to provision them with security resources.

- Requirements on page 3259
- Overview on page 3259
- Configuration on page 3259
- Verification on page 3265

Copyright © 2016, Juniper Networks, Inc.

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 to understand how this task fits into the overall configuration process.

Overview

This example shows how to configure security profiles for the following logical systems:

- The root-logical-system logical system. The security profile master-profile is assigned to the master, or root, logical system.
- The ls-product-design logical system. The security profile ls-design-profile is assigned to the logical system.
- The ls-marketing-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.
- The ls-accounting-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.
- The interconnect-logical-system, if you use one. You must assign a dummy, or null, security profile to it.

This configuration relies on the deployment shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

Configuration

- Configuring Logical System Security Profiles on page 3259

Configuring Logical System Security Profiles

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system security-profile master-profile policy maximum 65
set system security-profile master-profile policy reserved 60
set system security-profile master-profile zone maximum 22
set system security-profile master-profile zone reserved 17
set system security-profile master-profile flow-session maximum 3000
set system security-profile master-profile flow-session reserved 2100
set system security-profile master-profile nat-nopat-address maximum 115
set system security-profile master-profile nat-nopat-address reserved 100
set system security-profile master-profile nat-static-rule maximum 125
set system security-profile master-profile nat-static-rule reserved 100
set system security-profile master-profile idp
set system security-profile master-profile logical-system root-logical-system
set system security-profile ls-accnt-mrkt-profile policy maximum 65
```

Copyright © 2016, Juniper Networks, Inc.

```
set system security-profile ls-accnt-mrkt-profile policy reserved 60
set system security-profile ls-accnt-mrkt-profile zone maximum 22
set system security-profile ls-accnt-mrkt-profile zone reserved 17
set system security-profile ls-accnt-mrkt-profile flow-session maximum 2500
set system security-profile ls-accnt-mrkt-profile flow-session reserved 2000
set system security-profile ls-accnt-mrkt-profile nat-nopat-address maximum 125
set system security-profile ls-accnt-mrkt-profile nat-nopat-address reserved 100
set system security-profile ls-accnt-mrkt-profile nat-static-rule maximum 125
set system security-profile ls-accnt-mrkt-profile nat-static-rule reserved 100
set system security-profile ls-accnt-mrkt-profile logical-system ls-marketing-dept
set system security-profile ls-accnt-mrkt-profile logical-system ls-accounting-dept
set system security-profile ls-design-profile policy maximum 50
set system security-profile ls-design-profile policy reserved 40
set system security-profile ls-design-profile zone maximum 10
set system security-profile ls-design-profile zone reserved 5
set system security-profile ls-design-profile flow-session maximum 2500
set system security-profile ls-design-profile flow-session reserved 2000
set system security-profile ls-design-profile nat-nopat-address maximum 120
set system security-profile ls-design-profile nat-nopat-address reserved 100
set system security-profile ls-design-profile logical-system ls-product-design
set system security-profile interconnect-profile logical-system
  interconnect-logical-system
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

Create three security profiles.

1. Create the first security profile.

   a. Specify the number of maximum and reserved policies.

      ```
      [edit system security-profile]
      user@host# set master-profile policy maximum 65 reserved 60
      ```

   b. Specify the number of maximum and reserved zones.

      ```
      [edit system security-profile]
      user@host# set master-profile zone maximum 22 reserved 17
      ```

   c. Specify the number of maximum and reserved sessions.

      ```
      [edit system security-profile]
      user@host# set master-profile flow-session maximum 3000 reserved 2100
      ```

   d. Specify the number of maximum and reserved source NAT no-PAT addresses and static NAT rules.

      ```
      [edit system security-profile]
      user@host# set master-profile nat-nopat-address maximum 115 reserved 100
      user@host# set master-profile nat-static-rule maximum 125 reserved 100
      ```

   e. Enable intrusion detection and prevention (IIDP). You can enable IDP only for the master (root) logical system.

      ```
      [edit system security-profile]
      user@host# set idp
      ```


Copyright © 2016, Juniper Networks, Inc.

f. Bind the security profile to the logical system.

```
[edit system security-profile]
user@host# set master-profile logical-system root-logical-system
```

2. Create the second security profile.

a. Specify the number of maximum and reserved policies.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile policy maximum 65 reserved 60
```

b. Specify the number of maximum and reserved zones.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile zone maximum 22 reserved 17
```

c. Specify the number of maximum and reserved sessions.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile flow-session maximum 2500 reserved
  2000
```

d. Specify the number of maximum and reserved source NAT no-PAT addresses.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile nat-nopat-address maximum 125 reserved
  100
```

e. Specify the number of maximum and reserved static NAT rules.

```
[edit system security-profile]
user@host# set ls-accnt-mrkt-profile nat-static-rule maximum 125 reserved 100
```

f. Bind the security profile to two logical systems.

```
[edit system]
user@host# set security-profile ls-accnt-mrkt-profile logical-system
  ls-marketing-dept
user@host# set security-profile ls-accnt-mrkt-profile logical-system
  ls-accounting-dept
```

3. Create the third security profile.

a. Specify the number of maximum and reserved policies.

```
[edit system security-profile]
user@host# set ls-design-profile policy maximum 50 reserved 40
```

b. Specify the number of maximum and reserved zones.

```
[edit system security-profile]
user@host# set ls-design-profile zone maximum 10 reserved 5
```

c. Specify the number of maximum and reserved sessions.

```
[edit system security-profile]
user@host# set ls-design-profile flow-session maximum 2500 reserved 2000
```

d. Specify the number of maximum and reserved source NAT no-PAT addresses.

```
[edit system security-profile]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ls-design-profile nat-nopat-address maximum 120 reserved 100
```

4. Bind the security profile to a logical system.

```
user@host# set system security-profile ls-design-profile logical-system
  ls-product-design
```

5. Bind a null security profile to the interconnect logical system.

```
user@host# set system security-profile interconnect-profile logical-system
  interconnect-logical-system
```

Results From configuration mode, confirm your configuration by entering the **show system security-profile** command to see all security profiles configured.

To see individual security profiles, enter the **show system security-profile master-profile**, the **show system security-profile ls-accnt-mrkt-profile** and, the **show system security-profile ls-design-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show system security-profile
  interconnect-profile {
    logical-system interconnect-logical-system;
  }
    ls-accnt-mrkt-profile {
      policy {
        maximum 65;
        reserved 60;
      }
      zone {
        maximum 22;
        reserved 17;
      }
      flow-session {
        maximum 2500;
        reserved 2000;
      }
      nat-nopat-address {
        maximum 125;
        reserved 100;
      }
      nat-static-rule {
        maximum 125;
        reserved 100;
      }
      logical-system [ ls-marketing-dept ls-accounting-dept ];
    }
    ls-design-profile {
      policy {
        maximum 50;
        reserved 40;
      }
      zone {
        maximum 10;
        reserved 5;
      }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    flow-session {
      maximum 2500;
      reserved 2000;
    }
    nat-nopat-address {
      maximum 120;
      reserved 100;
    }
    nat-static-rule {
      maximum 125;
      reserved 100;
    }
    logical-system ls-product-design;
  }
  master-profile {
    policy {
      maximum 65;
      reserved 60;
    }
    zone {
      maximum 22;
      reserved 17;
    }
    flow-session {
      maximum 3000;
      reserved 2100;
    }
    nat-nopat-address {
      maximum 115;
      reserved 100;
    }
    nat-static-rule {
      maximum 125;
      reserved 100;
    }
    root-logical-system;
  }
```

```
user@host# show system security-profile master-profile
    policy {
      maximum 65;
      reserved 60;
      }
      zone {
        maximum 22;
        reserved 17;
      }
      flow-session {
        maximum 3000;
        reserved 2100;
      }
      nat-nopat-address {
        maximum 115;
        reserved 100;
      }
      nat-static-rule {
```

Copyright © 2016, Juniper Networks, Inc.

```
            maximum 125;
            reserved 100;
          }
        root-logical-system;
```

```
user@host# show system security-profile ls-accnt-mrkt-profile
  policy {
    maximum 65;
    reserved 60;
    }
    zone {
      maximum 22;
      reserved 17;
    }
    flow-session {
      maximum 2500;
      reserved 2000;
    }
    nat-nopat-address {
      maximum 125;
      reserved 100;
    }
    nat-static-rule {
      maximum 125;
      reserved 100;
    }
  logical-system [ ls-accounting-dept ls-marketing-dept ];
```

```
user@host# show system security-profile ls-design-profile
  policy {
    maximum 50;
    reserved 40;
    }
    zone {
      maximum 10;
      reserved 5;
    }
    flow-session {
      maximum 2500;
      reserved 2000;
    }
    nat-nopat-address {
      maximum 120;
      reserved 100;
    }
    nat-static-rule {
      maximum 125;
      reserved 100;
    }
  logical-system ls-product-design;
```

If you are done configuring the device, enter commit from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the security resources that you allocated for logical systems have been assigned to them, follow this procedure for each logical system and for all its resources.

- Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems on page 3265

Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems

Purpose Verify security resources for each logical system. Follow this process for all configured logical systems.

Action

1. Use Telnet to log in to each user logical system as its user logical system administrator.

   Run Telnet, specifying the IP address of your SRX Series device. For example:

   telnet 10.11.11.19

2. Enter the login ID and password for one of the user logical systems that you created

   login: **lsmarketingadmin1**
   password: **Talk2345**
   lsmarketingadmin1@host:ls-marketing-dept>

3. Enter the following statement to Identify the resources configured for the profile.

   lsmarketingadmin1@host:ls-marketing-dept> **show system security-profile ?**

4. Enter the following command at the resulting prompt. Do this for each feature configured for the profile.

   lsmarketingadmin1@host:ls-marketing-dept> **show system security-profile zone detail**
   logical system name : ls-marketing-dept
   security profile name : ls-accnt-mrkt-profile
   used amount : 0
   reserved amount : 17
   maximum quota : 22

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 147

Configuring Master Logical System Security Features

- Understanding Logical System Firewall Authentication on page 3267
- Example: Configuring Access Profiles (Master Administrators Only) on page 3269
- Example: Configuring Security Features for the Master Logical System on page 3271
- IDP in Logical Systems Overview on page 3276
- Understanding IDP Features in Logical Systems on page 3278
- Example: Configuring an IDP Policy for the Master Logical System on page 3281
- Understanding Logical System Application Identification Services on page 3286
- Understanding Logical System Application Firewall Services on page 3287
- Example: Configuring Application Firewall Services for a Master Logical System on page 3288
- Understanding Logical System Application Tracking Services on page 3292
- Understanding Route-Based VPN Tunnels in Logical Systems on page 3293
- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) on page 3294

Understanding Logical System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

The master administrator is responsible for configuring access profiles in the master logical system. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored. Access profiles configured at the master logical system are available to all user logical systems.

The master administrator configures the maximum and reserved numbers of firewall authentications for each user logical system. The user logical system administrator can then create firewall authentications in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile auth-entry**

Copyright © 2016, Juniper Networks, Inc.

command to view the number of authentication resources allocated to the user logical system.

To configure the access profile, the master administrator uses the **profile** configuration statement at the [**edit access**] hierarchy level in the master logical system. The access profile can also include the order of authentication methods, LDAP or RADIUS server options, and session options.

The user logical system administrator can then associate the access profile with a security policy in the user logical system. The user logical system administrator also specifies the type of authentication:

- With pass-through authentication, a host or a user from one zone tries to access resources on another zone using an FTP, a Telnet, or an HTTP client. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.
- With Web authentication, users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

The user logical system administrator configures the following properties for firewall authentication in the user logical system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the **firewall-authentication** configuration statement at the [**edit security policies from-zone** *zone-name* **to-zone** *zone-name* **policy** *policy-name* **then permit**] hierarchy level.

  Users or user groups in an access profile who are allowed access by the policy can optionally be specified with the client-match configuration statement. (If no users or user groups are specified, any user who is successfully authenticated is allowed access.)

  For pass-through authentication, the access profile can optionally be specified and Web redirect (redirecting the client system to a webpage for authentication) can be enabled.
- Type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the **firewall-authentication** configuration statement at the [**edit access**] hierarchy level.
- Host inbound traffic. Protocols, services, or both are allowed to access the logical system. The types of traffic are configured with the **host-inbound-traffic** configuration statement at the [**edit security zones security-zone** *zone-name*] or [**edit security zones security-zone** *zone-name* **interfaces** *interface-name*] hierarchy levels.

From a user logical system, the user logical system administrator can use the **show security firewall-authentication users** or **show security firewall-authentication history** commands to view the information about firewall users and history for the user logical system. From the master logical system, the master administrator can use the same

 Copyright © 2016, Juniper Networks, Inc.

commands to view information for the master logical system, a specific user logical system, or all logical systems.

Related Documentation

- Example: Configuring Access Profiles (Master Administrators Only) on page 3269
- Example: Configuring Firewall Authentication for a User Logical System on page 3313
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Firewall User Authentication Overview on page 5179

Example: Configuring Access Profiles (Master Administrators Only)

The master administrator is responsible for configuring access profiles in the master logical system. This example shows how to configure access profiles.

- Requirements on page 3269
- Overview on page 3269
- Configuration on page 3270

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.
- Read "Firewall User Authentication Overview" on page 5179.

Overview

This example configures an access profile for LDAP authentication for logical system users. This example creates the access profile described in Table 336.

NOTE: The master administrator creates the access profile.

Table 336: Access Profile Configuration

Name	Configuration Parameters

ldap1	• LDAP is used as the first (and only) authentication method.
	• Base distinguished name:
	  ◦ Organizational unit name (OU): people
	  ◦ Domain components (DC): example, com
	• A user's LDAP distinguished name is assembled through the use of a common name identifier, username, and base distinguished name. The common name identifier is user ID (UID).
	• The LDAP server address is 10.155.26.104 and is reached through port 389.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

NOTE: You must be logged in as the master administrator.

```
set access profile ldap1 authentication-order ldap
set access profile ldap1 ldap-options base-distinguished-name
  ou=people,dc=example,dc=com
set access profile ldap1 ldap-options assemble common-name uid
set access profile ldap1 ldap-server 10.155.26.104 port 389
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an access profile in the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Configure an access profile and set the authentication order.

   ```
   [edit access profile ldap1]
   admin@host# set authentication-order ldap
   ```

3. Configure LDAP options.

   ```
   [edit access profile ldap1]
   admin@host# set ldap-options base-distinguished-name
     ou=people,dc=example,dc=com
   admin@host# set ldap-options assemble common-name uid
   ```

4. Configure the LDAP server.

   ```
   [edit access profile ldap1]
   admin@host# set ldap-server 10.155.26.104 port 389
   ```

Results

From configuration mode, confirm your configuration by entering the **show access profile *profile-name*** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
admin@host# show access profile ldap1
authentication-order ldap;
ldap-options {
  base-distinguished-name ou=people,dc=example,dc=com;
  assemble {
```


Copyright © 2016, Juniper Networks, Inc.

```
        common-name uid;
    }
}
    ldap-server {
        10.155.26.104 port 389;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- Example: Configuring Firewall Authentication for a User Logical System on page 3313
- Understanding Logical System Firewall Authentication on page 3267
- User Logical System Configuration Overview on page 3225

Example: Configuring Security Features for the Master Logical System

This example shows how to configure security features, such as zones, policies, and firewall authentication, for the master logical system.

- Requirements on page 3271
- Overview on page 3271
- Configuration on page 3272
- Verification on page 3276

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 3241.
- Use the **show system security-profile** command to see the resources allocated to the master logical system.
- Configure logical interfaces for the master logical system. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 3342.
- Configure the access profile ldap1 in the master logical system. The ldap1 access profile is used for Web authentication of firewall users. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 3269.

Overview

In this example, you configure security features for the master logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242. This example configures the security features described in Table 337.

Copyright © 2016, Juniper Networks, Inc.

Table 337: root-logical-system Security Feature Configuration

Feature	Name	Configuration Parameter
Zones	ls-root-trust	Bind to interface ge-0/0/4.0.
	ls-root-untrust	Bind to interface lt-0/0/0.1
Address books	root-internal	• Address masters: 12.12.1.0/24 • Attach to zone ls-root-trust
	root-external	• Address design: 12.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address marketing: 13.1.1.0/24 • Address set userlsys: design, accounting, marketing • Attach to zone ls-root-untrust
Security policies	permit-to-userlsys	Permit the following traffic: • From zone: ls-root-trust • To zone: ls-root-untrust • Source address: masters • Destination address: userlsys • Application: any
	permit-authorized-users	Permit the following traffic: • From zone: ls-root-untrust • To zone: ls-root-trust • Source address: userlsys • Destination address: masters • Application: junos-http, junos-https
Firewall authentication		• Web authentication • Authentication success banner “WEB AUTH LOGIN SUCCESS” • Default access profile ldap1
HTTP daemon		Activate on interface ge-0/0/4.0

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security address-book root-internal address masters 12.12.1.0/24
set security address-book root-internal attach zone ls-root-trust
set security address-book root-external address design 12.1.1.0/24
set security address-book root-external address accounting 14.1.1.0/24
set security address-book root-external address marketing 13.1.1.0/24
set security address-book root-external address-set userlsys address design
```


Copyright © 2016, Juniper Networks, Inc.

```
set security address-book root-external address-set userlsys address accounting
set security address-book root-external address-set userlsys address marketing
set security address-book root-external attach zone ls-root-untrust
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
  permit-to-userlsys match source-address masters
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
  permit-to-userlsys match destination-address userlsys
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
  permit-to-userlsys match application any
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
  permit-to-userlsys then permit
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
  permit-authorized-users match source-address userlsys
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
  permit-authorized-users match destination-address masters
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
  permit-authorized-users match application junos-http
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
  permit-authorized-users match application junos-https
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
  permit-authorized-users then permit firewall-authentication web-authentication
set security zones security-zone ls-root-trust interfaces ge-0/0/4.0
set security zones security-zone ls-root-untrust interfaces lt-0/0/0.1
set system services web-management http interface ge-0/0/4.0
set access firewall-authentication web-authentication default-profile ldap1
set access firewall-authentication web-authentication banner success "WEB AUTH
  LOGIN SUCCESS"
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure zones and policies for the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Create security zones and assign interfaces to each zone.

   ```
   [edit security zones]
   admin@host# set security-zone ls-root-trust interfaces ge-0/0/4.0
   admin@host# set security-zone ls-root-untrust interfaces lt-0/0/0.1
   ```

3. Create address book entries.

   ```
   [edit security]
   admin@host# set address-book root-internal address masters 12.12.1.0/24
   admin@host# set address-book root-external address design 12.1.1.0/24
   admin@host# set address-book root-external address accounting 14.1.1.0/24
   admin@host# set address-book root-external address marketing 13.1.1.0/24
   admin@host# set address-book root-external address-set userlsys address design
   admin@host# set address-book root-external address-set userlsys address
     accounting
   ```

Copyright © 2016, Juniper Networks, Inc.

```
admin@host# set address-book root-external address-set userlsys address
  marketing
```

4. Attach address books to zones.

```
[edit security]
admin@host# set address-book root-internal attach zone ls-root-trust
admin@host# set address-book root-external attach zone ls-root-untrust
```

5. Configure a security policy that permits traffic from the ls-root-trust zone to the ls-root-untrust zone.

```
[edit security policies from-zone ls-root-trust to-zone ls-root-untrust]
admin@host# set policy permit-to-userlsys match source-address masters
admin@host# set policy permit-to-userlsys match destination-address userlsys
admin@host# set policy permit-to-userlsys match application any
admin@host# set policy permit-to-userlsys then permit
```

6. Configure a security policy that authenticates traffic from the ls-root-untrust zone to the ls-root-trust zone.

```
[edit security policies from-zone ls-root-untrust to-zone ls-root-trust]
admin@host# set policy permit-authorized-users match source-address userlsys
admin@host# set policy permit-authorized-users match destination-address masters
admin@host# set policy permit-authorized-users match application junos-http
admin@host# set policy permit-authorized-users match application junos-https
admin@host# set policy permit-authorized-users then permit firewall-authentication
  web-authentication
```

7. Configure the Web authentication access profile and define a success banner.

```
[edit access]
admin@host# set firewall-authentication web-authentication default-profile ldap1
admin@host# set firewall-authentication web-authentication banner success “WEB
  AUTH LOGIN SUCCESS”
```

8. Activate the HTTP daemon on the device.

```
[edit system]
admin@host# set services web-management http interface ge-0/0/4.0
```

Results From configuration mode, confirm your configuration by entering the **show security**, **show access**, and **show system services** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security
...
address-book {
  root-internal {
    address masters 12.12.1.0/24;
    attach {
      zone ls-root-trust;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    root-external {
      address design 12.1.1.0/24;
      address accounting 14.1.1.0/24;
      address marketing 13.1.1.0/24;
      address-set userlsys {
        address design;
        address accounting;
        address marketing;
      }
      attach {
        zone ls-root-untrust;
      }
    }
  }
    policies {
      from-zone ls-root-trust to-zone ls-root-untrust {
        policy permit-to-userlsys {
          match {
            source-address masters;
            destination-address userlsys;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone ls-root-untrust to-zone ls-root-trust {
        policy permit-authorized-users {
          match {
            source-address userlsys;
            destination-address masters;
            application [ junos-http junos-https ];
          }
          then {
            permit {
              firewall-authentication {
                web-authentication;
              }
            }
          }
        }
      }
    }
    zones {
      security-zone ls-root-trust {
        interfaces {
          ge-0/0/4.0;
        }
      }
      security-zone ls-root-untrust {
        interfaces {
          lt-0/0/0.1;
        }
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
[edit]
admin@host# show access
...
firewall-authentication {
  web-authentication {
    default-profile ldap1;
    banner {
      success "WEB AUTH LOGIN SUCCESS";
    }
  }
}
[edit]
admin@host# show system services
web-management {
  http {
    interface ge-0/0/4.0;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Policy Configuration on page 3276

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation

- Understanding Logical System Zones on page 3299
- Understanding Logical System Security Policies on page 3306
- Understanding Logical System Firewall Authentication on page 3267

IDP in Logical Systems Overview

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a logical system.

This topic includes the following sections:

- IDP Policies on page 3277
- IDP Installation and Licensing for Logical Systems on page 3277

Copyright © 2016, Juniper Networks, Inc.

IDP Policies

The master administrator configures IDP policies at the root level. Configuring an IDP policy for logical systems is similar to configuring an IDP policy on a device that is not configured for logical systems. This can include the configuration of custom attack objects.

NOTE: User logical system administrators cannot create or modify IDP policies for their user logical systems. Only the master administrator can create IDP policies and bind them to user logical systems through a logical systems security profile.

NOTE: The user logical system administrator can create security zones in the user logical system and assign interfaces to each security zone. Zones that are specific to user logical systems cannot be referenced in IDP policies configured by the master administrator. The master administrator can reference zones in the master logical system in an IDP policy configured for the master logical system.

The master administrator then specifies an IDP policy in the security profile that is bound to a logical system. To enable IDP in a logical system, the master administrator or user logical system administrator configures a security policy that defines the traffic to be inspected and specifies the **permit application-services idp** action.

Although the master administrator can configure multiple IDP policies, a logical system can have only one active IDP policy at a time. For user logical systems, the master administrator can either bind the same IDP policy to multiple user logical systems or bind a unique IDP policy to each user logical system. To specify the active IDP policy for the master logical system, the master administrator can *either* reference the IDP policy in the security profile that is bound to the master logical system or use the **active-policy** configuration statement at the [**edit security idp**] hierarchy level.

NOTE: A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the active-policy configuration statement. Use only one method to specify the active IDP policy for the master logical system.

IDP Installation and Licensing for Logical Systems

A single IDP security package is installed for all logical systems on the device. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all logical systems.

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any logical system on the device.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding IDP Features in Logical Systems on page 3278
- Example: Configuring an IDP Policy for a User Logical System on page 3321
- Example: Configuring an IDP Policy for the Master Logical System on page 3281
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- *IDP Policies Overview*

Understanding IDP Features in Logical Systems

This topic includes the following sections:

- Rulebases on page 3278
- Protocol Decoders on page 3278
- SSL Inspection on page 3279
- Inline Tap Mode on page 3279
- Multi-Detectors on page 3279
- Logging and Monitoring on page 3279

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The following IDP rulebases are supported for logical systems:

- The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.
- The application-level distributed denial-of-service (DDoS) rulebase defines parameters to protect servers such as DNS or HTTP. The application-level DDoS rulebase defines the source match condition for traffic that should be monitored and takes an action, such as drop the connection, drop the packet, or no action. It can also perform actions against future connections that use the same IP address.

NOTE: **Status monitoring for IPS and application-level DDoS is global to the device and not on a per logical system basis.**

Protocol Decoders

The Junos IDP module ships with a set of preconfigured protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks that they perform. The IDP protocol decoder configuration is global and applies to all logical systems. Only the master administrator at the root level can modify the settings at the [**edit security idp sensor-configuration**] hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

SSL Inspection

IDP SSL inspection uses the Secure Sockets Layer (SSL) protocol suite to enable inspection of HTTP traffic encrypted in SSL.

SSL inspection configuration is global and applies to all logical systems on a device. SSL inspection can only be configured by the master administrator at the root level with the **ssl-inspection** configuration statement at the [**edit security idp sensor-configuration**] hierarchy level.

Inline Tap Mode

The inline tap mode feature provides passive, inline detection of Application Layer threats for traffic matching security policies that have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module. This allows the packets to get to the next service module without waiting for IDP processing results.

Inline tap mode is enabled or disabled for all logical systems at the root level by the master administrator. To enable inline tap mode, use the **inline-tap** configuration statement at the [**edit security forwarding-process application-services maximize-idp-sessions**] hierarchy level. Delete the inline tap mode configuration to switch the device back to regular mode.

NOTE: **The device must be restarted when switching to inline tap mode or back to regular mode.**

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all logical systems.

Logging and Monitoring

Status monitoring options are available to the master administrator only. All status monitoring options under the **show security idp** and **clear security idp** CLI operational commands present global information, but not on a per logical system basis.

NOTE: **SNMP monitoring for IDP is not supported on logical systems.**

Copyright © 2016, Juniper Networks, Inc.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The logical systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the ls-product-design logical system:

  Oct 12 17:33:32 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In ls-product-design at 1286930013, SIG Attack log <4.0.0.1/34327->5.0.0.1/21> for TCP protocol and service SERVICE_IDP application NONE by rule 1 of rulebase IPS in policy Recommended. attack: repeat=0, action=IGNORE, threat-severity=MEDIUM, name=FTP:USER:ROOT, NAT <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:ls-product-design-untrust:ge-0/0/0.0->ls-product-design-trust:ge-0/0/1.0, packet-log-id: 65535 and misc-message -

- IP action logs. The following example shows an IP action log for the ls-product-design logical system:

  Oct 13 16:56:04 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In ls-product-design at 1287014163, TRAFFIC Attack log <25.0.0.1/34802->15.0.0.1/21> for TCP protocol and service SERVICE_NONE application NONE by rule 1 of rulebase IPS in policy Recommended. attack: repeat=0, action=TRAFFIC_IPACTION_NOTIFY, threat-severity=INFO, name=_, NAT <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:ls-product-design-trust:ge-0/0/1.0->ls-product-design-untrust:plt0.3, packet-log-id: 0 and misc-message -

- Application DDoS logs. The following example shows an application DDoS log for the ls-product-design logical system:

  Oct 11 16:29:57 8.0.0.254 RT_IDP: IDP_APPDDOS_APP_ATTACK_EVENT_LS: DDOS Attack in ls-product-design at 1286839797 on my-http, <ls-product-design-untrust:ge-0/0/0.0:4.0.0.1:33738->ls-product-design-trust:ge-0/0/1.0:5.0.0.1:80> for TCP protocol and service HTTP by rule 1 of rulebase DDOS in policy Recommended. attack: repeats 0 action DROP threat-severity INFO, connection-hit-rate 0, context-name http-url-parsed, hit-rate 6, value-hit-rate 6 time-scope PEER time-count 2 time-period 10 secs, context value: ascii: /abc.html hex: 2f 61 62 63 2e 68 74 6d 6c

Related Documentation

- *Understanding IDP Policy Rule Bases*
- *Understanding IDP Protocol Decoders*
- *IDP SSL Overview*
- *Understanding IDP Inline Tap Mode*
- *Understanding Multiple IDP Detector Support*
- *Understanding IDP Logging*

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring an IDP Policy for the Master Logical System

This example shows how to configure an IDP policy in a master logical system.

- Requirements on page 3281
- Overview on page 3281
- Configuration on page 3282
- Verification on page 3286

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.
- Read "IDP in Logical Systems Overview" on page 3276.
- Use the **show system security-profile** command to see the resources allocated to the master logical system.

Overview

In this example you configure a custom attack that is used in an IDP policy. The IDP policy is specified in a security profile that is applied to the master logical system. IDP is then enabled in a security policy configured in the master logical system.

You configure the features described in Table 338.

Table 338: IDP Configuration for the Master Logical System

Feature	Name	Configuration Parameters

Custom attack	http-bf	• Severity critical
		• Detect three attacks between source and destination addresses of sessions.
		• Stateful signature attack type with the following characteristics:
		  ◦ location http-url-parsed
		  ◦ pattern .*juniper.*
		  ◦ client to server traffic

IPS rulebase policy	root-idp-policy	Match: • application default • http-bf custom attacks Action: • drop-connection • notification log-attacks

Copyright © 2016, Juniper Networks, Inc.

Table 338: IDP Configuration for the Master Logical System *(continued)*

Feature	Name	Configuration Parameters
Logical system security profile	master-profile (previously configured and applied to root-logical-system)	Add IDP policy root-idp-policy.
Security policy	enable-idp	Enable IDP in a security policy that matches any traffic from the lsys-root-untrust zone to the lsys-root-trust zone.

NOTE: A logical system can have only one active IDP policy at a time. To specify the active IDP policy for the master logical system, the master administrator can reference the IDP policy in the security profile that is bound to the master logical system as shown in this example. Alternatively, the master administrator can use the **active-policy** configuration statement at the **[edit security idp]** hierarchy level.

A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the **active-policy** configuration statement. Use only one method to specify the active IDP policy for the master logical system.

Configuration

- Configuring a Custom Attack on page 3282
- Configuring an IDP Policy for the Master Logical System on page 3284
- Enabling IDP in a Security Policy on page 3285

Configuring a Custom Attack

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security idp custom-attack http-bf severity critical
set security idp custom-attack http-bf time-binding count 3
set security idp custom-attack http-bf time-binding scope peer
set security idp custom-attack http-bf attack-type signature context http-url-parsed
set security idp custom-attack http-bf attack-type signature pattern .*juniper.*
set security idp custom-attack http-bf attack-type signature direction client-to-server
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a custom attack object:

1. Log in to the master logical system as the master administrator and enter configuration mode.

```
[edit]
admin@host> configure
admin@host#
```

2. Create the custom attack object and set the severity level.

```
[edit security idp]
admin@host# set custom-attack http-bf severity critical
```

3. Configure attack detection parameters.

```
[edit security idp]
admin@host# set custom-attack http-bf time-binding count 3
admin@host# set custom-attack http-bf time-binding scope peer
```

4. Configure stateful signature parameters.

```
[edit security idp]
admin@host# set custom-attack http-bf attack-type signature context
  http-url-parsed
admin@host# set custom-attack http-bf attack-type signature pattern .*juniper.*
admin@host# set custom-attack http-bf attack-type signature direction
  client-to-server
```

Results

From configuration mode, confirm your configuration by entering the **show security idp custom-attack http-bf** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp custom-attack http-bf
severity critical;
  time-binding {
    count 3;
    scope peer;
  }
    attack-type {
      signature {
        context http-url-parsed;
        pattern .*juniper.*;
        direction client-to-server;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring an IDP Policy for the Master Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match application default
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match attacks custom-attacks
  http-bf
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then action drop-connection
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then notification log-attacks
set system security-profile master-profile idp-policy lsys1-idp-policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IDP policy:

1. Create the IDP policy and configure match conditions.

```
[edit security idp]
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match application
  default
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match attacks
  custom-attacks http-bf
```

2. Configure actions for the IDP policy.

```
[edit security idp]
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then action
  drop-connection
admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then notification
  log-attacks
```

3. Add the IDP policy to the security profile.

```
[edit system security-profile master-profile]
admin@host# set idp-policy lsys1-idp-policy
```

Results

From configuration mode, confirm your configuration by entering the **show security idp idp-policy root-idp-policy** and **show system security-profile master-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp idp-policy root-idp-policy
  rulebase-ips {
    rule 1 {
      match {
        application default;
        attacks {
          custom-attacks http-bf;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          action {
            drop-connection;
          }
          notification {
            log-attacks;
          }
        }
      }
    }
admin@host# show system security-profile master-profile
...
idp-policy lsys1-idp-policy;
```

If you are done configuring the device, enter **commit** from configuration mode.

Enabling IDP in a Security Policy

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
  match source-address any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
  match destination-address any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
  match application any
set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp
  then permit application-services idp
```

Step-by-Step Procedure

To enable IDP in a security policy:

1. Create the security policy and configure match conditions.

   ```
   [edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
   admin@host# set policy enable-idp match source-address any
   admin@host# set policy enable-idp match destination-address any
   admin@host# set policy enable-idp match application any
   ```

2. Enable IDP.

   ```
   [edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
   admin@host# set policy enable-idp then permit application-services idp
   ```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
  admin@host# show security policies
  from-zone lsys-root-untrust to-zone lsys-root-trust {
    policy enable-idp {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          application-services {
            idp;
          }
        }
      }
    }
  }
...
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Attack Matches

Purpose Verify that attacks are being matched in network traffic.

Action From operational mode, enter the **show security idp attack table** command.

```
admin@host> show security idp attack table
  IDP attack statistics:
   Attack name                    #Hits
   http-bf                        1
```

Related Documentation

- IDP in Logical Systems Overview on page 3276
- SRX Series Logical System Master Administrator Configuration Tasks Overview on page 3222

Understanding Logical System Application Identification Services

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following benefits:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.
- Improves performance by narrowing the scope of attack signatures for applications without decoders.

Copyright © 2016, Juniper Networks, Inc.

- Enables you to create detailed reports using AppTrack on applications passing through the device.

With logical systems, predefined and custom application signatures are global resources that are shared by all logical systems. The master administrator is responsible for downloading and installing predefined Juniper Networks application signatures and creating custom application and nested application signatures to identify applications that are not part of the predefined database.

Application identification is enabled by default.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each user logical system has its own ASC. A user logical system administrator can display the ASC entries for their logical system with the **show services application-identification application-system-cache** command. A user logical system administrator can use the **clear services application-identification application-system-cache** command to clear the ASC entries for their logical system.

The master administrator can display or clear ASC entries for any logical system. The master administrator can also display or clear global counters with the **show services application-identification counter** and **clear services application-identification counter** commands.

Related Documentation

- Understanding the Junos OS Application Identification Database on page 488
- Example: Scheduling the Application Signature Package Updates on page 499
- Example: Configuring Junos OS Application Identification Custom Application Signatures on page 512
- *Understanding IDP Application Identification*
- Understanding the Application System Cache on page 523
- Verifying Application System Cache Statistics on page 524

Understanding Logical System Application Firewall Services

An application firewall enables administrators of logical systems to create security policies for traffic based on application identification defined by application signatures. The application firewall provides additional security protection against dynamic-application traffic that might not be adequately controlled by standard network firewall policies. The application firewall controls information transmission by allowing or blocking traffic originating from particular applications.

To configure an application firewall, you define a rule set that contains rules specifying the action to be taken on identified dynamic applications. The rule set is configured independently and assigned to a security policy. Each rule set contains at least two rules, a matched rule (consisting of match criteria and action) and a default rule.

Copyright © 2016, Juniper Networks, Inc.

- A matched rule defines the action to be taken on matching traffic. When traffic matches an application and other criteria specified in the rule, the traffic is allowed or blocked based on the action specified in the rule.
- A default rule is applied when traffic does not match any other rule in the rule set.

The master administrator can download a predefined application signature database from the Juniper Networks Security Engineering website or can define application signatures using the Junos OS configuration CLI. For more information about application identification and application signatures, see *AppSecure Services Feature Guide for Security Devices*.

Configuring an application firewall on a logical system is the same process as configuring an application firewall on a device that is not configured with logical systems. However, the application firewall applies only to the logical system for which it is configured. The master administrator can configure, enable, and monitor application firewalls on the master logical system and all user logical systems on a device. User logical system administrators can configure, enable, and monitor application firewalls only on the user logical systems for which they have access.

Related Documentation

- Example: Configuring Application Firewall Services for a Master Logical System on page 3288
- Example: Configuring Application Firewall Services for a User Logical System on page 3327

Example: Configuring Application Firewall Services for a Master Logical System

This example describes how to configure application firewall services on the master, or root, logical system by a master administrator. Only the master administrator can configure, manage, and view configuration of the master logical system, in addition to all user logical systems.

After configuring application firewall rule sets and rules, the master administrator adds the application firewall rule set information to the security policy on the master logical system.

For information about configuring an application firewall within a security policy, see "Application Firewall Overview" on page 555.

- Requirements on page 3288
- Overview on page 3289
- Configuration on page 3289
- Verification on page 3291

Requirements

Before you begin:

- Verify that all interfaces, routing instances, and security zones have been configured on the master logical system.

Copyright © 2016, Juniper Networks, Inc.

See "Example: Configuring Security Features for the Master Logical System" on page 3271.

- Verify that application firewall resources (appfw-rule-set and appfw-rule) have been allocated in a security profile and bound to the master logical system through the [**system security-profile**] command. For application firewall resources, a security profile configuration allows 0 to 10,000 rule sets and 0 to 10,000 rules.

  NOTE: The master administrator allocates various global system resources through a security profile configuration which is then bound to the various logical systems on the device. The master administrator owns this function and configures the security profile for all user logical systems as well as the master logical system.

  For more information, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 3253.

- Log in to the master logical system as the master administrator.

  For information about master administrator role functions, see "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Overview

In this example you create application firewall services on the master logical system, called root-logical-system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example creates the following application firewall configuration:

- Rule set, root-rs1, with rules r1 and r2. When r1 is matched, Telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.
- Rule set, root-rs2, with rule r1. When r1 is matched, Facebook traffic is blocked by the firewall.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r1 match dynamic-application junos:TELNET
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r1 then permit
```

Copyright © 2016, Juniper Networks, Inc.

```
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
  rule r2 match dynamic-application-group junos:web
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
  rule r2 then permit
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
  default-rule deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
  rule r1 match dynamic-application junos:facebook
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
  rule r1 then deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
  default-rule permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure application firewall for a master logical system:

1. Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 3241 and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Configure an application firewall rule set for root-logical-system.

   ```
   [edit ]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
   ```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

   ```
   [edit]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
     rule r1 match dynamic-application telnet then permit
   ```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

   ```
   [edit]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
     default-rule deny
   ```

5. Repeat these steps to configure another rule set, root-rs2, if desired.

Results

From configuration mode, confirm your configuration by entering the **show security application-firewall rule-sets** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security application-firewall rule-sets all
```


Copyright © 2016, Juniper Networks, Inc.

```
...
  application-firewall {
    rule-sets root-rs1 {
      rule r1 {
        match {
          dynamic-application [junos:TELNET];
        }
        then {
          permit;
        }
      }
      default-rule {
        deny;
      }
    }
    rule-sets root-rs1 {
      rule r2 {
        match {
          dynamic-application-group [junos:web];
        }
        then {
          permit;
        }
      }
      rule-sets root-rs2 {
        rule r1 {
          match {
            dynamic-application [junos:FACEBOOK];
          }
          then {
            deny;
          }
        }
        default-rule {
          permit;
        }
      }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Application Firewall Configuration on page 3291

Verifying Application Firewall Configuration

Purpose View the application firewall configuration on the master logical system.

Action From operational mode, enter the **show security application-firewall rule-set logical-system root-logical-system rule-set all** command.

```
admin@host> show security application-firewall rule-set logical-system root-logical-system
rule-set all
```

```
Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Related Documentation

- SRX Series Logical System Master Administrator Configuration Tasks Overview on page 3222
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Understanding Logical System Application Firewall Services on page 3287
- Example: Configuring Security Features for the Master Logical System on page 3271

Understanding Logical System Application Tracking Services

AppTrack is an application tracking tool that provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack can be enabled and configured within any logical system. Configuring AppTrack in a logical system is the same as configuring AppTrack on a device that is not configured for logical systems. An AppTrack configuration only applies to the logical system in which it is configured. The name of the logical system is added to AppTrack logs. The master administrator can configure AppTrack for any logical system while a user logical system

Copyright © 2016, Juniper Networks, Inc.

administrator can only configure AppTrack for the logical system that they are logged in to.

NOTE: The system log configuration is global on the device and must be configured by the master administrator. The user logical system administrator cannot configure system logging for a logical system.

Counters keep track of the number of log messages sent and logs that have failed. AppTrack counters are global to the device. The master administrator as well as user logical system administrators can view AppTrack counters with the **show security application-tracking counters** command.

Related Documentation

- Understanding AppTrack on page 573
- Example: Configuring AppTrack on page 574
- Example: Configuring AppTrack for a User Logical System on page 3331

Understanding Route-Based VPN Tunnels in Logical Systems

A VPN connection can secure traffic that passes between a logical system and a remote site across a WAN. With route-based VPNs, you configure one or more security policies in a logical system to regulate the traffic flowing through a single IP Security (IPsec) tunnel. For each IPsec tunnel, there is one set of IKE and IPsec security associations (SAs) that must be configured at the root level by the master administrator.

NOTE: The external interface configured under the gateway configuration can only be a part of the root logical system.

NOTE: Only route-based VPNs are supported for logical systems. Policy-based VPNs are not supported.

In addition to configuring IKE and IPsec SAs for each VPN, the master administrator must also assign a secure tunnel (st0) interface to a user logical system. An st0 interface can only be assigned to a single user logical system. However, multiple user logical systems can each be assigned their own st0 interface.

NOTE: The st0 unit 0 interface should not be assigned to a logical system, as an SA cannot be set up for this interface.

The user logical system administrator can configure the IP address and other attributes of the st0 interface assigned to the user logical system. The user logical system administrator cannot delete an st0 interface assigned to their user logical system.

Copyright © 2016, Juniper Networks, Inc.

For route-based VPNs, a security policy refers to a destination address and not a specific VPN tunnel. For cleartext traffic in a user logical system to be sent to the VPN tunnel for encapsulation, the user logical system administrator must make the following configurations:

- Security policy that permits traffic to a specified destination.
- Static route to the destination with the st0 interface as the next hop.

When Junos OS looks up routes in the user logical system to find the interface to use to send traffic to the destination address, it finds a static route through the st0 interface. Traffic is routed to the VPN tunnel as long as the security policy action is permit.

The master logical system and a user logical system can share a route-based VPN tunnel. An st0 interface assigned to a user logical system can also be used by the master logical system. For the master logical system, the master administrator configures a security policy that permits traffic to the remote destination and a static route to the remote destination with the st0 interface as the next hop.

VPN monitoring is configured by the master administrator in the master logical system. For the VPN monitor source interface, the master administrator must specify the st0 interface; a physical interface for a user logical system cannot be specified.

Related Documentation

- Understanding Route-Based IPsec VPNs on page 6073
- User Logical System Configuration Overview on page 3225
- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) on page 3294
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)

The master administrator is responsible for assigning an st0 interface to a user logical system and configuring IKE and IPsec SAs at the root level for each VPN tunnel. This example shows how to assign an st0 interface to a user logical system and configure IKE and IPsec SA parameters.

- Requirements on page 3294
- Overview on page 3295
- Configuration on page 3295
- Verification on page 3298

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "“Understanding the Master Logical System and the Master Administrator Role” on page 3221.

Copyright © 2016, Juniper Networks, Inc.

- Read "Understanding Route-Based IPsec VPNs" on page 6073.

Overview

In this example you configure a VPN tunnel for the ls-product-design user logical system. This example configures the VPN tunnel parameters described in Table 339.

Table 339: Logical System VPN Tunnel Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	Assigned to ls-product-design logical system
IKE proposal	ike-phase1-proposal	• Preshared keys authentication • Diffie-Hellman group 2 • sha1 authentication algorithm • aes-128-cbc encryption algorithm
IKE policy		• Main mode • References IKE proposal ike-phase1-proposal • ASCII preshared key 395psksecr3t
IKE gateway	ike-gw	• External interface ge-0/0/3.0 • References IKE policy ike-phase1-policy • Address 2.2.2.2
IPsec proposal	ipsec-phase2-proposal	• ESP protocol • hmac-sha1-96 authentication algorithm • aes-128-cbc encryption algorithm
IPsec policy	vpn-policy1	• References ipsec-phase2-proposal • perfect-forward-secrecy keys group2
VPN	ike-vpn	• bind-interface st0.1 • References ike-gw gateway • References vpn-policy1 policy
VPN monitoring		For ike-vpn VPN: • source-interface st0.1 • destination-ip 4.0.0.1

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems ls-product-design interfaces st0 unit 1
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike-gw ike-policy ike-phase1-policy
set security ike gateway ike-gw address 2.2.2.2
set security ike gateway ike-gw external-interface ge-0/0/3.0
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group2
set security ipsec policy vpn-policy1 proposals ipsec-phase2-proposal
set security ipsec vpn ike-vpn bind-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor source-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
set security ipsec vpn ike-vpn ike gateway ike-gw
set security ipsec vpn ike-vpn ike ipsec-policy vpn-policy1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To assign a VPN tunnel interface to a user logical system and configure IKE and IPsec SAs:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   [edit]
   admin@host> configure
   admin@host#
   ```

2. Assign a VPN tunnel interface.

   ```
   [edit logical-systems ls-product-design]
   admin@host# set interfaces st0 unit 1
   ```

3. Configure an IKE proposal.

   ```
   [edit security ike]
   admin@host# set proposal ike-phase1-proposal authentication-method
     pre-shared-keys
   admin@host# set proposal ike-phase1-proposal dh-group group2
   admin@host# set proposal ike-phase1-proposal authentication-algorithm sha1
   admin@host# set proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
   ```

4. Configure an IKE policy.

   ```
   [edit security ike]
   admin@host# set policy ike-phase1-policy mode main
   admin@host# set policy ike-phase1-policy proposals ike-phase1-proposal
   admin@host# set policy ike-phase1-policy pre-shared-key ascii-text 395psksecr3t
   ```

5. Configure an IKE gateway.

   ```
   [edit security ike]
   admin@host# set gateway ike-gw external-interface ge-0/0/3.0
   ```


Copyright © 2016, Juniper Networks, Inc.

```
admin@host# set gateway ike-gw ike-policy ike-phase1-policy
admin@host# set gateway ike-gw address 2.2.2.2
```

6. Configure an IPsec proposal.

```
[edit security ipsec]
admin@host# set proposal ipsec-phase2-proposal protocol esp
admin@host# set proposal ipsec-phase2-proposal authentication-algorithm
  hmac-sha1-96
admin@host# set proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
```

7. Configure an IPsec policy.

```
[edit security ipsec]
admin@host# set policy vpn-policy1 proposals ipsec-phase2-proposal
admin@host# set policy vpn-policy1 perfect-forward-secrecy keys group2
```

8. Configure the VPN.

```
[edit security ipsec]
admin@host# set vpn ike-vpn bind-interface st0.1
admin@host# set vpn ike-vpn ike gateway ike-gw
admin@host# set vpn ike-vpn ike ipsec-policy vpn-policy1
```

9. Configure VPN monitoring.

```
[edit security ipsec]
admin@host# set vpn ike-vpn vpn-monitor source-interface st0.1
admin@host# set vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, and **show security ipsec** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
admin@host# show interfaces
  st0 {
    unit 1;
  }
[edit]
admin@host# show security ike
  proposal ike-phase1-proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
  }
    policy ike-phase1-policy {
      mode main;
      proposals ike-phase1-proposal;
      pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
    }
    gateway ike-gw {
      ike-policy ike-phase1-policy;
      address 2.2.2.2;
      external-interface ge-0/0/3.0;
    }
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
admin@host# show security ipsec
  proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
  }
    policy vpn-policy1 {
      perfect-forward-secrecy {
        keys group2;
      }
      proposals ipsec-phase2-proposal;
    }
    vpn ike-vpn {
      bind-interface st0.1;
      vpn-monitor {
        source-interface st0.1;
        destination-ip 4.0.0.1;
      }
      ike {
        gateway ike-gw;
        ipsec-policy vpn-policy1;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose Verify that the IKE and IPsec SA configuration is correct.

Action From operational mode, enter the **show security ike** and **show security ipsec** commands.

Related Documentation

- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335
- Understanding Route-Based VPN Tunnels in Logical Systems on page 3293
- User Logical System Configuration Overview on page 3225

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 148

Configuring User Logical System Security Features

- Understanding Logical System Zones on page 3299
- Example: Configuring Zones for a User Logical System on page 3301
- Understanding Logical System Screen Options on page 3304
- Example: Configuring Screen Options for a User Logical System on page 3304
- Understanding Logical System Security Policies on page 3306
- Example: Configuring Security Policies in a User Logical System on page 3308
- Understanding Logical System Firewall Authentication on page 3311
- Example: Configuring Firewall Authentication for a User Logical System on page 3313
- IDP in Logical Systems Overview on page 3317
- Understanding IDP Features in Logical Systems on page 3318
- Example: Configuring an IDP Policy for a User Logical System on page 3321
- Example: Enabling IDP in a User Logical System Security Policy on page 3323
- Understanding Logical System Application Identification Services on page 3326
- Example: Configuring Application Firewall Services for a User Logical System on page 3327
- Understanding Logical System Application Tracking Services on page 3330
- Example: Configuring AppTrack for a User Logical System on page 3331
- Understanding Route-Based VPN Tunnels in Logical Systems on page 3334
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

Understanding Logical System Zones

Security zones are logical entities to which one or more interfaces are bound. Security zones can be configured on the master logical system by the master administrator or on user logical systems by the user logical system administrator. On a logical system, the administrator can configure multiple security zones, dividing the network into network segments to which various security options can be applied.

The master administrator configures the maximum and reserved numbers of security zones for each user logical system. The user logical system administrator can then create

Copyright © 2016, Juniper Networks, Inc.

security zones in the user logical system and assign interfaces to each security zone. From a user logical system, the user logical system administrator can use the **show system security-profile zones** command to view the number of security zones allocated to the user logical system and the **show interfaces** command to view the interfaces allocated to the user logical system.

NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of security zones applied to the master logical system. The number of zones configured in the master logical system count toward the maximum number of zones available on the device.

The master and user administrator can configure the following properties of a security zone in a logical system:

- Interfaces that are part of a security zone.
- Screen options—For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.
- TCP-Reset—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the synchronize flag set.
- Host inbound traffic—This feature specifies the kinds of traffic that can reach the device from systems that are directly connected to its interfaces. You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)

There are no preconfigured security zones in the master logical system or user logical system.

The management functional zone (MGT) can only be configured for the master logical system. There is only one management interface per device and that interface is allocated to the master logical system.

The **all** interface can only be assigned to a zone in the master logical system by the master administrator.

The user logical system administrator can configure and view all attributes for a security zone in a user logical system. All attributes of a security zone in a user logical system are also visible to the master administrator.

Related Documentation

- Example: Configuring Zones for a User Logical System on page 3301
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Understanding Logical System Interfaces and Routing Instances on page 3341

Copyright © 2016, Juniper Networks, Inc.

- Security Zones and Interfaces Overview on page 1049

Example: Configuring Zones for a User Logical System

This example shows how to configure zones for a user logical system.

- Requirements on page 3301
- Overview on page 3301
- Configuration on page 3302

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profile zones** command to see the zone resources allocated to the logical system.
- Logical interfaces for the user logical system must be configured. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 3360.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example creates the zones and address books described in Table 340.

Table 340: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 12.1.1.0/24 • Attach to zone ls-product-design-trust
	product-design-external	• Address marketing: 13.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security address-book product-design-internal address product-designers 12.1.1.0/24
set security address-book product-design-internal attach zone ls-product-design-trust
set security address-book product-design-external address marketing 13.1.1.0/24
set security address-book product-design-external address accounting 14.1.1.0/24
set security address-book product-design-external address others 12.12.1.0/24
set security address-book product-design-external address-set otherlsys address
  marketing
set security address-book product-design-external address-set otherlsys address
  accounting
set security address-book product-design-external attach zone ls-product-design-untrust
set security zones security-zone ls-product-design-trust tcp-rst
set security zones security-zone ls-product-design-trust interfaces ge-0/0/5.1
set security zones security-zone ls-product-design-untrust interfaces lt-0/0/0.3
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security zone and assign it to an interface.

   ```
   [edit security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust
     interfaces ge-0/0/5.1
   ```

3. Configure the TCP-Reset parameter for the zone.

   ```
   [edit security zones security-zone ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set tcp-rst
   ```

4. Configure a security zone and assign it to an interface.

   ```
   [edit security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust
     interfaces lt-0/0/0.3
   ```

5. Create global address book entries.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
     address product-designers 12.1.1.0/24
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address marketing 13.1.1.0/24
   ```


Copyright © 2016, Juniper Networks, Inc.

```
lsdesignadmin1@host:ls-product-design# set address-book product-design-external
  address accounting 14.1.1.0/24
lsdesignadmin1@host:ls-product-design# set address-book product-design-external
  address others 12.12.1.0/24
lsdesignadmin1@host:ls-product-design# set address-book product-design-external
  address-set otherlsys address marketing
lsdesignadmin1@host:ls-product-design# set address-book product-design-external
  address-set otherlsys address accounting
```

6. Attach address books to zones.

```
[edit security]
lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
  attach zone ls-product-design-trust
lsdesignadmin1@host:ls-product-design# set address-book product-design-external
  attach zone ls-product-design-untrust
```

Results From configuration mode, confirm your configuration by entering the **show security** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security
  address-book {
    product-design-internal {
      address product-designers 12.1.1.0/24;
      attach {
        zone ls-product-design-trust;
      }
    }
    product-design-external {
      address marketing 13.1.1.0/24;
      address accounting 14.1.1.0/24;
      address others 12.12.1.0/24;
      address-set otherlsys {
        address marketing;
        address accounting;
      }
      attach {
        zone ls-product-design-untrust;
      }
    }
  }
    zones {
      security-zone ls-product-design-trust {
        tcp-rst;
        interfaces {
          ge-0/0/5.1;
        }
      }
      security-zone ls-product-design-untrust {
        interfaces {
          lt-0/0/0.3;
        }
      }
    }
```

Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- Understanding Logical System Zones on page 3299
- User Logical System Configuration Overview on page 3225

Understanding Logical System Screen Options

Junos OS screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone. Junos OS then applies firewall policies, which can contain content filtering and IDP components, to the traffic that passes the screen filters.

All screen options available on the device are available in each logical system. Each user logical system administrator can configure screen options for their user logical system. The master administrator can configure screen options for the master logical system as well as all user logical systems.

The user logical system administrator can configure and view all screen options in a user logical system. All screen options in a user logical system are visible to the master administrator.

Related Documentation

- Example: Configuring Screen Options for a User Logical System on page 3304
- User Logical System Configuration Overview on page 3225
- Attack Detection and Prevention Overview on page 835

Example: Configuring Screen Options for a User Logical System

This example shows how to configure screen options for a user logical system.

- Requirements on page 3304
- Overview on page 3305
- Configuration on page 3305

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Configure zones for the user logical system. See "Example: Configuring Zones for a User Logical System" on page 3301.

Copyright © 2016, Juniper Networks, Inc.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

You can limit the number of concurrent sessions to the same destination IP address in a user logical system. Setting a destination-based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. When the number of concurrent connection requests to an IP address surpasses the limit, Junos OS blocks further connection attempts to that IP address. This example creates the screen options described in Table 341.

Table 341: User Logical System Screen Options Configuration

Name	Configuration Parameters
limit-destination-sessions	• Limits concurrent connection requests to destination IPs to 80. • Applied to ls-product-design-untrust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security screen ids-option limit-destination-sessions limit-session destination-ip-based
  80
set security zones security-zone ls-product-design-untrust screen
  limit-destination-sessions
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure destination-based session limits in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a screen option for a destination-based session limit.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set screen ids-option
     limit-destination-sessions limit-session destination-ip-based 80
   ```

3. Set the security zone for the screen option.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set zones security-zone
     ls-product-design-untrust screen limit-destination-sessions
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security screen** and **show security zone** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
lsdesignadmin1@host:ls-product-design# show security screen
ids-option limit-destination-sessions {
    limit-session {
        destination-ip-based 80;
    }
}
lsdesignadmin1@host:ls-product-design# show security zones
security-zone ls-product-design-trust {
    ...
}
    security-zone ls-product-design-untrust {
        screen limit-destination-sessions;
        ...
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- User Logical System Configuration Overview on page 3225
- Understanding Logical System Screen Options on page 3304

Understanding Logical System Security Policies

- Security Policies in Logical Systems on page 3306
- Application Timeouts on page 3307
- Security Policy Allocation on page 3307

Security Policies in Logical Systems

Security policies enforce rules for what traffic can pass through the firewall and actions that need to take place on the traffic as it passes through the firewall. From the perspective of security policies, traffic enters one security zone and exits another security zone.

By default, a logical system denies all traffic in all directions, including intra-zone and inter-zone directions. Through the creation of security policies, the logical system administrator can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations.

Security policies can be configured in the master logical system and in user logical systems. Configuring a security policy in a logical system is the same as configuring a security policy on a device that is not configured for logical systems. Any security policies, policy rules, address books, applications and application sets, and schedulers created

 Copyright © 2016, Juniper Networks, Inc.

within a logical system are only applicable to that logical system. Only predefined applications and application sets, such as **junos-ftp**, can be shared between logical systems.

NOTE: In a logical system, you cannot specify global as either the from-zone or the to-zone in a security policy.

The user logical system administrator can configure and view all attributes for security policies in a user logical system. All attributes of a security policy in a user logical system are also visible to the master administrator.

Application Timeouts

The application timeout value set for an application determines the session timeout. Application timeout behavior is the same in a logical system as at the root level. However, user logical system administrators can use predefined applications in security policies but cannot modify the timeout value of predefined applications. This is because the predefined applications are shared by the master logical system and all user logical systems, so the user logical system administrator is not allowed to change its behavior. Application timeout values are stored in the application entry database and in the corresponding logical system TCP and UDP port-based timeout tables.

If the application that is matched for the traffic has a timeout value, that timeout value is used. Otherwise, the lookup proceeds in the following order until an application timeout value is found:

1. The logical system TCP and UDP port-based timeout table is searched for a timeout value.
2. The root TCP and UDP port-based timeout table is searched for a timeout value.
3. The protocol-based default timeout table is searched for a timeout value.

Security Policy Allocation

The master administrator configures the maximum and reserved numbers of security policies for each user logical system. The user logical system administrator can then create security policies in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile policy** command to view the number of security policies allocated to the user logical system.

NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of security policies applied to the master logical system. The number of policies configured in the master logical system count toward the maximum number of policies available on the device.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Example: Configuring Security Policies in a User Logical System on page 3308
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- User Logical System Configuration Overview on page 3225
- Security Policies Overview on page 1085
- Understanding Policy Application Timeout Configuration and Lookup on page 1181

Example: Configuring Security Policies in a User Logical System

This example shows how to configure security policies for a user logical system.

- Requirements on page 3308
- Overview on page 3308
- Configuration on page 3309
- Verification on page 3311

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profiles policy** command to see the security policy resources allocated to the logical system.
- Configure zones and address books. See "Example: Configuring Zones for a User Logical System" on page 3301.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example configures the security policies described in Table 342.

Table 342: User Logical System Security Policies Configuration

Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: product-designers • Destination address: otherlsys • Application: any

Copyright © 2016, Juniper Networks, Inc.

Table 342: User Logical System Security Policies Configuration *(continued)*

Name	Configuration Parameters
permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-designers • Application: any

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match source-address product-designers
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys then permit
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match application any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

   ```
   [edit security policies from-zone ls-product-design-trust to-zone
     ls-product-design-untrust]
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
     source-address product-designers
   ```

Copyright © 2016, Juniper Networks, Inc.

```
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
  destination-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
  application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys then
  permit
```

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

```
[edit security policies from-zone ls-product-design-untrust to-zone
  ls-product-design-trust]
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
  source-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
  destination-address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
  application any
lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys then
  permit
```

Results From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
    from-zone ls-product-design-untrust to-zone ls-product-design-trust {
        policy permit-all-from-otherlsys {
            match {
                source-address otherlsys;
                destination-address product-designers;
                application any;
            }
            then {
                permit;
            }
        }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

 Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Policy Configuration on page 3311

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation

- Understanding Logical System Security Policies on page 3306
- User Logical System Configuration Overview on page 3225
- Troubleshooting Security Policies on page 1162

Understanding Logical System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

The master administrator is responsible for configuring access profiles in the master logical system. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored. Access profiles configured at the master logical system are available to all user logical systems.

The master administrator configures the maximum and reserved numbers of firewall authentications for each user logical system. The user logical system administrator can then create firewall authentications in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile auth-entry** command to view the number of authentication resources allocated to the user logical system.

To configure the access profile, the master administrator uses the **profile** configuration statement at the [**edit access**] hierarchy level in the master logical system. The access profile can also include the order of authentication methods, LDAP or RADIUS server options, and session options.

The user logical system administrator can then associate the access profile with a security policy in the user logical system. The user logical system administrator also specifies the type of authentication:

- With pass-through authentication, a host or a user from one zone tries to access resources on another zone using an FTP, a Telnet, or an HTTP client. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent

Copyright © 2016, Juniper Networks, Inc.

traffic from the user or host is allowed or denied based on the result of this authentication.

- With Web authentication, users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

The user logical system administrator configures the following properties for firewall authentication in the user logical system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the **firewall-authentication** configuration statement at the [**edit security policies from-zone** *zone-name* **to-zone** *zone-name* **policy** *policy-name* **then permit**] hierarchy level.

  Users or user groups in an access profile who are allowed access by the policy can optionally be specified with the client-match configuration statement. (If no users or user groups are specified, any user who is successfully authenticated is allowed access.)

  For pass-through authentication, the access profile can optionally be specified and Web redirect (redirecting the client system to a webpage for authentication) can be enabled.

- Type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the **firewall-authentication** configuration statement at the [**edit access**] hierarchy level.

- Host inbound traffic. Protocols, services, or both are allowed to access the logical system. The types of traffic are configured with the **host-inbound-traffic** configuration statement at the [**edit security zones security-zone** *zone-name*] or [**edit security zones security-zone** *zone-name* **interfaces** *interface-name*] hierarchy levels.

From a user logical system, the user logical system administrator can use the **show security firewall-authentication users** or **show security firewall-authentication history** commands to view the information about firewall users and history for the user logical system. From the master logical system, the master administrator can use the same commands to view information for the master logical system, a specific user logical system, or all logical systems.

Related Documentation

- Example: Configuring Access Profiles (Master Administrators Only) on page 3269
- Example: Configuring Firewall Authentication for a User Logical System on page 3313
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Firewall Authentication for a User Logical System

This example shows how to configure firewall authentication for a user logical system.

- Requirements on page 3313
- Overview on page 3313
- Configuration on page 3314
- Verification on page 3316

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profiles auth-entry** command to see the firewall authentication entries allocated to the logical system.
- Access profiles must be configured in the master logical system by the master administrator. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 3269.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

In this example, users in the ls-marketing-dept and ls-accounting-dept logical systems are required to authenticate when initiating certain connections to the product designers subnet. This example configures the firewall authentication described in Table 343.

NOTE: This example uses the access profile configured in "Example: Configuring Access Profiles (Master Administrators Only)" on page 3269and address book entries configured in "Example: Configuring Zones for a User Logical System" on page 3301.

Copyright © 2016, Juniper Networks, Inc.

Table 343: User Logical System Firewall Authentication Configuration

Feature	Name	Configuration Parameters
Security policy	permit-authorized-users NOTE: Policy lookup is performed in the order that the policies are configured. The first policy that matches the traffic is used. If you have previously configured a policy that permits traffic for the same from-zone, to-zone, source address, and destination address but with application **any**, the policy configured in this example would never be matched. (See "Example: Configuring Security Policies in a User Logical System" on page 3308.) Therefore, this policy should be reordered so that it is checked first.	Permit firewall authentication for the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-engineers • Application: junos-h323 The ldap1 access profile is used for pass-through authentication.
Firewall authentication		• Pass-through authentication • HTTP login prompt "welcome" • Default access profile ldap1

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match application junos-h323
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users then permit firewall-authentication pass-through
    access-profile ldap1
set access firewall-authentication pass-through default-profile ldap1
set access firewall-authentication pass-through http banner login "welcome"
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure firewall authentication in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

```
lsdesignadmin1@host:ls-product-design> configure
lsdesignadmin1@host:ls-product-design#
```

2. Configure a security policy that permits firewall authentication.

```
[edit security policies from-zone ls-product-design-untrust to-zone
    ls-product-design-trust]
```


Copyright © 2016, Juniper Networks, Inc.

```
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
  source-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
  destination -address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
  application junos-h323
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users then
  permit firewall-authentication pass-through access-profile ldap1
```

3. Reorder the security policies.

```
[edit]
lsdesignadmin1@host:ls-product-design# insert security policies from-zone
  ls-product-design-untrust to-zone ls-product-design-trust policy
  permit-authorized-users before policy permit-all-from-otherlsys
```

4. Configure firewall authentication.

```
[edit access firewall-authentication]
lsdesignadmin1@host:ls-product-design# set pass-through http banner login
  "welcome"
lsdesignadmin1@host:ls-product-design# set pass-through default-profile ldap1
```

Results From configuration mode, confirm your configuration by entering the **show security policies** and **show access firewall-authentication** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
  policy permit-all-to-otherlsys {
    match {
      source-address product-designers;
      destination-address otherlsys;
      application any;
    }
    then {
      permit;
    }
  }
}
  from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy permit-authorized-users {
      match {
        source-address otherlsys;
        destination-address product-designers;
        application junos-h323;
      }
      then {
        permit {
          firewall-authentication {
            pass-through {
              access-profile ldap1;
            }
          }
        }
```

```
        }
      }
      policy permit-all-from-otherlsys {
        match {
          source-address otherlsys;
          destination-address product-designers;
          application any;
        }
        then {
          permit;
        }
      }
    }
lsdesignadmin1@host:ls-product-design# show access firewall-authentication
pass-through {
  default-profile ldap1;
  http {
    banner {
      login welcome;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses on page 3316

Verifying Firewall User Authentication and Monitoring Users and IP Addresses

Purpose Display firewall authentication user history and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Action From operational mode, enter these **show** commands.

```
lsdesignadmin1@host:ls-product-design> show security firewall-authentication history
lsdesignadmin1@host:ls-product-design> show security firewall-authentication history
  identifier id
lsdesignadmin1@host:ls-product-design> show security firewall-authentication users
lsdesignadmin1@host:ls-product-design> show security firewall-authentication users
  identifier id
```

Related Documentation

- Example: Configuring Access Profiles (Master Administrators Only) on page 3269
- Understanding Logical System Firewall Authentication on page 3267
- User Logical System Configuration Overview on page 3225
- Example: Configuring Pass-Through Authentication on page 5207

Copyright © 2016, Juniper Networks, Inc.

IDP in Logical Systems Overview

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a logical system.

This topic includes the following sections:

- IDP Policies on page 3317
- IDP Installation and Licensing for Logical Systems on page 3318

IDP Policies

The master administrator configures IDP policies at the root level. Configuring an IDP policy for logical systems is similar to configuring an IDP policy on a device that is not configured for logical systems. This can include the configuration of custom attack objects.

NOTE: User logical system administrators cannot create or modify IDP policies for their user logical systems. Only the master administrator can create IDP policies and bind them to user logical systems through a logical systems security profile.

NOTE: The user logical system administrator can create security zones in the user logical system and assign interfaces to each security zone. Zones that are specific to user logical systems cannot be referenced in IDP policies configured by the master administrator. The master administrator can reference zones in the master logical system in an IDP policy configured for the master logical system.

The master administrator then specifies an IDP policy in the security profile that is bound to a logical system. To enable IDP in a logical system, the master administrator or user logical system administrator configures a security policy that defines the traffic to be inspected and specifies the **permit application-services idp** action.

Although the master administrator can configure multiple IDP policies, a logical system can have only one active IDP policy at a time. For user logical systems, the master administrator can either bind the same IDP policy to multiple user logical systems or bind a unique IDP policy to each user logical system. To specify the active IDP policy for the master logical system, the master administrator can *either* reference the IDP policy in the security profile that is bound to the master logical system or use the **active-policy** configuration statement at the [**edit security idp**] hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

NOTE: A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the active-policy configuration statement. Use only one method to specify the active IDP policy for the master logical system.

IDP Installation and Licensing for Logical Systems

A single IDP security package is installed for all logical systems on the device. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all logical systems.

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any logical system on the device.

Related Documentation

- Understanding IDP Features in Logical Systems on page 3278
- Example: Configuring an IDP Policy for a User Logical System on page 3321
- Example: Configuring an IDP Policy for the Master Logical System on page 3281
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- *IDP Policies Overview*

Understanding IDP Features in Logical Systems

This topic includes the following sections:

- Rulebases on page 3318
- Protocol Decoders on page 3319
- SSL Inspection on page 3319
- Inline Tap Mode on page 3319
- Multi-Detectors on page 3319
- Logging and Monitoring on page 3320

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The following IDP rulebases are supported for logical systems:

- The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.
- The application-level distributed denial-of-service (DDoS) rulebase defines parameters to protect servers such as DNS or HTTP. The application-level DDoS rulebase defines the source match condition for traffic that should be monitored and takes an action,

Copyright © 2016, Juniper Networks, Inc.

such as drop the connection, drop the packet, or no action. It can also perform actions against future connections that use the same IP address.

NOTE: Status monitoring for IPS and application-level DDoS is global to the device and not on a per logical system basis.

Protocol Decoders

The Junos IDP module ships with a set of preconfigured protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks that they perform. The IDP protocol decoder configuration is global and applies to all logical systems. Only the master administrator at the root level can modify the settings at the [**edit security idp sensor-configuration**] hierarchy level.

SSL Inspection

IDP SSL inspection uses the Secure Sockets Layer (SSL) protocol suite to enable inspection of HTTP traffic encrypted in SSL.

SSL inspection configuration is global and applies to all logical systems on a device. SSL inspection can only be configured by the master administrator at the root level with the **ssl-inspection** configuration statement at the [**edit security idp sensor-configuration**] hierarchy level.

Inline Tap Mode

The inline tap mode feature provides passive, inline detection of Application Layer threats for traffic matching security policies that have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module. This allows the packets to get to the next service module without waiting for IDP processing results.

Inline tap mode is enabled or disabled for all logical systems at the root level by the master administrator. To enable inline tap mode, use the **inline-tap** configuration statement at the [**edit security forwarding-process application-services maximize-idp-sessions**] hierarchy level. Delete the inline tap mode configuration to switch the device back to regular mode.

NOTE: The device must be restarted when switching to inline tap mode or back to regular mode.

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded

Copyright © 2016, Juniper Networks, Inc.

along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all logical systems.

Logging and Monitoring

Status monitoring options are available to the master administrator only. All status monitoring options under the **show security idp** and **clear security idp** CLI operational commands present global information, but not on a per logical system basis.

NOTE: SNMP monitoring for IDP is not supported on logical systems.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The logical systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the ls-product-design logical system:

  Oct 12 17:33:32 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In ls-product-design at 1286930013, SIG Attack log <4.0.0.1/34327->5.0.0.1/21> for TCP protocol and service SERVICE_IDP application NONE by rule 1 of rulebase IPS in policy Recommended. attack: repeat=0, action=IGNORE, threat-severity=MEDIUM, name=FTP:USER:ROOT, NAT <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:ls-product-design-untrust:ge-0/0/0.0->ls-product-design-trust:ge-0/0/1.0, packet-log-id: 65535 and misc-message -

- IP action logs. The following example shows an IP action log for the ls-product-design logical system:

  Oct 13 16:56:04 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In ls-product-design at 1287014163, TRAFFIC Attack log <25.0.0.1/34802->15.0.0.1/21> for TCP protocol and service SERVICE_NONE application NONE by rule 1 of rulebase IPS in policy Recommended. attack: repeat=0, action=TRAFFIC_IPACTION_NOTIFY, threat-severity=INFO, name=_, NAT <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0, intf:ls-product-design-trust:ge-0/0/1.0->ls-product-design-untrust:plt0.3, packet-log-id: 0 and misc-message -

- Application DDoS logs. The following example shows an application DDoS log for the ls-product-design logical system:

  Oct 11 16:29:57 8.0.0.254 RT_IDP: IDP_APPDDOS_APP_ATTACK_EVENT_LS: DDOS Attack in ls-product-design at 1286839797 on my-http, <ls-product-design-untrust:ge-0/0/0.0:4.0.0.1:33738->ls-product-design-trust:ge-0/0/1.0:5.0.0.1:80> for TCP protocol and service HTTP by rule 1 of rulebase DDOS in policy Recommended. attack: repeats 0 action DROP threat-severity INFO, connection-hit-rate 0, context-name http-url-parsed, hit-rate 6, value-hit-rate 6 time-scope PEER time-count 2 time-period 10 secs, context value: ascii: /abc.html hex: 2f 61 62 63 2e 68 74 6d 6c

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *Understanding IDP Policy Rule Bases*
- *Understanding IDP Protocol Decoders*
- *IDP SSL Overview*
- *Understanding IDP Inline Tap Mode*
- *Understanding Multiple IDP Detector Support*
- *Understanding IDP Logging*

Example: Configuring an IDP Policy for a User Logical System

The master administrator can *either* download predefined IDP policies to the device or configure custom IDP policies at the root level using custom or predefined attack objects. The master administrator is responsible for assigning an IDP policy to a user logical system. This example shows how to assign a predefined IDP policy to a user logical system.

- Requirements on page 3321
- Overview on page 3321
- Configuration on page 3322
- Verification on page 3323

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.
- Read *IDP Policies Overview*.
- Assign the ls-design-profile security policy to the ls-product-design user logical system. See "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 3258.
- Download predefined IDP policy templates to the device. See *Downloading and Using Predefined IDP Policy Templates (CLI Procedure)*.

NOTE: **Activating a predefined IDP policy with the active-policy configuration statement at the [edit security idp] hierarchy level only applies to the master logical system. For a user logical system, the master administrator specifies the active IDP policy in the security profile that is bound to the user logical system.**

Overview

The predefined IDP policy named Recommended contains attack objects recommended by Juniper Networks. All rules in the policy have their actions set to take the recommended

Copyright © 2016, Juniper Networks, Inc.

action for each attack object. You add the Recommended IDP policy to the ls-design-profile, which is bound to the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system security-profile ls-design-profile idp-policy Recommended
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To add a predefined IDP policy to a security profile for a user logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   [edit]
   admin@host> configure
   admin@host#
   ```

2. Add the IDP policy to the security profile.

   ```
   [edit system security-profile]
   admin@host# set ls-design-profile idp-policy Recommended
   ```

Results

From configuration mode, confirm your configuration by entering the **show security idp** and **show system security-profile ls-design-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp
  idp-policy Recommended {
    ...
  }
[edit]
admin@host# show system security-profile ls-design-profile
  policy {
    ...
  }
  idp-policy Recommended;
logical-system ls-product-design;
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying the Configuration

Purpose Verify the IDP policy assigned to the logical system.

Action From operational mode, enter the **show security idp logical-system policy-association** command. Ensure that the IDP policy in the security profile that is bound to the logical system is correct.

```
admin@host> show security idp logical-system policy-association
  Logical system      IDP policy
  ls-product-design   Recommended
```

Related Documentation

- Example: Enabling IDP in a User Logical System Security Policy on page 3323
- IDP in Logical Systems Overview on page 3276
- User Logical System Configuration Overview on page 3225

Example: Enabling IDP in a User Logical System Security Policy

This example shows how to enable IDP in a security policy in a user logical system.

- Requirements on page 3323
- Overview on page 3323
- Configuration on page 3324
- Verification on page 3325

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profiles idp-policy** command to see the security policy resources allocated to the logical system.
- Configure an IDP security policy for the user logical system as the master administrator. See "Example: Configuring an IDP Policy for a User Logical System" on page 3321.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

You enable IDP in a security policy that matches any traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone. Enabling IDP in a security policy directs matching traffic to be checked against the IDP rulebases.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This example uses the IDP policy configured and assigned to the ls-product-design user logical system by the master administrator in "Example: Configuring an IDP Policy for a User Logical System" on page 3321.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match source-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match destination-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match application any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp then permit application-services idp
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a security policy to enable IDP in a user logical system:

1. Log in to the logical system as the user logical system administrator and enter configuration mode.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design>configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security policy that matches traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

   ```
   [edit security policies from-zone ls-product-design-untrust to-zone
       ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match source-address
       any
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match
       destination-address any
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match application
       any
   ```

3. Configure the security policy to enable IDP for matching traffic.

   ```
   [edit security policies from-zone ls-product-design-untrust to-zone
       ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy enable-idp then permit
       application-services idp
   ```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security policies
  from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy enable-idp {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          application-services {
            idp;
          }
        }
      }
    }
    ...
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Attack Matches

Purpose Verify that attacks are being matched in network traffic.

Action From operational mode, enter the **show security idp attack table** command.

```
admin@host> show security idp attack table
  IDP attack statistics:
    Attack name                    #Hits
    FTP:USER:ROOT                     1
```

Related Documentation

- Example: Configuring an IDP Policy for a User Logical System on page 3321
- IDP in Logical Systems Overview on page 3276
- User Logical System Configuration Overview on page 3225

Copyright © 2016, Juniper Networks, Inc.

Understanding Logical System Application Identification Services

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following benefits:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.
- Improves performance by narrowing the scope of attack signatures for applications without decoders.
- Enables you to create detailed reports using AppTrack on applications passing through the device.

With logical systems, predefined and custom application signatures are global resources that are shared by all logical systems. The master administrator is responsible for downloading and installing predefined Juniper Networks application signatures and creating custom application and nested application signatures to identify applications that are not part of the predefined database.

Application identification is enabled by default.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each user logical system has its own ASC. A user logical system administrator can display the ASC entries for their logical system with the **show services application-identification application-system-cache** command. A user logical system administrator can use the **clear services application-identification application-system-cache** command to clear the ASC entries for their logical system.

The master administrator can display or clear ASC entries for any logical system. The master administrator can also display or clear global counters with the **show services application-identification counter** and **clear services application-identification counter** commands.

Related Documentation

- Understanding the Junos OS Application Identification Database on page 488
- Example: Scheduling the Application Signature Package Updates on page 499
- Example: Configuring Junos OS Application Identification Custom Application Signatures on page 512
- *Understanding IDP Application Identification*
- Understanding the Application System Cache on page 523
- Verifying Application System Cache Statistics on page 524

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Application Firewall Services for a User Logical System

This example describes how to configure application firewall services on a user logical system by a user logical system administrator. User logical system administrators can manage and monitor their own system application firewall rule sets and rules and manage the dynamic applications allowed or blocked on their respective logical systems.

After configuring application firewall rule sets and rules, user logical system administrators add the application firewall rule set information to the security policy on their individual logical systems.

For information about configuring an application firewall within a security policy, see "Application Firewall Overview" on page 555.

- Requirements on page 3327
- Overview on page 3327
- Configuration on page 3328
- Verification on page 3329

Requirements

Before you begin:

- Verify that the security zones are configured for the user logical system.
- Verify that the master administrator has allocated application firewall resources (appfw-rule-set and appfw-rule) in the security profile bound to the user logical system.

  For more information, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 3253.
- Log in to the logical system as the user logical system administrator.

  For information about user logical system administrator role functions, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 3227.

Overview

In this example you configure application firewall services on the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example creates the following application firewall configuration:

- Rule set, ls-product-design-rs1, with rules r1 and r2. When r1 is matched, Telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.
- Rule set, ls-product-design-rs2, with rule r1. When r1 is matched, Facebook traffic is blocked by the firewall.

Copyright © 2016, Juniper Networks, Inc.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security application-firewall rule-sets ls-product-design-rs1 rule r1 match
  dynamic-application junos:TELNET
set security application-firewall rule-sets ls-product-design-rs1 rule r1 then permit
set security application-firewall rule-sets ls-product-design-rs1 rule r2 match
  dynamic-application-group junos:web
set security application-firewall rule-sets ls-product-design-rs1 rule r2 then permit
set security application-firewall rule-sets ls-product-design-rs1 default-rule deny
set security application-firewall rule-sets ls-product-design-rs2 rule r1 match
  dynamic-application junos:facebook
set security application-firewall rule-sets ls-product-design-rs2 rule r1 then deny
set security application-firewall rule-sets ls-product-design-rs2 default-rule permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure application firewall for a user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure an application firewall rule set for this logical system.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
     ls-product-design-rs1
   ```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
     ls-product-design-rs1 rule r1 match dynamic-application telnet then permit
   ```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
     ls-product-design-rs1 default-rule deny
   ```

5. Repeat these steps to configure another rule set, ls-product-design-rs2, if desired.

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security application-firewall rule-set all** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security application-firewall rule-set all
...
  application-firewall {
    rule-sets ls-product-design-rs1 {
      rule r1 {
        match {
          dynamic-application [junos:TELNET];
        }
        then {
          permit;
        }
      }
      default-rule {
        deny;
      }
    }
    rule-sets ls-product-design-rs1 {
      rule r2 {
        match {
          dynamic-application-group [junos:web];
        }
        then {
          permit;
        }
      }
      rule-sets ls-product-design-rs2 {
        rule r1 {
          match {
            dynamic-application [junos:FACEBOOK];
          }
          then {
            deny;
          }
        }
        default-rule {
          permit;
        }
      }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Application Firewall Configuration on page 3330

Copyright © 2016, Juniper Networks, Inc.

Verifying Application Firewall Configuration

Purpose View the application firewall configuration on the user logical system.

Action From operational mode, enter the **show security application-firewall rule-set all** command.

```
lsdesignadmin1@host:ls-product-design> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Related Documentation

- User Logical System Configuration Overview on page 3225
- Understanding Logical System Application Firewall Services on page 3287

Understanding Logical System Application Tracking Services

AppTrack is an application tracking tool that provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack can be enabled and configured within any logical system. Configuring AppTrack in a logical system is the same as configuring AppTrack on a device that is not configured for logical systems. An AppTrack configuration only applies to the logical system in which it is configured. The name of the logical system is added to AppTrack logs. The master

 Copyright © 2016, Juniper Networks, Inc.

administrator can configure AppTrack for any logical system while a user logical system administrator can only configure AppTrack for the logical system that they are logged in to.

NOTE: **The system log configuration is global on the device and must be configured by the master administrator. The user logical system administrator cannot configure system logging for a logical system.**

Counters keep track of the number of log messages sent and logs that have failed. AppTrack counters are global to the device. The master administrator as well as user logical system administrators can view AppTrack counters with the **show security application-tracking counters** command.

Related Documentation

- Understanding AppTrack on page 573
- Example: Configuring AppTrack on page 574
- Example: Configuring AppTrack for a User Logical System on page 3331

Example: Configuring AppTrack for a User Logical System

This example shows how to configure the AppTrack tracking tool so you can analyze the bandwidth usage of your network.

- Requirements on page 3331
- Overview on page 3331
- Configuration on page 3332
- Verification on page 3333

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- (Master administrator) Configure system logging in the master logical system. See *Network Monitoring and Troubleshooting Guide for Security Devices*.

Overview

This example shows how to enable application tracking for the security zone ls-product-design-trust in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

The first message is generated at session start and update messages are sent every 5 minutes after that or until the session ends. A final message is sent at session end.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security zones security-zone ls-product-design-trust application-tracking
set security application-tracking first-update
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure AppTrack for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Enable AppTrack for the security zone.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set zones security-zone
     ls-product-design-trust application-tracking
   ```

3. Generate update messages at session start and at 5-minute intervals.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set application-tracking first-update
   ```

Results

From configuration mode, confirm your configuration by entering the **show security** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security
...
  application-tracking {
    first-update;
  }
...
  zones {
    security-zone ls-product-design-trust {
      ...
      application-tracking;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying AppTrack Operation on page 3333
- Verifying Security Flow Session Statistics on page 3333
- Verifying Application System Cache Statistics on page 3333
- Verifying the Status of Application Identification Counter Values on page 3333

Verifying AppTrack Operation

Purpose View the AppTrack counters periodically to monitor tracking.

Action From operational mode, enter the **show application-tracking counters** command.

Verifying Security Flow Session Statistics

Purpose Compare byte and packet counts in logged messages with the session statistics from the **show security flow session** command output.

Action From operational mode, enter the **show security flow session** command.

Verifying Application System Cache Statistics

Purpose Compare cache statistics such as IP address, port, protocol, and service for an application from the **show services application-identification application-system-cache** command output.

Action From operational mode, enter the **show services application-identification application-system-cache** command.

Verifying the Status of Application Identification Counter Values

Purpose Compare session statistics for application identification counter values from the **show services application-identification counter** command output.

Action From operational mode, enter the **show services application-identification counter** command.

Related Documentation

- Understanding Logical System Application Tracking Services on page 3292
- User Logical System Configuration Overview on page 3225

Copyright © 2016, Juniper Networks, Inc.

Understanding Route-Based VPN Tunnels in Logical Systems

A VPN connection can secure traffic that passes between a logical system and a remote site across a WAN. With route-based VPNs, you configure one or more security policies in a logical system to regulate the traffic flowing through a single IP Security (IPsec) tunnel. For each IPsec tunnel, there is one set of IKE and IPsec security associations (SAs) that must be configured at the root level by the master administrator.

NOTE: The external interface configured under the gateway configuration can only be a part of the root logical system.

NOTE: Only route-based VPNs are supported for logical systems. Policy-based VPNs are not supported.

In addition to configuring IKE and IPsec SAs for each VPN, the master administrator must also assign a secure tunnel (st0) interface to a user logical system. An st0 interface can only be assigned to a single user logical system. However, multiple user logical systems can each be assigned their own st0 interface.

NOTE: The st0 unit 0 interface should not be assigned to a logical system, as an SA cannot be set up for this interface.

The user logical system administrator can configure the IP address and other attributes of the st0 interface assigned to the user logical system. The user logical system administrator cannot delete an st0 interface assigned to their user logical system.

For route-based VPNs, a security policy refers to a destination address and not a specific VPN tunnel. For cleartext traffic in a user logical system to be sent to the VPN tunnel for encapsulation, the user logical system administrator must make the following configurations:

- Security policy that permits traffic to a specified destination.
- Static route to the destination with the st0 interface as the next hop.

When Junos OS looks up routes in the user logical system to find the interface to use to send traffic to the destination address, it finds a static route through the st0 interface. Traffic is routed to the VPN tunnel as long as the security policy action is permit.

The master logical system and a user logical system can share a route-based VPN tunnel. An st0 interface assigned to a user logical system can also be used by the master logical system. For the master logical system, the master administrator configures a security policy that permits traffic to the remote destination and a static route to the remote destination with the st0 interface as the next hop.

Copyright © 2016, Juniper Networks, Inc.

VPN monitoring is configured by the master administrator in the master logical system. For the VPN monitor source interface, the master administrator must specify the st0 interface; a physical interface for a user logical system cannot be specified.

Related Documentation

- Understanding Route-Based IPsec VPNs on page 6073
- User Logical System Configuration Overview on page 3225
- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) on page 3294
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

Example: Configuring a Route-Based VPN Tunnel in a User Logical System

This example shows how to configure a route-based VPN tunnel in a user logical system.

- Requirements on page 3335
- Overview on page 3335
- Configuration on page 3336
- Verification on page 3337

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Ensure that an st0 interface is assigned to the user logical system and IKE and IPsec SAs are configured at the root level by the master administrator. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 3294.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

You configure the route-based VPN parameters described in Table 344.

Table 344: User Logical System Route-Based VPN Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	• IPv4 protocol family (inet) • IP address 10.11.11.150/24
Static route		• Destination 192.168.168.0/24 • Next hop st0.1

Copyright © 2016, Juniper Networks, Inc.

Table 344: User Logical System Route-Based VPN Configuration *(continued)*

Feature	Name	Configuration Parameters
Security policy	through-vpn	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: any • Destination address: 192.168.168.0/24 • Application: any

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces st0 unit 1 family inet address 10.11.11.150/24
set routing-options static route 192.168.168.0/24 next-hop st0.1
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match source-address any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match destination-address 192.168.168.0/24
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a route-based VPN tunnel in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design>configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure the VPN tunnel interface.

   ```
   [edit interfaces]
   lsdesignadmin1@host:ls-product-design# set st0 unit 1 family inet address
     10.11.11.150/24
   ```

3. Create a static route to the remote destination.

   ```
   [edit routing-options]
   lsdesignadmin1@host:ls-product-design# set static route 192.168.168.0/24 next-hop
     st0.1
   ```

4. Configure a security policy to permit traffic to the remote destination.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies from-zone ls-product-design-trust to-zone
  ls-product-design-untrust]
lsdesignadmin1@host:ls-product-design# set policy through-vpn match
  source-address any
lsdesignadmin1@host:ls-product-design# set policy through-vpn match
  destination-address 192.168.168.0/24
lsdesignadmin1@host:ls-product-design# set policy through-vpn match application
  any
lsdesignadmin1@host:ls-product-design# set policy through-vpn then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces st0**, **show routing-options**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
lsdesignadmin1@host:ls-product-design# show interfaces st0
  unit 1 {
    family inet {
      address 10.11.11.150/24;
    }
  }
  lsdesignadmin1@host:ls-product-design# show routing-options
    static {
      route 192.168.168.0/24 next-hop st0.1;
    }
  [edit]
  lsdesignadmin1@host:ls-product-design# show security policies
    from-zone ls-product-design-trust to-zone ls-product-design-untrust {
      policy through-vpn {
        match {
          source-address any;
          destination-address 192.168.168.0/24;
          application any;
        }
        then {
          permit;
        }
      }
      ...
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Before starting the verification process, you need to send traffic from a host in the user logical system to a host in the 192.168.168.0/24 network. For example, initiate a ping from a host in the 12.1.1.1/24 subnet in the ls-product-design user logical system to the host 192.168.168.10.

- Verifying the IKE Phase 1 Status on page 3338
- Verifying the IPsec Phase 2 Status on page 3338

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

For sample outputs and meanings, see the "Verification" section of "Example: Configuring a Route-Based VPN" on page 6074.

Verifying the IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

For sample outputs and meanings, see the "Verification" section of "Example: Configuring a Route-Based VPN" on page 6074.

Related Documentation

- Example: Configuring a Route-Based VPN on page 6074.
- Understanding Route-Based VPN Tunnels in Logical Systems on page 3293
- User Logical System Configuration Overview on page 3225

Copyright © 2016, Juniper Networks, Inc.

PART 47

Configuring Routing and Interfaces Features

- Configuring Master Logical System Routing and Interfaces on page 3341
- Configuring User Logical System Routing, Interfaces, and NAT Features on page 3355

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 149

Configuring Master Logical System Routing and Interfaces

- Understanding Logical System Interfaces and Routing Instances on page 3341
- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 3342
- Example: Configuring OSPF Routing Protocol for the Master Logical System on page 3350

Understanding Logical System Interfaces and Routing Instances

Logical interfaces on the device are allocated among the user logical systems by the master administrator. The user logical system administrator configures the attributes of the interfaces, including IP addresses, and assigns them to routing instances and zones.

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Interfaces and routing instances can be configured in the master logical system and in user logical systems. Configuring an interface or routing instance in a logical system is the same as configuring an interface or routing instance on a device that is not configured for logical systems. Any routing instance created within a logical system is only applicable to that logical system.

The default routing instance, master, refers to the main inet.0 routing table in the logical system. The master routing instance is reserved and cannot be specified as a routing instance. Routes are installed in the master routing instance by default, unless a routing instance is specified. Configure global routing options and protocols for the master routing instance by including statements at the [**edit protocols**] and [**edit routing-options**] hierarchy levels in the logical system.

You can configure only virtual router routing instance type in a user logical system. Only one virtual private LAN service (VPLS) routing instance type can be configured on the device and it must be in the interconnect logical system.

Copyright © 2016, Juniper Networks, Inc.

The user logical system administrator can configure and view all attributes for an interface or routing instance in a user logical system. All attributes of an interface or routing instance in a user logical system are also visible to the master administrator.

Multicast is a "one source, many destinations" method of traffic distribution, which means the destinations needing to receive the information from a particular source receive the traffic stream. The master and user logical system administrators can configure a logical system to support multicast applications. The same multicast configurations to configure a device as a node in a multicast network can be used in a logical system.

Related Documentation

- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 3360
- User Logical System Configuration Overview on page 3225
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)

This topic covers configuration of interfaces, static routes, and routing instances for the master and interconnect logical systems. It also covers configuration of logical tunnel interfaces for user logical systems.

- Requirements on page 3342
- Overview on page 3343
- Configuration on page 3344
- Verification on page 3350

Requirements

The example uses an SRX5600 device running Junos operating system (Junos OS) with logical systems.

Before you begin:

- Read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 to understand how and where this procedure fits in the overall master administrator configuration process.
- Read "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210

Copyright © 2016, Juniper Networks, Inc.

Overview

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, lt-0/0/0.4, and lt-0/0/0.6. The example configures a routing instance called vr-ic and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a virtual private LAN service (VPLS) routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the master and user logical systems are configured with Ethernet as the encapsulation type.

  - lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.
  - lt-0/0/0.2 connects to lt-0/0/0.3 on the ls-product-design logical system.
  - lt-0/0/0.4 connects to lt-0/0/0.5 on the ls-marketing-dept logical system.
  - lt-0/0/0.6 connects to lt-0/0/0.7 on the ls-accounting-dept logical system.

- For the master logical system, called root-logical-system, the example configures ge-0/0/4.0 and assigns it to the vr1-root routing instance. The example configures lt-0/0/0.1 to connect to lt-0/0/0.0 on the interconnect logical system and assigns it to the vr1-root routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr1-root routing instance.
- For the ls-product-design logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.
- For the ls-marketing-dept logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.
- For the ls-accounting-dept logical system, the example configures lt-0/0/0.7 to connect to lt-0/0/0.6 on the interconnect logical system.

Figure 134 shows the topology for this deployment including virtual routers and their interfaces for all logical systems.

Copyright © 2016, Juniper Networks, Inc.

Figure 134: Configuring Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

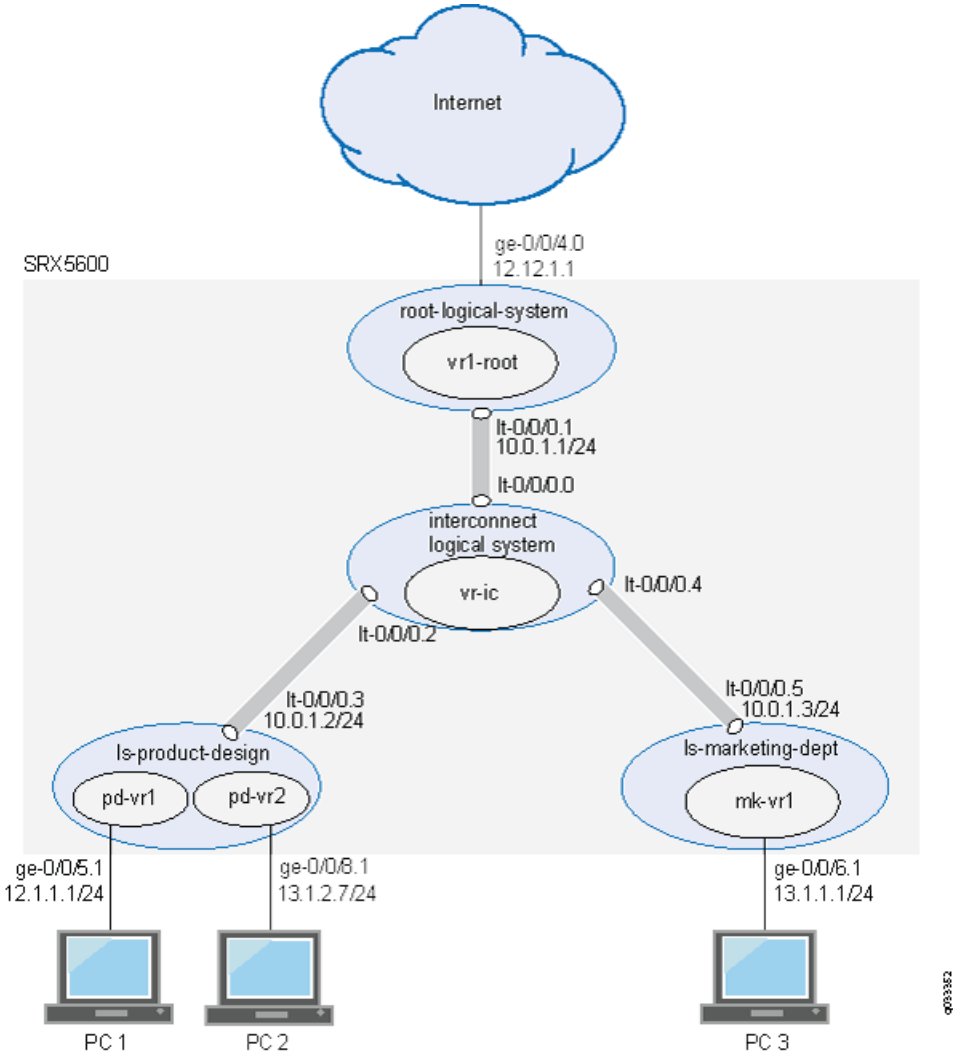

Configuration

This topic explains how to configure interfaces for logical systems.

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System on page 3344
- Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System on page 3346
- Configuring Logical Tunnel Interfaces for the User Logical Systems on page 3348

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation
  ethernet-vpls
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 2 encapsulation
  ethernet-vpls
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 4 encapsulation
  ethernet-vpls
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 6 encapsulation
  ethernet-vpls
set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 6 peer-unit 7
set logical-systems interconnect-logical-system routing-instances vr-ic instance-type
  vpls
set logical-systems interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.0
set logical-systems interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.2
set logical-systems interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.4
set logical-systems interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.6
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Configure the lt-0/0/0 interfaces.

```
[edit logical-systems]
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation
  ethernet-vpls
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 0 peer-unit 1
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 2 encapsulation
  ethernet-vpls
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 2 peer-unit 3
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 4 encapsulation
  ethernet-vpls
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 4 peer-unit 5
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 6 encapsulation
  ethernet-vpls
user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 6 peer-unit 7
```

2. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

```
[edit logical-systems]
user@host# set interconnect-logical-system routing-instances vr-ic instance-type
  vpls
user@host# set interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.2
user@host# set interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.4
user@host# set interconnect-logical-system routing-instances vr-ic interface
  lt-0/0/0.6
```

Results From configuration mode, confirm your configuration by entering the **show logical-systems interconnect-logical-system** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

If you are done configuring the device, enter **commit** from configuration mode.

```
user@host# show logical-systems interconnect-logical-system
  interfaces {
    lt-0/0/0 {
      unit 0 {
        encapsulation ethernet-vpls;
        peer-unit 1;
      }
      unit 2 {
        encapsulation ethernet-vpls;
        peer-unit 3;
      }
      unit 4 {
        encapsulation ethernet-vpls;
        peer-unit 5;
      }
      unit 6 {
        encapsulation ethernet-vpls;
        peer-unit 7;
      }
    }
  }
    routing-instances {
      vr-ic {
        instance-type vpls;
        interface lt-0/0/0.0;
        interface lt-0/0/0.2;
        interface lt-0/0/0.4;
        interface lt-0/0/0.6;
      }
    }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 vlan-id 600
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/4 unit 0 family inet address 12.12.1.1/24
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/6 vlan-tagging
set interfaces ge-0/0/7 vlan-tagging
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 10.0.1.1/24
set routing-instances vr1-root instance-type virtual-router
set routing-instances vr1-root interface ge-0/0/4.0
set routing-instances vr1-root interface lt-0/0/0.1
set routing-instances vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the master logical system interfaces:

1. Configure the master (root) logical and lt-0/0/0.1 interfaces.

```
[edit interfaces]
user@host# set ge-0/0/4 vlan-tagging
user@host# set ge-0/0/4 unit 0 vlan-id 600
user@host# set ge-0/0/4 unit 0 family inet address 12.12.1.1/24
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 10.0.1.1/24
```

2. Configure the interfaces for other logical systems to support VLAN tagging.

```
[edit interfaces]
user@host# set ge-0/0/5 vlan-tagging
user@host# set ge-0/0/6 vlan-tagging
user@host# set ge-0/0/7 vlan-tagging
```

3. Configure a routing instance for the master logical system, assign its interfaces to it, and configure static routes for it.

```
[edit routing-instances]
user@host# set vr1-root instance-type virtual-router
user@host# set vr1-root interface ge-0/0/4.0
user@host# set vr1-root interface lt-0/0/0.1
user@host# set vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
user@host# set vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
user@host# set vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
  ge-0/0/4 {
    vlan-tagging;
```

Copyright © 2016, Juniper Networks, Inc.

```
    unit 0 {
      vlan-id 600;
      family inet {
        address 12.12.1.1/24;
      }
    }
  }
    ge-0/0/5 {
      vlan-tagging;
    }
    ge-0/0/6 {
      vlan-tagging;
    }
    ge-0/0/7 {
      vlan-tagging;
    }
    lt-0/0/0 {
      unit 1 {
        encapsulation ethernet;
        peer-unit 0;
        family inet {
          address 10.0.1.1/24;
        }
      }
    }

[edit]
user@host# show routing-instances
  vr1-root {
    instance-type virtual-router;
    interface ge-0/0/4.0;
    interface lt-0/0/0.1;
    routing-options {
      static {
        route 14.1.1.0/24 next-hop 10.0.1.4;
        route 12.1.1.0/24 next-hop 10.0.1.2;
        route 13.1.1.0/24 next-hop 10.0.1.3;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Logical Tunnel Interfaces for the User Logical Systems

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 family inet address
  10.0.1.2/24
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
```


Copyright © 2016, Juniper Networks, Inc.

```
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address
  10.0.1.3/24
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address
  10.0.1.4/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Configure the lt-0/0/0 interface for the first user logical system:

```
[edit logical-systems]
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
user@host# set ls-product-design interfaces lt-0/0/0 unit 3 family inet address
  10.0.1.2/24
```

2. Configure the lt-0/0/0 interface for the second user logical system.

```
[edit logical-systems]
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address
  10.0.1.3/24 face
```

3. Configure the lt-0/0/0 interface for the third user logical system.

```
[edit logical-systems]
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address
  10.0.1.4/24
```

Results

From configuration mode, confirm your configuration by entering the **show logical-systems ls-product-design interfaces lt-0/0/0**, **show logical-systems ls-marketing-dept interfaces lt-0/0/0** , and **show logical-systems ls-accounting-dept interfaces lt-0/0/0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host#  show logical-systems ls-product-design interfaces lt-0/0/0

lt-0/0/0 {
  unit 3 {
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.0.1.2/24;
    }
  }
}
user@host# show logical-systems ls-marketing-dept interfaces lt-0/0/0
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
```

Copyright © 2016, Juniper Networks, Inc.

```
        family inet {
          address 10.0.1.3/24;
        }
      }
    }
  }
user@host# show logical-systems ls-accounting-dept interfaces lt-0/0/0
lt-0/0/0 {
  unit 7 {
    encapsulation ethernet;
    peer-unit 6;
    family inet {
      address 10.0.1.4/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Static Routes Configured for the Master Administrator Are Correct on page 3350

Verifying That the Static Routes Configured for the Master Administrator Are Correct

Purpose Verify if you can send data from the master logical system to the other logical systems.

Action From operational mode, use the **ping** command.

Related Documentation

- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210

Example: Configuring OSPF Routing Protocol for the Master Logical System

This example shows how to configure OSPF for the master logical system.

- Requirements on page 3351
- Overview on page 3351
- Configuration on page 3351
- Verification on page 3353

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 3241.
- Configure logical interfaces ge-0/0/4.0 and lt-0/0/0.1 for the master logical system and assign them to the vr1-root routing instance. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 3342.

Overview

In this example, you configure OSPF for the master logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example enables OSPF routing on the ge-0/0/4.0 and lt-0/0/0.1 interfaces in the master logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the vr1-root routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.
- ospf-redist-static—Static routes.
- ospf-to-ospf—Routes learned from OSPF.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances vr1-root protocols ospf export ospf-redist-direct
set routing-instances vr1-root protocols ospf export ospf-redist-static
set routing-instances vr1-root protocols ospf export ospf-to-ospf
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF for the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

```
admin@host> configure
admin@host#
```

2. Create routing policies that accept routes.

```
[edit policy-options]
admin@host# set policy-statement ospf-redist-direct from protocol direct
admin@host# set policy-statement ospf-redist-direct then accept
admin@host# set policy-statement ospf-redist-static from protocol static
admin@host# set policy-statement ospf-redist-static then accept
admin@host# set policy-statement ospf-to-ospf from protocol ospf
admin@host# set policy-statement ospf-to-ospf then accept
```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

```
[edit routing-instances]
admin@host# set vr1-root protocols ospf export ospf-redist-direct
admin@host# set vr1-root protocols ospf export ospf-redist-static
admin@host# set vr1-root protocols ospf export ospf-to-ospf
```

4. Enable OSPF on the logical interfaces.

```
[edit routing-instances]
admin@host# set vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
admin@host# set vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
```

Results

From configuration mode, confirm your configuration by entering the **show policy-options** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  admin@host# show policy-options
  policy-statement ospf-redist-direct {
    from protocol direct;
    then accept;
  }
  policy-statement ospf-redist-static {
    from protocol static;
    then accept;
  }
  policy-statement ospf-to-ospf {
    from protocol ospf;
    then accept;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    [edit]
  admin@host# show routing-instances
    vr1-root {
      ...
      protocols {
        ospf {
          export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
          area 0.0.0.1 {
            interface lt-0/0/0.1;
            interface ge-0/0/4.0;
          }
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying OSPF Interfaces on page 3353
- Verifying OSPF Neighbors on page 3353
- Verifying OSPF Routes on page 3353

Verifying OSPF Interfaces

Purpose Verify OSPF-enabled interfaces.

Action From the CLI, enter the **show ospf interface instance vr1-root** command.

```
admin@host> show ospf interface instance vr1-root
Interface           State   Area            DR ID           BDR ID          Nbrs
lt-0/0/0.1          DR      0.0.0.0         10.0.1.1        0.0.0.0            0
ge-0/0/4.0          DR      0.0.0.1         10.0.1.1        0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose Verify OSPF neighbors.

Action From the CLI, enter the **show ospf neighbor instance vr1-root** command.

```
admin@host> show ospf neighbor instance vr1-root
Address   Interface  State   ID     Pri Dead
10.0.1.2  plt0.3    Full    0.0.0.0    128   39
```

Verifying OSPF Routes

Purpose Verify OSPF routes.

Action From the CLI, enter the **show ospf route instance vr1-root** command.

```
admin@host> show ospf route instance vr1-root
```

Copyright © 2016, Juniper Networks, Inc.

```
Topology default Route Table:

Prefix              Path  Route      NH       Metric NextHop       Nexthop
                    Type  Type       Type             Interface     Address/LSP
10.0.1.0/24         Intra Network    IP            1 lt-0/0/0.1
12.12.1.0/24        Intra Network    IP            1 ge-0/0/4.0
```

Related Documentation

- Understanding Logical System Interfaces and Routing Instances on page 3341
- Example: Configuring OSPF Routing Protocol for a User Logical System on page 3362
- *Understanding OSPF Configuration*
- *Verifying an OSPF Configuration*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 150

Configuring User Logical System Routing, Interfaces, and NAT Features

- Understanding Logical System Network Address Translation on page 3355
- Example: Configuring Network Address Translation for a User Logical System on page 3356
- Understanding Logical System Interfaces and Routing Instances on page 3359
- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 3360
- Example: Configuring OSPF Routing Protocol for a User Logical System on page 3362

Understanding Logical System Network Address Translation

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

Any combination of static, destination, or source NAT can be configured in the root or user logical systems. Configuring NAT in a logical system is the same as configuring NAT in a root system. The master administrator can configure and monitor NAT in the master logical system as well as any user logical system.

For each user logical system, the master administrator can configure the maximum and reserved numbers for the following NAT resources:

- Source NAT pools and destination NAT pools
- IP addresses in source NAT pools with and without port address translation
- Rules for source, destination, and static NAT
- Persistent NAT bindings
- IP addresses that support port overloading

From a user logical system, the user logical system administrator can use the operational command **show system security-profile** with a NAT option to view the number of NAT resources allocated to the user logical system.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of NAT resources applied to the master logical system. The number of resources configured in the master logical system count toward the maximum number of NAT resources available on the device.

From a user logical system, the user logical system administrator can use the **show security nat** command to view the information about NAT for the user logical system. From the master logical system, the master administrator can use the same command to view information for the master logical system, a specific user logical system, or all logical systems.

Related Documentation

- Example: Configuring Network Address Translation for a User Logical System on page 3356
- User Logical System Configuration Overview on page 3225
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Introduction to NAT on page 4839

Example: Configuring Network Address Translation for a User Logical System

This example shows how to configure static NAT for a user logical system.

- Requirements on page 3356
- Overview on page 3356
- Configuration on page 3357
- Verification on page 3358

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profile nat-static-rule** command to see the static NAT resources allocated to the logical system.
- Configure security policies. See "Example: Configuring Security Policies in a User Logical System" on page 3308.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

Copyright © 2016, Juniper Networks, Inc.

Devices in the ls-product-design-untrust zone access a specific host in the ls-product-design-trust zone by way of the address 12.1.1.200/32. For packets that enter the ls-product-design logical system from the ls-product-design-untrust zone with the destination IP address 12.1.1.200/32, the destination IP address is translated to the 12.1.1.100/32. This example configures the static NAT described in Table 345.

Table 345: User Logical System Static NAT Configuration

Feature	Name	Configuration Parameters
Static NAT rule set	rs1	• Rule r1 to match packets from the ls-product-design-untrust zone with destination address 12.1.1.200/32. • Destination IP address in matching packets is translated to 12.1.1.100/32.
Proxy ARP		Address 12.1.1.200 on interface lt-0/0/0.3.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set rs1 from zone ls-product-design-untrust
set security nat static rule-set rs1 rule r1 match destination-address 12.1.1.200/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 12.1.1.100/32
set security nat proxy-arp interface lt-0/0/0.3 address 12.1.1.200/32
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure NAT in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a static NAT rule set.

   ```
   [edit security nat static]
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 from zone
     ls-product-design-untrust
   ```

3. Configure a rule that matches packets and translates the destination address in the packets.

   ```
   [edit security nat static]
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 match
     destination-address 12.1.1.200/32
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 then static-nat prefix
     12.1.1.100/32
   ```

4. Configure proxy ARP.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat]
lsdesignadmin1@host:ls-product-design# set proxy-arp interface lt-0/0/0.3 address
  12.1.1.200/32
```

Results From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security nat
static {
  rule-set rs1 {
    from zone ls-product-design-untrust;
    rule r1 {
      match {
        destination-address 12.1.1.200/32;
      }
      then {
        static-nat prefix 12.1.1.100/32;
      }
    }
  }
}
  proxy-arp {
    interface lt-0/0/0.3 {
      address {
        12.1.1.200/32;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static NAT Configuration on page 3358
- Verifying NAT Application to Traffic on page 3358

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- User Logical System Configuration Overview on page 3225
- Understanding Logical System Network Address Translation on page 3355
- Static NAT Configuration Overview on page 4949

Understanding Logical System Interfaces and Routing Instances

Logical interfaces on the device are allocated among the user logical systems by the master administrator. The user logical system administrator configures the attributes of the interfaces, including IP addresses, and assigns them to routing instances and zones.

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Interfaces and routing instances can be configured in the master logical system and in user logical systems. Configuring an interface or routing instance in a logical system is the same as configuring an interface or routing instance on a device that is not configured for logical systems. Any routing instance created within a logical system is only applicable to that logical system.

The default routing instance, master, refers to the main inet.0 routing table in the logical system. The master routing instance is reserved and cannot be specified as a routing instance. Routes are installed in the master routing instance by default, unless a routing instance is specified. Configure global routing options and protocols for the master routing instance by including statements at the [**edit protocols**] and [**edit routing-options**] hierarchy levels in the logical system.

You can configure only virtual router routing instance type in a user logical system. Only one virtual private LAN service (VPLS) routing instance type can be configured on the device and it must be in the interconnect logical system.

The user logical system administrator can configure and view all attributes for an interface or routing instance in a user logical system. All attributes of an interface or routing instance in a user logical system are also visible to the master administrator.

Multicast is a “one source, many destinations” method of traffic distribution, which means the destinations needing to receive the information from a particular source receive the traffic stream. The master and user logical system administrators can configure a logical system to support multicast applications. The same multicast configurations to configure a device as a node in a multicast network can be used in a logical system.

Related Documentation

- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 3360
- User Logical System Configuration Overview on page 3225
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Interfaces and Routing Instances for a User Logical System

This example shows how to configure interfaces and routing instances for a user logical system.

- Requirements on page 3360
- Overview on page 3360
- Configuration on page 3360

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Determine which logical interfaces and, optionally, which logical tunnel interfaces are allocated to your user logical system by the master administrator. The master administrator configures the logical tunnel interfaces. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example configures the interfaces and routing instances described in Table 346.

Table 346: User Logical System Interface and Routing Instance Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/5.1	• IP address 12.1.1.1/24 • VLAN ID 700

Routing instance	pd-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/5.1 and lt-0/0/0.3
		• Static routes:
		  ◦ 13.1.1.0/24 next-hop 10.0.1.3
		  ◦ 14.1.1.0/24 next-hop 10.0.1.4
		  ◦ 12.12.1.0/24 next-hop 10.0.1.1

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/5 unit 1 family inet address 12.1.1.1/24
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/5 unit 1 vlan-id 700
set routing-instances pd-vr1 instance-type virtual-router
set routing-instances pd-vr1 interface ge-0/0/5.1
set routing-instances pd-vr1 interface lt-0/0/0.3
set routing-instances pd-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances pd-vr1 routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
set routing-instances pd-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an interface and a routing instance in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure the logical interface for a user logical system.

   ```
   [edit interfaces]
   lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 family inet address
     12.1.1.1/24
   lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 vlan-id 700
   ```

3. Configure the routing instance and assign interfaces.

   ```
   [edit routing-instances]
   lsdesignadmin1@host:ls-product-design# set pd-vr1 instance-type virtual-router
   lsdesignadmin1@host:ls-product-design# set pd-vr1 interface ge-0/0/5.1
   lsdesignadmin1@host:ls-product-design# set pd-vr1 interface lt-0/0/0.3
   ```

4. Configure static routes.

   ```
   [edit routing-instances]
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
     13.1.1.0/24 next-hop 10.0.1.3
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
     14.1.1.0/24 next-hop 10.0.1.4
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
     12.12.1.0/24 next-hop 10.0.1.1
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

NOTE: The master administrator configures the lt-0/0/0.3 interface. Thus, the lt-0/0/0.3 configuration appears in the show interfaces output even though you did not configure this item.

```
lsdesignadmin1@host:ls-product-design# show interfaces
ge-0/0/5 {
```

Copyright © 2016, Juniper Networks, Inc.

```
    unit 1 {
      vlan-id 700;
      family inet {
        address 12.1.1.1/24;
      }
    }
  }
    lt-0/0/0 {
      unit 3 {
        encapsulation ethernet;
        peer-unit 2;
        family inet {
          address 10.0.1.2/24;
        }
      }
    }
lsdesignadmin1@host:ls-product-design# show routing-instances
pd-vr1 {
  instance-type virtual-router;
  interface ge-0/0/5.1;
  interface lt-0/0/0.3;
  routing-options {
    static {
      route 13.1.1.0/24 next-hop 10.0.1.3;
      route 14.1.1.0/24 next-hop 10.0.1.4;
      route 12.12.1.0/24 next-hop 10.0.1.1;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- User Logical System Configuration Overview on page 3225
- Understanding Logical System Interfaces and Routing Instances on page 3341

Example: Configuring OSPF Routing Protocol for a User Logical System

This example shows how to configure OSPF for a user logical system.

- Requirements on page 3362
- Overview on page 3363
- Configuration on page 3363
- Verification on page 3365

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 3225.

Copyright © 2016, Juniper Networks, Inc.

- Configure logical interface ge-0/0/5.1. Assign ge-0/0/5.1 and lt-0/0/0.3 to the pd-vr1 routing instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 3360.

Overview

In this example, you configure OSPF for the ls-product-design user logical system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.

This example enables OSPF routing on the ge-0/0/5.1 and lt-0/0/0.3 interfaces in the ls-product-design user logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the pd-vr1 routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.
- ospf-redist-static—Static routes.
- ospf-to-ospf—Routes learned from OSPF.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances pd-vr1 protocols ospf export ospf-redist-direct
set routing-instances pd-vr1 protocols ospf export ospf-redist-static
set routing-instances pd-vr1 protocols ospf export ospf-to-ospf
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface ge-0/0/5.1
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface lt-0/0/0.3
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure OSPF for the user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Create routing policies that accept routes.

   ```
   [edit policy-options]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct
  from protocol direct
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct
  then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static
  from protocol static
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static
  then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf from
  protocol ospf
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf then
  accept
```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
  ospf-redist-direct
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
  ospf-redist-static
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
  ospf-to-ospf
```

4. Enable OSPF on the logical interfaces.

```
[edit routing-instances]
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1
  interface ge-0/0/5.1
lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1
  interface lt-0/0/0.3
```

Results From configuration mode, confirm your configuration by entering the **show policy-options** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show policy-options
  policy-statement ospf-redist-direct {
    from protocol direct;
    then accept;
  }
  policy-statement ospf-redist-static {
    from protocol static;
    then accept;
  }
  policy-statement ospf-to-ospf {
    from protocol ospf;
    then accept;
  }
[edit]
lsdesignadmin1@host:ls-product-design# show routing-instances
```


Copyright © 2016, Juniper Networks, Inc.

```
pd-vr1 {
  ...
  protocols {
    ospf {
      export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
      area 0.0.0.1 {
        interface lt-0/0/0.3;
        interface ge-0/0/5.1;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying OSPF Interfaces on page 3365
- Verifying OSPF Neighbors on page 3365
- Verifying OSPF Routes on page 3365

Verifying OSPF Interfaces

Purpose Verify OSPF-enabled interfaces.

Action From the CLI, enter the **show ospf interface instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf interface instance pd-vr1
Interface           State   Area            DR ID           BDR ID          Nbrs
lt-0/0/0.3          DR      0.0.0.0         10.0.1.2        0.0.0.0            0
ge-0/0/5.1          DR      0.0.0.1         10.0.1.2        0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose Verify OSPF neighbors.

Action From the CLI, enter the **show ospf neighbor instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf neighbor instance pd-vr1
Address   Interface  State   ID     Pri Dead
10.0.1.1  plt0.1    Full    0.0.0.0    128   39
```

Verifying OSPF Routes

Purpose Verify OSPF routes.

Action From the CLI, enter the **show ospf route instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf route instance pd-vr1
Topology default Route Table:

Prefix             Path  Route      NH       Metric NextHop       Nexthop
                   Type  Type       Type            Interface     Address/LSP
```

Copyright © 2016, Juniper Networks, Inc.

```
10.0.1.0/24         Intra Network    IP           1 lt-0/0/0.3
12.12.1.0/24        Intra Network    IP           1 ge-0/0/5.1
```

Related Documentation

- Understanding Logical System Interfaces and Routing Instances on page 3341
- Example: Configuring OSPF Routing Protocol for the Master Logical System on page 3350
- *Understanding OSPF Configuration*
- *Verifying an OSPF Configuration*

Copyright © 2016, Juniper Networks, Inc.

PART 48

Configuring Logical Systems in Chassis Cluster

- Configuring Logical Systems When Device is in Chassis Cluster Mode on page 3369

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 151

Configuring Logical Systems When Device is in Chassis Cluster Mode

- Understanding Logical Systems in the Context of Chassis Cluster on page 3369
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 3370
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 3403

Understanding Logical Systems in the Context of Chassis Cluster

The behavior of a chassis cluster whose nodes consist of SRX Series devices running logical systems is the same as that of a cluster whose SRX Series nodes in the cluster are not running logical systems. No difference exists between events that cause a node to fail over. In particular, if a link associated with a single logical system fails, then the device fails over to another node in the cluster.

The master administrator configures the chassis cluster (including both primary and secondary nodes) before he or she creates and configures the logical systems. Each node in the cluster has the same configuration, as is the case for nodes in a cluster not running logical systems. All logical system configurations are synchronized and replicated between both nodes in the cluster.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related Documentation

- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 3370
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 3403
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- *Chassis Cluster Overview*

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only)

This example shows how to configure logical systems in a basic active/passive chassis cluster.

NOTE: The master administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the master administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

- Requirements on page 3370
- Overview on page 3371
- Configuration on page 3374
- Verification on page 3397

Requirements

Before you begin:

- Obtain two high-end SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.
- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line. See *Connecting SRX Series Devices to Create a Chassis Cluster*.
- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*.

NOTE: For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the master administrator. Log in to the device as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively. See "Understanding Licenses for Logical Systems on SRX Series Devices" on page 3209.

Overview

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.
- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The master administrator configures the following logical systems on the primary device (node 0):

- Master logical system—The master administrator configures a security profile to provision portions of the system's security resources to the master logical system and configures the resources of the master logical system.
- User logical systems LSYS1 and LSYS2 and their administrators—The master administrator also configures security profiles to provision portions of the system's security resources to user logical systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.
- Interconnect logical system LSYS0 that connects logical systems on the device—The master administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 and "User Logical System Configuration Overview" on page 3225 for more information about features that can be configured for logical systems.

If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See "Configuring Proxy ARP (CLI Procedure)" on page 4857.

Topology

Figure 135 shows the topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 135: Logical Systems in a Chassis Cluster

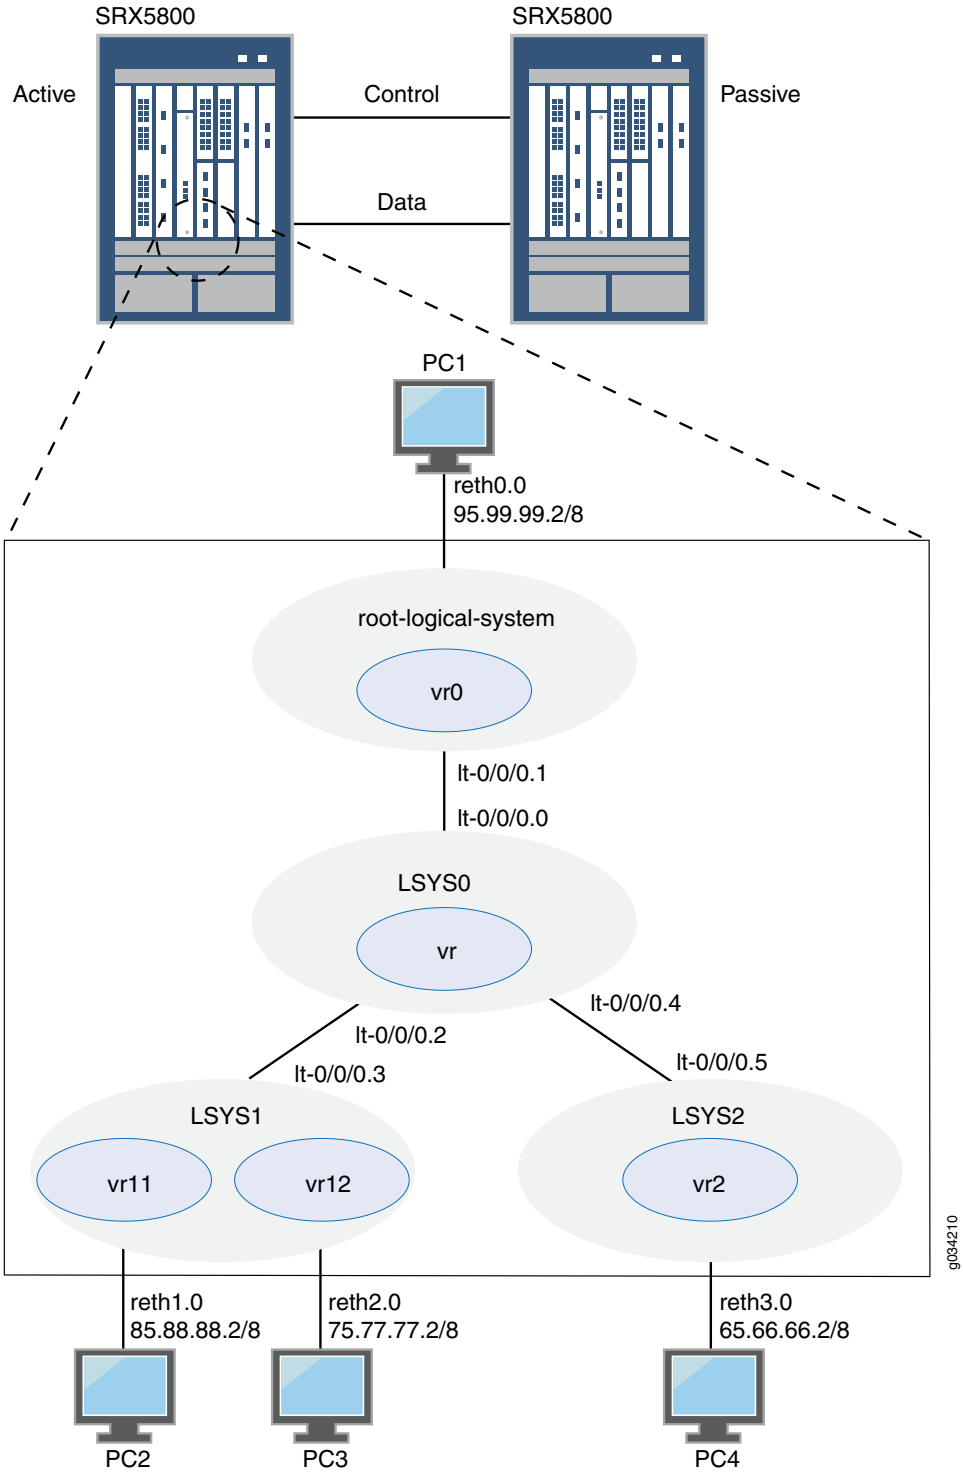

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Chassis Cluster Configuration (Master Administrator) on page 3374
- Logical System Configuration (Master Administrator) on page 3378
- User Logical System Configuration (User Logical System Administrator) on page 3387

Chassis Cluster Configuration (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups “${node}”
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 95.99.99.1/8
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a chassis cluster:

NOTE: Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a commit command.

1. Configure control ports for the clusters.

```
[edit chass cluster]
user@host# set control-ports fpc 0 port 0
user@host# set control-ports fpc 6 port 0
```

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

```
[edit interfaces]
user@host# set fab0 fabric-options member-interfaces ge-1/1/0
user@host# set fab1 fabric-options member-interfaces ge-7/1/0
```

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

```
[edit]
user@host# set groups node0 system host-name SRX5800-1
user@host# set groups node0 interfaces fxp0 unit 0 family inet address
  10.157.90.24/9
user@host# set groups node0 system backup-router 10.157.64.1 destination
  0.0.0.0/0
user@host# set groups node1 system host-name SRX5800-2
user@host# set groups node1 interfaces fxp0 unit 0 family inet address
  10.157.90.23/19
user@host# set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set apply-groups “${node}”
```

4. Configure redundancy groups for chassis clustering.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 0 node 0 priority 200
user@host# set redundancy-group 0 node 1 priority 100
user@host# set redundancy-group 1 node 0 priority 200
user@host# set redundancy-group 1 node 1 priority 100
```

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options redundant-parent reth0
user@host# set ge-1/0/1 gigether-options redundant-parent reth1
user@host# set ge-1/0/2 gigether-options redundant-parent reth2
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 95.99.99.1/8
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

Results From operational mode, confirm your configuration by entering the **show configuration** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
```

 Copyright © 2016, Juniper Networks, Inc.

```
        }
        redundancy-group 0 {
            node 0 priority 200;
            node 1 priority 100;
        }
        redundancy-group 1 {
            node 0 priority 200;
            node 1 priority 100;
        }
    }
}
interfaces {
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
```

```
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet {
                address 95.99.99.1/8;
            }
        }
    }
    reth1 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
}
```

Logical System Configuration (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
```


Copyright © 2016, Juniper Networks, Inc.

```
set system login user lsys2admin class lsys2
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 2.1.1.1/24
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match destination-address any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match application any
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust then permit
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 2.1.1.3/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To create logical systems and user logical system administrators and configure the master and interconnect logical systems:

1. Create the interconnect and user logical systems.

   ```
   [edit logical-systems]
   user@host# set LSYS0
   user@host# set LSYS1
   user@host# set LSYS2
   ```

2. Configure user logical system administrators.

   a. Configure the user logical system administrator for LSYS1.

      ```
      [edit system login]
      user@host# set class lsys1 logical-system LSYS1
      user@host# set class lsys1 permissions all
      user@host# set user lsys1admin full-name lsys1-admin
      user@host# set user lsys1admin class lsys1
      user@host# set user lsys1admin authentication plain-text-password
      ```


Copyright © 2016, Juniper Networks, Inc.

b. Configure the user logical system administrator for LSYS2.

```
[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password
```

3. Configure security profiles and assign them to logical systems.

a. Configure a security profile and assign it to the root logical system.

```
[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system
```

b. Assign a dummy security profile containing no resources to the interconnect logical system LSYS0.

```
[edit system security-profile]
user@host# set SP0 logical-system LSYS0
```

c. Configure a security profile and assign it to LSYS1.

```
[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1
```

d. Configure a security profile and assign it to LSYS2.

```
[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2
```

4. Configure the master logical system.

a. Configure logical tunnel interfaces.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 2.1.1.1/24
```

Copyright © 2016, Juniper Networks, Inc.

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services
  all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services
  all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address
  any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit

[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address
  any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit

[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address
  any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit

[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
user@host# set policy root-Trust_to_root-Trust then permit
```

5. Configure the interconnect logical system.

a. Configure logical tunnel interfaces.

```
[edit logical-systems LSYS0 interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set lt-0/0/0 unit 2 peer-unit 3
user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 4 peer-unit 5
```

b. Configure the VPLS routing instance.

```
[edit logical-systems LSYS0 routing-instances]
user@host# set vr instance-type vpls
user@host# set vr interface lt-0/0/0.0
user@host# set vr interface lt-0/0/0.2
user@host# set vr interface lt-0/0/0.4
```

6. Configure logical tunnel interfaces for the user logical systems.

a. Configure logical tunnel interfaces for LSYS1.

```
[edit logical-systems LSYS1 interfaces ]
user@host# set lt-0/0/0 unit 3 encapsulation ethernet
user@host# set lt-0/0/0 unit 3 peer-unit 2
user@host# set lt-0/0/0 unit 3 family inet address 2.1.1.3/24
```

b. Configure logical tunnel interfaces for LSYS2.

```
[edit logical-systems LSYS2 interfaces ]
user@host# set lt-0/0/0 unit 5 encapsulation ethernet
user@host# set lt-0/0/0 unit 5 peer-unit 4
user@host# set lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Results From configuration mode, confirm the configuration for LSYS0 by entering the **show logical-systems LSYS0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS0
interfaces {
  lt-0/0/0 {
    unit 0 {
      encapsulation ethernet-vpls;
      peer-unit 1;
    }
    unit 2 {
      encapsulation ethernet-vpls;
      peer-unit 3;
    }
    unit 4 {
      encapsulation ethernet-vpls;
      peer-unit 5;
    }
  }
}
  routing-instances {
    vr {
      instance-type vpls;
      interface lt-0/0/0.0;
      interface lt-0/0/0.2;
      interface lt-0/0/0.4;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
```

From configuration mode, confirm the configuration for the master logical system by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet {
      address 2.1.1.1/24;
    }
  }
}
  ge-1/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-1/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-1/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-1/0/3 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  ge-7/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-7/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-7/0/3 {
    gigether-options {
```


Copyright © 2016, Juniper Networks, Inc.

```
        redundant-parent reth3;
      }
    }
    fab0 {
      fabric-options {
        member-interfaces {
          ge-1/1/0;
        }
      }
    }
    fab1 {
      fabric-options {
        member-interfaces {
          ge-7/1/0;
        }
      }
    }
    reth0 {
      redundant-ether-options {
        redundancy-group 1;
      }
      unit 0 {
        family inet {
          address 95.99.99.1/8;
        }
      }
    }
    reth1 {
      redundant-ether-options {
        redundancy-group 1;
      }
    }
    reth2 {
      redundant-ether-options {
        redundancy-group 1;
      }
    }
    reth3 {
      redundant-ether-options {
        redundancy-group 1;
      }
    }
[edit]
user@host# show routing-instances
vr0 {
    instance-type virtual-router;
    interface lt-0/0/0.1;
    interface reth0.0;
    routing-options {
      static {
        route 85.0.0.0/8 next-hop 2.1.1.3;
        route 75.0.0.0/8 next-hop 2.1.1.3;
        route 65.0.0.0/8 next-hop 2.1.1.5;
      }
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security
policies {
    from-zone root-trust to-zone root-untrust {
        policy root-Trust_to_root-Untrust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-untrust to-zone root-trust {
        policy root-Untrust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-untrust to-zone root-untrust {
        policy root-Untrust_to_root-Untrust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone root-trust to-zone root-trust {
        policy root-Trust_to_root-Trust {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
}
    zones {
        security-zone root-trust {
            host-inbound-traffic {
                system-services {
```


Copyright © 2016, Juniper Networks, Inc.

```
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          reth0.0;
        }
      }
      security-zone root-untrust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          lt-0/0/0.1;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

User Logical System Configuration (User Logical System Administrator)

CLI Quick Configuration

To quickly configure user logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet address 85.88.88.1/8
set interfaces reth2 unit 0 family inet address 75.77.77.1/8
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set routing-instances vr11 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet vr11vr12v4
set routing-instances vr12 routing-options static route 85.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 95.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 65.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 2.1.1.0/24 next-table vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr12.inet.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
```

Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  then permit
```

Enter the following commands while logged in as the user logical system administrator for LSYS2:

```
set interfaces reth3 unit 0 family inet address 65.66.66.1/8
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr2 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr2 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match source-address any
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust then permit
```

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The master administrator can also configure a user logical system at the [edit logical-systems *logical-system*] hierarchy level.

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the LSYS1 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys1-admin@host:LSYS1# set reth1 unit 0 family inet address 85.88.88.1/8
lsys1-admin@host:LSYS1# set reth2 unit 0 family inet address 75.77.77.1/8
```

2. Configure routing.

```
[edit routing-instances]
lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
```

Copyright © 2016, Juniper Networks, Inc.

```
lsys1-admin@host:LSYS1# set vr11 interface reth1.0
lsys1-admin@host:LSYS1# set vr11 routing-options static route 65.0.0.0/8 next-hop
  2.1.1.5
lsys1-admin@host:LSYS1# set vr11 routing-options static route 95.0.0.0/8 next-hop
  2.1.1.1
lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr12 interface reth2.0
lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet
  vr11vr12v4
lsys1-admin@host:LSYS1# set vr12 routing-options static route 85.0.0.0/8 next-table
  vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 95.0.0.0/8 next-table
  vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 65.0.0.0/8 next-table
  vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 2.1.1.0/24 next-table
  vr11.inet.0

[edit routing-options]
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr11.inet.0
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr12.inet.0
```

3. Configure zones and security policies.

```
[edit security zones]
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
  system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
  protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
  system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
  protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0

[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
  source-address any
```


Copyright © 2016, Juniper Networks, Inc.

```
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust then permit
```

Step-by-Step Procedure

To configure the LSYS2 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys2-admin@host:LSYS2# set reth3 unit 0 family inet address 65.66.66.1/8
```

2. Configure routing.

```
[edit routing-instances]
lsys2-admin@host:LSYS2# set vr2 instance-type virtual-router
lsys2-admin@host:LSYS2# set vr2 interface lt-0/0/0.5
lsys2-admin@host:LSYS2# set vr2 interface reth3.0
lsys2-admin@host:LSYS2# set vr2 routing-options static route 75.0.0.0/8 next-hop
  2.1.1.3
lsys2-admin@host:LSYS2# set vr2 routing-options static route 85.0.0.0/8 next-hop
  2.1.1.3
lsys2-admin@host:LSYS2# set vr2 routing-options static route 95.0.0.0/8 next-hop
  2.1.1.1
```

3. Configure zones and security policies.

```
[edit security zones]
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic
  system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic
  protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust interfaces reth3.0
lsys2-admin@host:LSYS2# set security zones security-zone lsys2-untrust
  host-inbound-traffic system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust host-inbound-traffic
  protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust interfaces lt-0/0/0.5

[edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust then permit

[edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]
```

Copyright © 2016, Juniper Networks, Inc.

```
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust then permit

[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust then permit

[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match source-address
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust then permit
```

Results From configuration mode, confirm the configuration for LSYS1 by entering the **show interfaces**, **show routing-instances**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 2;
            family inet {
                address 2.1.1.3/24;
            }
        }
    }
    reth1 {
        unit 0 {
            family inet {
                address 85.88.88.1/8;
            }
        }
    }
    reth2 {
        unit 0 {
            family inet {
                address 75.77.77.1/8;
            }
        }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-instances
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.3;
        interface reth1.0;
        routing-options {
            static {
                route 65.0.0.0/8 next-hop 2.1.1.5;
                route 95.0.0.0/8 next-hop 2.1.1.1;
            }
        }
    }
    vr12 {
        instance-type virtual-router;
        interface reth2.0;
        routing-options {
            interface-routes {
                rib-group inet vr11vr12v4;
            }
            static {
                route 85.0.0.0/8 next-table vr11.inet.0;
                route 95.0.0.0/8 next-table vr11.inet.0;
                route 65.0.0.0/8 next-table vr11.inet.0;
                route 2.1.1.0/24 next-table vr11.inet.0;
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-options
rib-groups {
    vr11vr12v4 {
        import-rib [ vr11.inet.0 vr12.inet.0 ];
    }
}
[edit]
lsys1-admin@host:LSYS1# show security
security {
    policies {
        from-zone lsys1-trust to-zone lsys1-untrust {
            policy lsys1trust-to-lsys1untrust {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone lsys1-untrust to-zone lsys1-trust {
```

```
    policy lsys1untrust-to-lsys1trust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone lsys1-untrust to-zone lsys1-untrust {
    policy lsys1untrust-to-lsys1untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone lsys1-trust to-zone lsys1-trust {
    policy lsys1trust-to-lsys1trust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
}
zones {
  security-zone lsys1-trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      reth1.0;
      lt-0/0/0.3;
    }
  }
  security-zone lsys1-untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
            protocols {
              all;
            }
          }
          interfaces {
            reth2.0;
          }
        }
      }
    }
```

From configuration mode, confirm the configuration for LSYS2 by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsys2-admin@host:LSYS2# show interfaces
[edit]
  interfaces {
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 2.1.1.5/24;
      }
    }
  }
  reth3 {
    unit 0 {
      family inet {
        address 65.66.66.1/8;
      }
    }
  }
}
[edit]
lsys2-admin@host:LSYS2# show routing-instances
routing-instances {
  vr2 {
    instance-type virtual-router;
    interface lt-0/0/0.5;
    interface reth3.0;
    routing-options {
      static {
        route 75.0.0.0/8 next-hop 2.1.1.3;
        route 85.0.0.0/8 next-hop 2.1.1.3;
        route 95.0.0.0/8 next-hop 2.1.1.1;
      }
    }
  }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
  policies {
```

Copyright © 2016, Juniper Networks, Inc.

```
    from-zone lsys2-trust to-zone lsys2-untrust {
      policy lsys2trust-to-lsys2untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys2-untrust to-zone lsys2-trust {
      policy lsys2untrust-to-lsys2trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys2-untrust to-zone lsys2-untrust {
      policy lsys2untrust-to-lsys2untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys2-trust to-zone lsys2-trust {
      policy lsys2trust-to-lsys2trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
  }
  zones {
    security-zone lsys2-trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
```


Copyright © 2016, Juniper Networks, Inc.

```
            all;
          }
        }
        interfaces {
          reth3.0;
        }
      }
      security-zone lsys2-untrust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          lt-0/0/0.5;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Chassis Cluster Status on page 3397
- Troubleshooting Chassis Cluster with Logs on page 3398
- Verifying Logical System Licenses on page 3398
- Verifying Logical System License Usage on page 3398
- Verifying Intra-Logical System Traffic on a Logical System on page 3399
- Verifying Intra-Logical System Traffic Within All Logical Systems on page 3399
- Verifying Traffic Between User Logical Systems on page 3400

Verifying Chassis Cluster Status

Purpose Verify the chassis cluster status, failover status, and redundancy group information.

Action From operational mode, enter the **show chassis cluster status** command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                 Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                  200         primary        no       no
    node1                  100         secondary      no       no

Redundancy group: 1 , Failover count: 1
```

```
node0                   200         primary        no       no
node1                   100         secondary      no       no
```

Troubleshooting Chassis Cluster with Logs

Purpose Identify any chassis cluster issues by looking at the logs on both nodes.

Action From operational mode, enter these **show log** commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses

Purpose Verify information about logical system licenses.

Action From operational mode, enter the **show system license status logical-system all** command.

```
{primary:node0}
user@host> show system license status logical-system all
node0:
--------------------------------------------------------------------------
Logical system license status:

logical system name                 license status
root-logical-system                 enabled
LSYS0                               enabled
LSYS1                               enabled
LSYS2                               enabled
```

Verifying Logical System License Usage

Purpose Verify information about logical system license usage.

NOTE: The actual number of licenses used is only displayed on the primary node.

Action From operational mode, enter the **show system license** command.

```
{primary:node0}
user@host> show system license
License usage:
                                Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  logical-system                        4           25           0    permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
  Valid for device: JN110B54BAGB
  Features:
```


Copyright © 2016, Juniper Networks, Inc.

```
    logical-system-25 - Logical System Capacity
      permanent
```

Verifying Intra-Logical System Traffic on a Logical System

Purpose Verify information about currently active security sessions within a logical system.

Action From operational mode, enter the **show security flow session logical-system LSYS1** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1782, Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881

  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
Total sessions: 1

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14388, Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Verifying Intra-Logical System Traffic Within All Logical Systems

Purpose Verify information about currently active security sessions on all logical systems.

Action From operational mode, enter the **show security flow session logical-system all** command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
```

Copyright © 2016, Juniper Networks, Inc.

```
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1776, Valid
Logical system: LSYS1
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881

  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
Total sessions: 1

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14382, Valid
Logical system: LSYS1
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Verifying Traffic Between User Logical Systems

Purpose Verify information about currently active security sessions between logical systems.

Action From operational mode, enter the **show security flow session logical-system *logical-system-name*** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000094, Policy name: root-Untrust_to_root-Trust/5, State: Active,
Timeout: 1768, Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 23, Bytes:
1351
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 22, Bytes: 1880
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000002, Policy name: root-Untrust_to_root-Trust/5, State: Backup,
Timeout: 14384, Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 0, Bytes: 0
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1790, Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes:
2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14398, Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------
```

```
Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000088, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1782, Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 40, Bytes: 2252

  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 32, Bytes:
2114

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1782, Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes:
2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000001, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14382, Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14390, Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Related Documentation

- Understanding Logical Systems in the Context of Chassis Cluster on page 3369
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 3403
- *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*
- *Chassis Cluster Overview*

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only)

This example shows how to configure logical systems in a basic active/passive chassis cluster with IPv6 addresses.

NOTE: The master administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the master administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

- Requirements on page 3403
- Overview on page 3404
- Configuration on page 3407
- Verification on page 3431

Requirements

Before you begin:

- Obtain two high-end SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.
- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line.
- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*.

NOTE: For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the master administrator. Log in to the device as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively. See "Understanding Licenses for Logical Systems on SRX Series Devices" on page 3209.

Overview

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.
- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The master administrator configures the following logical systems on the primary device (node 0):

- Master logical system—The master administrator configures a security profile to provision portions of the system's security resources to the master logical system and configures the resources of the master logical system.
- User logical systems LSYS1 and LSYS2 and their administrators—The master administrator also configures security profiles to provision portions of the system's security resources to user logical systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.
- Interconnect logical system LSYS0 that connects logical systems on the device—The master administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 and "User Logical System Configuration Overview" on page 3225 for more information about features that can be configured for logical systems.

If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See "Configuring Proxy ARP (CLI Procedure)" on page 4857.

Topology

Figure 136 shows the topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 136: Logical Systems in a Chassis Cluster (IPv6)

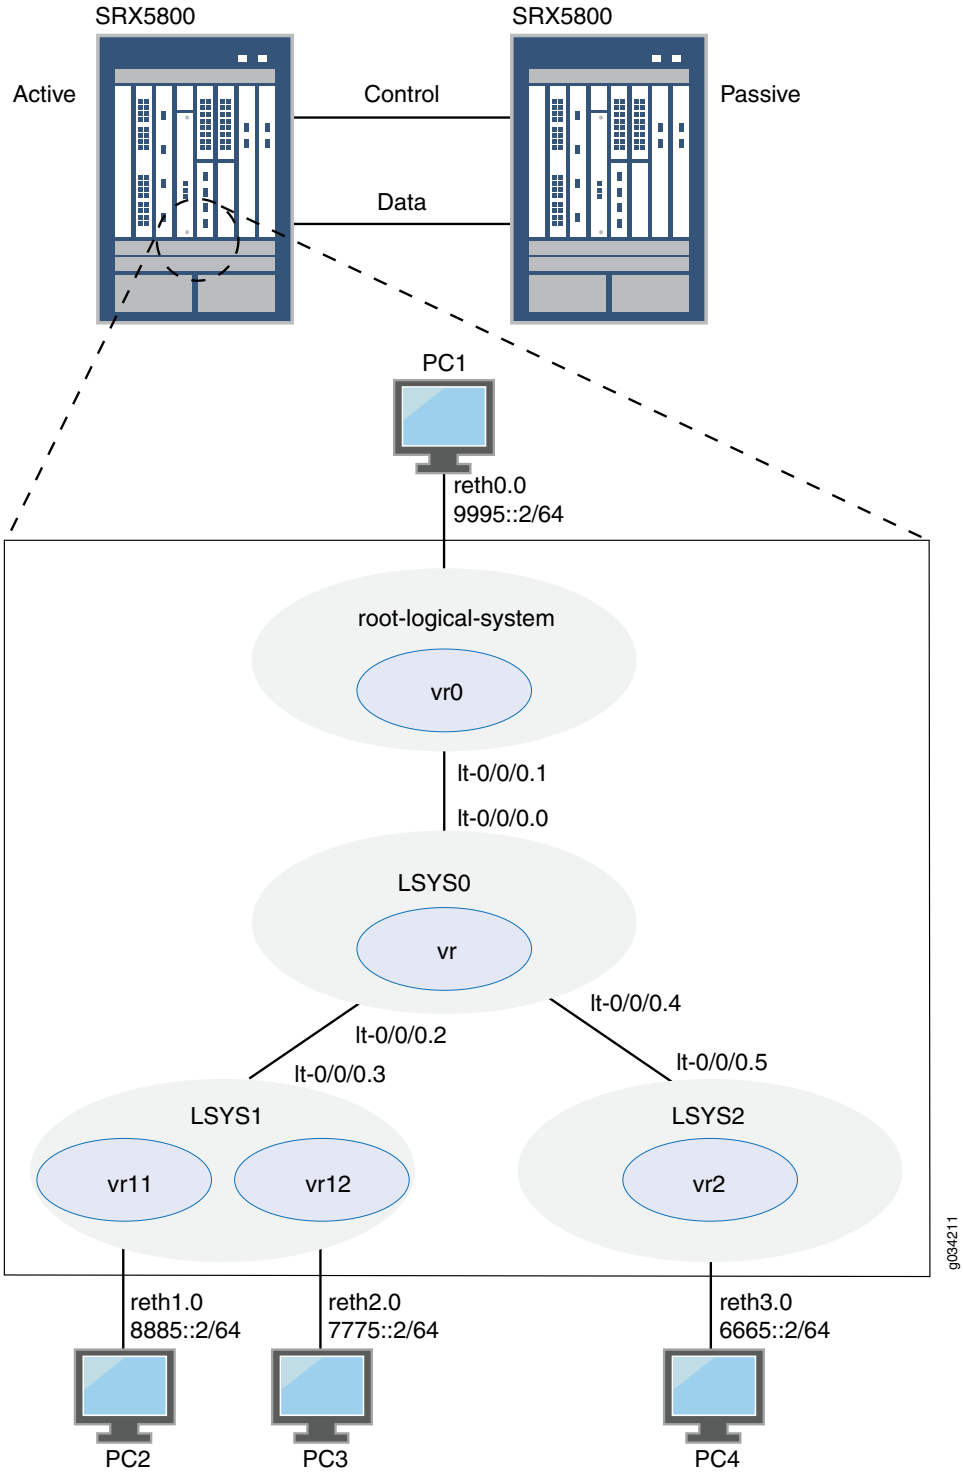

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Chassis Cluster Configuration with IPv6 Addresses (Master Administrator) on page 3407
- Logical System Configuration with IPv6 Addresses (Master Administrator) on page 3411
- User Logical System Configuration with IPv6 (User Logical System Administrator) on page 3420

Chassis Cluster Configuration with IPv6 Addresses (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups “${node}”
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet6 address 9995::1/64
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a chassis cluster:

NOTE: Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a commit command.

1. Configure control ports for the clusters.

```
[edit chassis cluster]
user@host# set control-ports fpc 0 port 0
user@host# set control-ports fpc 6 port 0
```

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

```
[edit interfaces]
user@host# set fab0 fabric-options member-interfaces ge-1/1/0
user@host# set fab1 fabric-options member-interfaces ge-7/1/0
```

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

```
[edit]
user@host# set groups node0 system host-name SRX5800-1
user@host# set groups node0 interfaces fxp0 unit 0 family inet address
  10.157.90.24/9
user@host# set groups node0 system backup-router 10.157.64.1 destination
  0.0.0.0/0
user@host# set groups node1 system host-name SRX5800-2
user@host# set groups node1 interfaces fxp0 unit 0 family inet address
  10.157.90.23/19
user@host# set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
user@host# set apply-groups “${node}”
```

4. Configure redundancy groups for chassis clustering.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 0 node 0 priority 200
user@host# set redundancy-group 0 node 1 priority 100
user@host# set redundancy-group 1 node 0 priority 200
user@host# set redundancy-group 1 node 1 priority 100
```

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options redundant-parent reth0
user@host# set ge-1/0/1 gigether-options redundant-parent reth1
user@host# set ge-1/0/2 gigether-options redundant-parent reth2
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet6 address 9995::1/64
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

Results From operational mode, confirm your configuration by entering the **show configuration** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        redundancy-group 0 {
            node 0 priority 200;
            node 1 priority 100;
        }
        redundancy-group 1 {
            node 0 priority 200;
            node 1 priority 100;
        }
    }
}
interfaces {
    ge-1/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-1/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-1/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-1/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    ge-7/0/0 {
        gigether-options {
            redundant-parent reth0;
        }
    }
    ge-7/0/1 {
        gigether-options {
            redundant-parent reth1;
        }
    }
    ge-7/0/2 {
        gigether-options {
            redundant-parent reth2;
        }
    }
    ge-7/0/3 {
        gigether-options {
            redundant-parent reth3;
        }
    }
    fab0 {
        fabric-options {
            member-interfaces {
                ge-1/1/0;
            }
        }
    }
    fab1 {
        fabric-options {
```


Copyright © 2016, Juniper Networks, Inc.

```
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet6 {
                address 9995::1/64;
            }
        }
    }
    reth1 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
}
```

Logical System Configuration with IPv6 Addresses (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
```

Copyright © 2016, Juniper Networks, Inc.

```
set system login user lsys2admin class lsys2
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet6 address 2111::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 8885::/64 next-hop
  2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop
  2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 6665::/64 next-hop
  2111::5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match destination-address any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy
  root-Trust_to_root-Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust match application any
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone root-untrust to-zone root-trust policy
  root-Untrust_to_root-Trust then permit
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy
  root-Untrust_to_root-Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
  then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 2111::3/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 2111::5/64
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To create logical systems and user logical system administrators and configure the master and interconnect logical systems:

1. Create the interconnect and user logical systems.

   ```
   [edit logical-systems]
   user@host# set LSYS0
   user@host# set LSYS1
   user@host# set LSYS2
   ```

2. Configure user logical system administrators.

   a. Configure the user logical system administrator for LSYS1.

      ```
      [edit system login]
      user@host# set class lsys1 logical-system LSYS1
      user@host# set class lsys1 permissions all
      user@host# set user lsys1admin full-name lsys1-admin
      ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set user lsys1admin class lsys1
user@host# set user lsys1admin authentication plain-text-password
```

b. Configure the user logical system administrator for LSYS2.

```
[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password
```

3. Configure security profiles and assign them to logical systems.

a. Configure a security profile and assign it to the root logical system.

```
[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system
```

b. Assign a dummy security profile containing no resources to the interconnect logical system LSYS0.

```
[edit system security-profile]
user@host# set SP0 logical-system LSYS0
```

c. Configure a security profile and assign it to LSYS1.

```
[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1
```

d. Configure a security profile and assign it to LSYS2.

```
[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2
```

4. Configure the master logical system.

a. Configure logical tunnel interfaces.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet6 address 2111::1/64
```

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options rib vr0.inet6.0 static route 8885::/64
  next-hop 2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop
  2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 6665::/64
  next-hop 2111::5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services
  all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services
  all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address
  any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit

[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address
  any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit

[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address
  any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit

[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy root-Trust_to_root-Trust then permit
```

5. Configure the interconnect logical system.

   a. Configure logical tunnel interfaces.

   ```
   [edit logical-systems LSYS0 interfaces]
   user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 0 peer-unit 1
   user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 2 peer-unit 3
   user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 4 peer-unit 5
   ```

   b. Configure the VPLS routing instance.

   ```
   [edit logical-systems LSYS0 routing-instances]
   user@host# set vr instance-type vpls
   user@host# set vr interface lt-0/0/0.0
   user@host# set vr interface lt-0/0/0.2
   user@host# set vr interface lt-0/0/0.4
   ```

6. Configure logical tunnel interfaces for the user logical systems.

   a. Configure logical tunnel interfaces for LSYS1.

   ```
   [edit logical-systems LSYS1 interfaces ]
   user@host# set lt-0/0/0 unit 3 encapsulation ethernet
   user@host# set lt-0/0/0 unit 3 peer-unit 2
   user@host# set lt-0/0/0 unit 3 family inet6 address 2111::3/64
   ```

   b. Configure logical tunnel interfaces for LSYS2.

   ```
   [edit logical-systems LSYS2 interfaces ]
   user@host# set lt-0/0/0 unit 5 encapsulation ethernet
   user@host# set lt-0/0/0 unit 5 peer-unit 4
   user@host# set lt-0/0/0 unit 5 family inet6 address 2111::5/64
   ```

Results From configuration mode, confirm the configuration for LSYS0 by entering the **show logical-systems LSYS0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS0
interfaces {
  lt-0/0/0 {
    unit 0 {
      encapsulation ethernet-vpls;
      peer-unit 1;
    }
    unit 2 {
      encapsulation ethernet-vpls;
      peer-unit 3;
    }
    unit 4 {
      encapsulation ethernet-vpls;
      peer-unit 5;
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
  routing-instances {
    vr {
      instance-type vpls;
      interface lt-0/0/0.0;
      interface lt-0/0/0.2;
      interface lt-0/0/0.4;
    }
  }
```

From configuration mode, confirm the configuration for the master logical system by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet6 {
      address 2111::1/64;
    }
  }
}
  ge-1/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-1/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-1/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-1/0/3 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  ge-7/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/1 {
    gigether-options {
      redundant-parent reth1;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  ge-7/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-7/0/3 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  fab0 {
    fabric-options {
      member-interfaces {
        ge-1/1/0;
      }
    }
  }
  fab1 {
    fabric-options {
      member-interfaces {
        ge-7/1/0;
      }
    }
  }
  reth0 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet6 {
        address 9995::1/64;
      }
    }
  }
  reth1 {
    redundant-ether-options {
      redundancy-group 1;
    }
  }
  reth2 {
    redundant-ether-options {
      redundancy-group 1;
    }
  }
  reth3 {
    redundant-ether-options {
      redundancy-group 1;
    }
  }
[edit]
user@host# show routing-instances
vr0 {
  instance-type virtual-router;
  interface lt-0/0/0.1;
```


Copyright © 2016, Juniper Networks, Inc.

```
    interface reth0.0;
    routing-options {
      rib vr0.inet6.0 {
        static {
          route 8885::/64 next-hop 2111::3;
          route 7775::/64 next-hop 2111::3;
          route 6665::/64 next-hop 2111::5;
        }
      }
    }
}
[edit]
user@host# show security
policies {
    from-zone root-trust to-zone root-untrust {
      policy root-Trust_to_root-Untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone root-untrust to-zone root-trust {
      policy root-Untrust_to_root-Trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone root-untrust to-zone root-untrust {
      policy root-Untrust_to_root-Untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone root-trust to-zone root-trust {
      policy root-Trust_to_root-Trust {
        match {
          source-address any;
          destination-address any;
          application any;
```

```
        }
        then {
          permit;
        }
      }
    }
  }
    zones {
      security-zone root-trust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          reth0.0;
        }
      }
      security-zone root-untrust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          lt-0/0/0.1;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

User Logical System Configuration with IPv6 (User Logical System Administrator)

CLI Quick Configuration

To quickly configure user logical systems, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet6 address 8885::1/64
set interfaces reth2 unit 0 family inet6 address 7775::1/64
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options rib vr11.inet6.0 static route 6665::/64 next-hop
  2111::5
```


Copyright © 2016, Juniper Networks, Inc.

```
set routing-instances vr11 routing-options rib vr11.inet6.0 static route 9995::/64 next-hop
  2111::1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet6 vr11vr12v6
set routing-instances vr12 rrouting-options rib vr12.inet6.0 static route 8885::/64
  next-table vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 9995::/64 next-table
  vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 6665::/64 next-table
  vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 2111::/64 next-table
  vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr12.inet6.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
  lsys1trust-to-lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
  lsys1untrust-to-lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
  lsys1untrust-to-lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
  then permit
```

Enter the following commands while logged in as the user logical system administrator for LSYS2:

```
set interfaces reth3 unit 0 family inet6 address 6665::1/64
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 7775::/64 next-hop
  2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 8885::/64 next-hop
  2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 9995::/64 next-hop
  2111::1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
  lsys2trust-to-lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
  lsys2untrust-to-lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
  lsys2untrust-to-lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
  lsys2trust-to-lsys2trust then permit
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The master administrator can also configure a user logical system at the [edit logical-systems *logical-system*] hierarchy level.

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the LSYS1 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys1-admin@host:LSYS1# set reth1 unit 0 family inet6 address 8885::1/64
lsys1-admin@host:LSYS1# set reth2 unit 0 family inet6 address 7775::1/64
```

2. Configure routing.

```
[edit routing-instances]
lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
lsys1-admin@host:LSYS1# set vr11 interface reth1.0
lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route
  6665::/64 next-hop 2111::5
lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route
  9995::/64 next-hop 2111::1
lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr12 interface reth2.0
lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet6
  vr11vr12v6
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
  8885::/64 next-table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
  9995::/64 next-table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
  6665::/64 next-table vr11.inet6.0
lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
  2111::/64 next-table vr11.inet6.0

[edit routing-options]
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr11.inet6.0
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr12.inet6.0
```

3. Configure zones and security policies.

```
[edit security zones]
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
  system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
  protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
```

Copyright © 2016, Juniper Networks, Inc.

```
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
  system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
  protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0

[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
  source-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match application
  any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match source-address
  any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match
  destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust then permit
```

Step-by-Step Procedure

To configure the LSYS2 user logical system:

1. Configure interfaces.

```
[edit interfaces]
lsys2-admin@host:LSYS2# set reth3 unit 0 family inet6 address 6665::1/64
```

2. Configure routing.

```
[edit routing-instances]
lsys2-admin@host:LSYS2# set vr2 instance-type virtual-router
lsys2-admin@host:LSYS2# set vr2 interface lt-0/0/0.5
lsys2-admin@host:LSYS2# set vr2 interface reth3.0
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route
  7775::/64 next-hop 2111::3
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route
  8885::/64 next-hop 2111::3
```


Copyright © 2016, Juniper Networks, Inc.

```
lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route
  9995::/64 next-hop 2111::1
```

3. Configure zones and security policies.

```
[edit security zones]
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic
  system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust host-inbound-traffic
  protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-trust interfaces reth3.0
lsys2-admin@host:LSYS2# set security zones security-zone lsys2-untrust
  host-inbound-traffic system-services all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust host-inbound-traffic
  protocols all
lsys2-admin@host:LSYS2# set security-zone lsys2-untrust interfaces lt-0/0/0.5

[edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2untrust then permit

[edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2trust then permit

[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match
  source-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2untrust-to-lsys2untrust then permit

[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match source-address
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match
  destination-address any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust match application
  any
lsys2-admin@host:LSYS2# set policy lsys2trust-to-lsys2trust then permit
```

Results From configuration mode, confirm the configuration for LSYS1 by entering the **show interfaces**, **show routing-instances**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
    lt-0/0/0 {
        unit 3 {
            encapsulation ethernet;
            peer-unit 2;
            family inet6 {
                address 2111::3/64;
            }
        }
    }
    reth1 {
        unit 0 {
            family inet6 {
                address 8885::1/64;
            }
        }
    }
    reth2 {
        unit 0 {
            family inet6 {
                address 7775::1/64;
            }
        }
    }
}
[edit]
lsys1-admin@host:LSYS1# show routing-instances
routing-instances {
    vr11 {
        instance-type virtual-router;
        interface lt-0/0/0.3;
        interface reth1.0;
        routing-options {
            rib vr11.inet6.0 {
                static {
                    route 6665::/64 next-hop 2111::5;
                    route 9995::/64 next-hop 2111::1;
                }
            }
        }
    }
    vr12 {
        instance-type virtual-router;
        interface reth2.0;
        routing-options {
            interface-routes {
                rib-group inet6 vr11vr12v6;
            }
            rib vr12.inet6.0 {
                static {
                    route 8885::/64 next-table vr11.inet6.0;
                    route 9995::/64 next-table vr11.inet6.0;
                    route 6665::/64 next-table vr11.inet6.0;
                    route 2111::/64 next-table vr11.inet6.0;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
}
[edit]
lsys1-admin@host:LSYS1# show routing-options
rib-groups {
  vr11vr12v6 {
    import-rib [ vr11.inet6.0 vr12.inet6.0 ];
  }
}
[edit]
lsys1-admin@host:LSYS1# show security
security {
  policies {
    from-zone lsys1-trust to-zone lsys1-untrust {
      policy lsys1trust-to-lsys1untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-untrust to-zone lsys1-trust {
      policy lsys1untrust-to-lsys1trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-untrust to-zone lsys1-untrust {
      policy lsys1untrust-to-lsys1untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-trust to-zone lsys1-trust {
      policy lsys1trust-to-lsys1trust {
        match {
          source-address any;
```

```
            destination-address any;
            application any;
          }
          then {
            permit;
          }
        }
      }
    }
    zones {
      security-zone lsys1-trust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          reth1.0;
          lt-0/0/0.3;
        }
      }
      security-zone lsys1-untrust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          reth2.0;
        }
      }
    }
  }
```

From configuration mode, confirm the configuration for LSYS2 by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys2-admin@host:LSYS2# show interfaces
interfaces {
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet6 {
        address 2111::5/64;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    reth3 {
      unit 0 {
        family inet6 {
          address 6665::1/64;
        }
      }
    }
}
[edit]
lsys2-admin@host:LSYS2# show routing-instances
routing-instances {
    vr2 {
      instance-type virtual-router;
      interface lt-0/0/0.5;
      interface reth3.0;
      routing-options {
        rib vr2.inet6.0 {
          static {
            route 7775::/64 next-hop 2111::3;
            route 8885::/64 next-hop 2111::3;
            route 9995::/64 next-hop 2111::1;
          }
        }
      }
    }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
    policies {
      from-zone lsys2-trust to-zone lsys2-untrust {
        policy lsys2trust-to-lsys2untrust {
          match {
            source-address any;
            destination-address any;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone lsys2-untrust to-zone lsys2-trust {
        policy lsys2untrust-to-lsys2trust {
          match {
            source-address any;
            destination-address any;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone lsys2-untrust to-zone lsys2-untrust {
```

```
        policy lsys2untrust-to-lsys2untrust {
          match {
            source-address any;
            destination-address any;
            application any;
          }
          then {
            permit;
          }
        }
      }
      from-zone lsys2-trust to-zone lsys2-trust {
        policy lsys2trust-to-lsys2trust {
          match {
            source-address any;
            destination-address any;
            application any;
          }
          then {
            permit;
          }
        }
      }
    }
    zones {
      security-zone lsys2-trust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          reth3.0;
        }
      }
      security-zone lsys2-untrust {
        host-inbound-traffic {
          system-services {
            all;
          }
          protocols {
            all;
          }
        }
        interfaces {
          lt-0/0/0.5;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying Chassis Cluster Status (IPv6) on page 3431
- Troubleshooting Chassis Cluster with Logs (IPv6) on page 3431
- Verifying Logical System Licenses (IPv6) on page 3431
- Verifying Logical System License Usage (IPv6) on page 3432
- Verifying Intra-Logical System Traffic on a Logical System (IPv6) on page 3432
- Verifying Intra-Logical System Traffic Within All Logical Systems (IPv6) on page 3433
- Verifying Traffic Between User Logical Systems (IPv6) on page 3434

Verifying Chassis Cluster Status (IPv6)

Purpose Verify the chassis cluster status, failover status, and redundancy group information.

Action From operational mode, enter the **show chassis cluster status** command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                  Priority          Status    Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no

Redundancy group: 1 , Failover count: 1
    node0                   200         primary        no       no
    node1                   100         secondary      no       no
```

Troubleshooting Chassis Cluster with Logs (IPv6)

Purpose Use these logs to identify any chassis cluster issues. You should run these logs on both nodes.

Action From operational mode, enter these **show log** commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses (IPv6)

Purpose Verify information about logical system licenses.

Action From operational mode, enter the **show system license status logical-system all** command.

```
{primary:node0}
user@host> show system license status logical-system all
```

Copyright © 2016, Juniper Networks, Inc.

```
node0:
--------------------------------------------------------------------------
Logical system license status:

logical system name                license status
root-logical-system                enabled
LSYS0                              enabled
LSYS1                              enabled
LSYS2                              enabled
```

Verifying Logical System License Usage (IPv6)

Purpose Verify information about logical system license usage.

NOTE: The actual number of licenses used is only displayed on the primary node.

Action From operational mode, enter the **show system license** command.

```
{primary:node0}
user@host> show system license
License usage:
                                 Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  logical-system                        4           25           0    permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
  Valid for device: JN110B54BAGB
  Features:
    logical-system-25 - Logical System Capacity
      permanent
```

Verifying Intra-Logical System Traffic on a Logical System (IPv6)

Purpose Verify information about currently active security sessions within a logical system.

Action From operational mode, enter the **show security flow session logical-system LSYS1** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1784, Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
```


Copyright © 2016, Juniper Networks, Inc.

```
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14392, Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Intra-Logical System Traffic Within All Logical Systems (IPv6)

Purpose Verify information about currently active security sessions on all logical systems.

Action From operational mode, enter the **show security flow session logical-system all** command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1776, Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14384, Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Traffic Between User Logical Systems (IPv6)

Purpose Verify information about currently active security sessions between logical systems.

Action From operational mode, enter the **show security flow session logical-system *logical-system-name*** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1792, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14388, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
```


Copyright © 2016, Juniper Networks, Inc.

```
Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1788, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14380, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1784, Valid
Logical system: LSYS1
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701

Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1784, Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
```

```
Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14378, Valid
Logical system: LSYS1
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14376, Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Related Documentation

- Understanding Logical Systems in the Context of Chassis Cluster on page 3369
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 3370
- *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*
- *Chassis Cluster Overview*

Copyright © 2016, Juniper Networks, Inc.

PART 49

Configuring IPv6 for Logical Systems

- Configuring IPv6 Addresses for Logical Systems on page 3439

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 152

Configuring IPv6 Addresses for Logical Systems

- IPv6 Addresses in Logical Systems Overview on page 3439
- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 3440
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 3441
- Example: Configuring IPv6 Zones for a User Logical System on page 3449
- Example: Configuring IPv6 Security Policies for a User Logical System on page 3452
- Example: Configuring IPv6 Dual-Stack Lite for a User Logical System on page 3456

IPv6 Addresses in Logical Systems Overview

IP version 6 (IPv6) increases the size of an IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of its address space. IPv6 has a much larger address space than the soon-to-be exhausted IPv4 address space.

IPv6 addresses can be configured in logical systems for the following features:

- Interfaces
- Firewall authentication
- Flows
- Routing (BGP only)
- Zones and security policies
- Screen options
- Network Address Translation (except for interface NAT)
- Administrative operations such as Telnet, SSH, HTTPS, and other utilities
- Chassis clusters

Copyright © 2016, Juniper Networks, Inc.

NOTE: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

Related Documentation

- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 3441
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 3403
- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 3440

Understanding IPv6 Dual-Stack Lite in Logical Systems

IPv6 dual-stack lite (DS-Lite) allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content. A DS-Lite softwire initiator at the customer edge encapsulates IPv4 packets into IPv6 packets while a softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4 NAT translations.

A specific softwire concentrator and the set of softwire initiators that connect with that softwire concentrator can belong to only one logical system. The master administrator configures the maximum and reserved numbers of softwire initiators that can be connected to a softwire concentrator in a logical system using the **dslite-softwire-initiator** configuration statement at the [**edit system security-profile resources**] hierarchy level. The default maximum value is the system maximum; the default reserved value is 0.

NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of softwire initiators that can connect to a softwire concentrator configured for the master logical system. The number of softwire initiators configured in the master logical system count toward the maximum number of softwire initiators available on the device.

The user logical system administrator can configure softwire concentrators for their user logical system and the master administrator can configure softwire concentrators for the master logical system at the [**edit security softwires**] hierarchy level. The master administrator can also configure softwire concentrators for a user logical system at the [**edit logical-systems *logical-system* security softwires**] hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The softwire concentrator IPv6 address can match an IPv6 address configured on either a physical interface or a loopback interface.

Related Documentation

- Example: Configuring IPv6 Dual-Stack Lite for a User Logical System on page 3456
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Understanding IPv6 Dual-Stack Lite on page 1773

Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only)

This topic covers configuration of IPv6 interfaces, static routes, and routing instances for the master and interconnect logical systems. It also covers configuration of IPv6 logical tunnel interfaces for user logical systems.

- Requirements on page 3441
- Overview on page 3441
- Configuration on page 3443
- Verification on page 3449

Requirements

Before you begin:

- See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 3222 to understand how and where this procedure fits in the overall master administrator configuration process.
- See "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242.
- See "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 3210.

Overview

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, and lt-0/0/0.4. The example configures a routing instance called vr and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a VPLS routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer

Copyright © 2016, Juniper Networks, Inc.

lt-0/0/0 interfaces in the master and user logical systems are configured with Ethernet as the encapsulation type.

- lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.
- lt-0/0/0.2 connects to lt-0/0/0.3 on the LSYS1 logical system.
- lt-0/0/0.4 connects to lt-0/0/0.5 on the LSYS2 logical system.

- For the master logical system, called root-logical-system, the example configures ge-5/0/0 and assigns it to the vr0 routing instance. The example configures lt-0/0/0.1 to connect to lt-0/0/0.0 on the interconnect logical system and assigns it to the vr0 routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr0 routing instance.
- For the LSYS1 logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.
- For the LSYS2 logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.

Figure 137 shows the topology for this deployment including virtual routers and their interfaces for all IPv6 logical systems.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 137: Configuring IPv6 Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

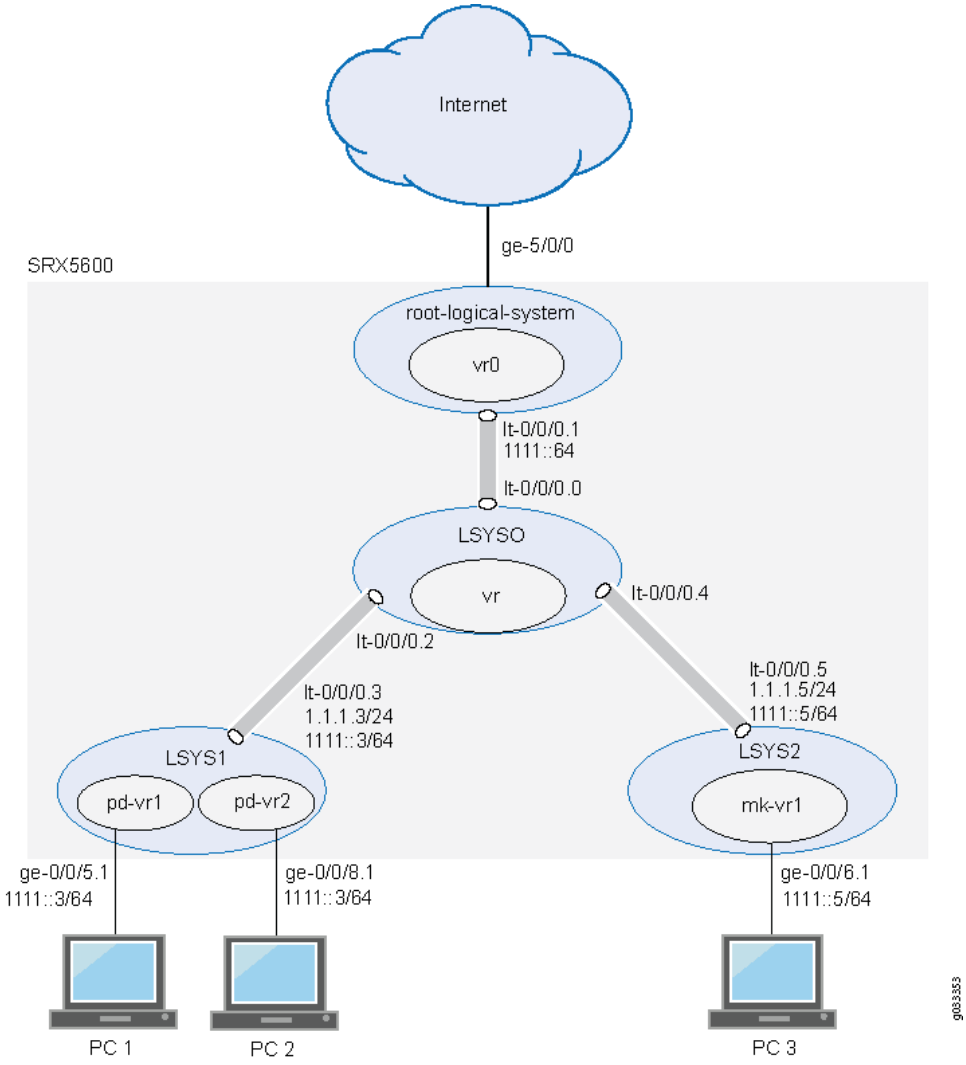

Configuration

This topic explains how to configure interfaces for logical systems.

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System on page 3444
- Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System on page 3445
- Configuring Logical Tunnel Interfaces for the User Logical Systems on page 3447

Copyright © 2016, Juniper Networks, Inc.

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set forwarding-options family inet6 mode flow-based
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Enable flow-based forwarding for IPv6 traffic.

   ```
   [edit security]
   user@host# set forwarding-options family inet6 mode flow-based
   ```

2. Configure the lt-0/0/0 interfaces.

   ```
   [edit logical-systems LSYS0 interfaces]
   user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 0 peer-unit 1
   user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 2 peer-unit 3
   user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
   user@host# set lt-0/0/0 unit 4 peer-unit 5
   ```

3. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

   ```
   [edit logical-systems LSYS0 routing-instances]
   user@host# set vr instance-type vpls
   user@host# set vr interface lt-0/0/0.0
   user@host# set vr interface lt-0/0/0.2
   user@host# set vr interface lt-0/0/0.4
   ```

Results

From configuration mode, confirm your configuration by entering the **show logical-systems interconnect-logical-system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show logical-systems LSYS0
interfaces {
  lt-0/0/0 {
    unit 0 {
      encapsulation ethernet-vpls;
      peer-unit 1;
    }
    unit 2 {
      encapsulation ethernet-vpls;
      peer-unit 3;
      }
      unit 4 {
        encapsulation ethernet-vpls;
        peer-unit 5;
      }
    }
  }
    routing-instances {
      vr {
        instance-type vpls;
        interface lt-0/0/0.0;
        interface lt-0/0/0.2;
        interface lt-0/0/0.4;
      }
    }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-5/0/0 vlan-tagging
set interfaces ge-5/0/0 unit 0 vlan-id 600
set interfaces lt-0/0/0 unit 1 encapsulation Ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 1.1.1.1/24
set interfaces lt-0/0/0 unit 1 family inet6 address 1111::1/64
set interfaces ge-5/0/0 unit 0 family inet address 99.99.99.1/24
set interfaces ge-5/0/0 unit 0 family inet6 address 9999::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface ge-5/0/0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 7777::/64 next-hop
  1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 8888::/64 next-hop
  1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 6666::/64 next-hop
  1111::5
set routing-instances vr0 routing-options static route 77.77.77.0/24 next-hop 1.1.1.3
set routing-instances vr0 routing-options static route 88.88.88.0/24 next-hop 1.1.1.3
set routing-instances vr0 routing-options static route 66.66.66.0/24 next-hop 1.1.1.5
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the master logical system interfaces:

1. Configure the master (root) logical system and lt-0/0/0.1 interfaces.

```
[edit interfaces]
user@host# set ge-5/0/0 vlan-tagging
user@host# set ge-5/0/0 unit 0 vlan-id 600
user@host# set lt-0/0/0 unit 1 encapsulation Ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 1.1.1.1/24
user@host# set lt-0/0/0 unit 1 family inet6 address 1111::1/64
user@host# set ge-5/0/0 unit 0 family inet address 99.99.99.1/24
user@host# set ge-5/0/0 unit 0 family inet6 address 9999::1/64
```

2. Configure a routing instance for the master logical system, assign its interfaces to it, and configure static routes for it.

```
[edit interfaces routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface ge-5/0/0.0
user@host# set vr0 routing-options rib vr0.inet6.0 static route 7777::/64 next-hop
  1111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 8888::/64 next-hop
  1111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 6666::/64 next-hop
  1111::5
user@host# set vr0 routing-options static route 77.77.77.0/24 next-hop 1.1.1.3
user@host# set vr0 routing-options static route 88.88.88.0/24 next-hop 1.1.1.3
user@host# set vr0 routing-options static route 66.66.66.0/24 next-hop 1.1.1.5
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-5/0/0 {
   vlan-tagging;
   unit 0 {
      vlan-id 600;
      family inet {
         address 99.99.99.1/24;
      }
      family inet 6{
         address 9999::1/64;
      }
   }
}
lt-0/0/0 {
   unit 1 {
      encapsulation ethernet;
```


Copyright © 2016, Juniper Networks, Inc.

```
    peer-unit 0;
    family inet {
      address 1.1.1.1/24;
    }
    family inet 6{
      address 1111::1/64;
    }
  }
}
```

```
[edit]
user@host# show routing-instances
vr0 {
  instance-type virtual-router;
  interface ge-5/0/0.0;
  interface lt-0/0/0;
  routing-options {
    rib vr0.inet6.0 {
    static {
      route 8888::/64 next-hop 1111::3;
      route 7777::/64 next-hop 1111::3;
      route 6666::/64 next-hop 1111::5;
    }
    }
    static {
      route 77.77.77.0/24 next-hop 1.1.1.3;
      route 88.88.88.0/24 next-hop 1.1.1.3;
      route 66.66.66.0/24 next-hop 1.1.1.5;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Logical Tunnel Interfaces for the User Logical Systems

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 1.1.1.3/24
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 1111::3/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 1.1.1.5/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 1111::5/64
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Configure the lt-0/0/0 interface for the first user logical system:

Copyright © 2016, Juniper Networks, Inc.

```
[edit logical-systems LSYS1 interfaces lt-0/0/0 unit 3]
user@host# set encapsulation ethernet
user@host# set peer-unit 2
user@host# set family inet address 1.1.1.3/24
user@host# set family inet6 address 1111::3/64
```

2. Configure the lt-0/0/0 interface for the second user logical system.

```
[edit logical-systems LSYS2 interfaces lt-0/0/0 unit 5]
user@host# set encapsulation ethernet
user@host# set peer-unit 4
user@host# set family inet address 1.1.1.5/24
user@host# set family inet6 address 1111::5/64
```

Results From configuration mode, confirm your configuration by entering the **show logical-systems LSYS1 interfaces lt-0/0/0**, and **show logical-systems LSYS2 interfaces lt-0/0/0** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host#  show logical-systems LSYS1 interfaces lt-0/0/0

lt-0/0/0 {
   unit 3 {
      encapsulation ethernet;
      peer-unit 2;
      family inet {
         address 1.1.1.3/24;
      }
      family inet 6{
         address 1111::3/64;
      }
   }
}

user@host#  show logical-systems LSYS2 interfaces lt-0/0/0

lt-0/0/0 {
   unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet {
         address 1.1.1.5/24;
      }
      family inet 6{
         address 1111::5/64;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying That the Static Routes Configured for the Master Administrator Are Correct

Purpose Confirm that the configuration is working properly. Verify if you can send data from the master logical system to the other logical systems.

Action From operational mode, use the **ping** command.

Related Documentation

- Understanding the Master Logical System and the Master Administrator Role on page 3221
- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210
- Example: Configuring IPv6 Zones for a User Logical System on page 3449
- Example: Configuring IPv6 Security Policies for a User Logical System on page 3452

Example: Configuring IPv6 Zones for a User Logical System

This example shows how to configure IPv6 zones for a user logical system.

- Requirements on page 3449
- Overview on page 3450
- Configuration on page 3450

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator.

  See "User Logical System Configuration Overview" on page 3225.

- Ensure that forwarding options for inet6 is flow-based. Otherwise, you must configure it and reset the device.

  Use the **show security forwarding-options** command to check the configuration.

NOTE: Only the user logical system administrator can configure the forwarding options.

Copyright © 2016, Juniper Networks, Inc.

Overview

This example configures the ls-product-design user logical system described in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242

This example creates the IPv6 zones and address books described in Table 347.

Table 347: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 3002::1/96 • Attach to zone ls-product-design-trust
	product-design-external	• Address marketing: 3003::1/24 • Address accounting: 3004::1/24 • Address others: 3002::2/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-system lsys1 security address-book product-design-internal address
  product-designers 3002::1/96
set logical-system lsys1 security address-book product-design-internal attach zone
  ls-product-design-trust
set logical-system lsys1 security address-book product-design-external address marketing
  3003::1/24
set logical-system lsys1 security address-book product-design-external address accounting
  3004::1/24
set logical-system lsys1 security address-book product-design-external address others
  3002::2/24
set logical-system lsys1 security address-book product-design-external address-set
  otherlsys address marketing
set logical-system lsys1 security address-book product-design-external address-set
  otherlsys address accounting
set logical-system lsys1 security address-book product-design-external attach zone
  ls-product-design-untrust
set logical-system lsys1 security zones security-zone ls-product-design-trust tcp-rst
set logical-system lsys1 security zones security-zone ls-product-design-trust interfaces
  ge-0/0/5.1
```


Copyright © 2016, Juniper Networks, Inc.

```
set logical-system lsys1 security zones security-zone ls-product-design-untrust interfaces
    lt-0/0/0.3
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPv6 zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security zone and assign it to an interface.

   ```
   [edit logical-system lsys1 security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust
     interfaces ge-0/0/5.1
   ```

3. Configure the TCP-Reset parameter for the zone.

   ```
   [edit logical-system lsys1 security zones security-zone ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set tcp-rst
   ```

4. Configure a security zone and assign it to an interface.

   ```
   [edit logical-system lsys1 security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust
     interfaces lt-0/0/0.3
   ```

5. Create global address book entries.

   ```
   [edit logical-system lsys1 security]
   lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
     address product-designers 3002::1/96
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address marketing 3003::1/24
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address accounting 3004::1/24
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address others 3002::2/24
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address-set otherlsys address marketing
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
     address-set otherlsys address accounting
   ```

6. Attach address books to zones.

   ```
   [edit logical-system lsys1 security]
   lsdesignadmin1@host:ls-product-design#set address-book product-design-internal
     attach zone ls-product-design-trust
   lsdesignadmin1@host:ls-product-design#set address-book product-design-external
     attach zone ls-product-design-untrust
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security zones** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security zones
address-book {
    product-design-internal {
        address product-designers 3002::1/96;
        attach {
            zone ls-product-design-trust;
        }
    }
    product-design-external {
        address marketing 3003::1/24;
        address accounting 3004::1/24;
        address others 3002::2/24;
        address-set otherlsys {
            address marketing;
            address accounting;
        }
        attach {
            zone ls-product-design-untrust;
        }
    }
}
zones {
    security-zone ls-product-design-trust {
        tcp-rst;
        interfaces {
            ge-0/0/5.1;
        }
    }
    security-zone ls-product-design-untrust {
        interfaces {
            lt-0/0/0.3;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- Understanding Logical System Zones on page 3299
- User Logical System Configuration Overview on page 3225
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 3441
- Example: Configuring IPv6 Security Policies for a User Logical System on page 3452

Example: Configuring IPv6 Security Policies for a User Logical System

This example shows how to configure IPv6 security policies for a user logical system.

- Requirements on page 3453
- Overview on page 3453

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 3453
- Verification on page 3455

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator.

  See "User Logical System Configuration Overview" on page 3225.

- Use the **show system security-profiles policy** command to see the security policy resources allocated to the logical system.
- Configure zones and address books.

  See "Example: Configuring IPv6 Zones for a User Logical System" on page 3449

Overview

This example shows how to configure the security policies described in Table 348.

Table 348: User Logical System Security Policies Configuration

Policy Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: product-designers • Destination address: otherlsys • Application: any
permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-designers • Application: any

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
    ls-product-design-untrust policy permit-all-to-otherlsys match source-address
    product-designers
```

Copyright © 2016, Juniper Networks, Inc.

```
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
  ls-product-design-untrust policy permit-all-to-otherlsys match destination-address
  otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
  ls-product-design-untrust policy permit-all-to-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
  ls-product-design-untrust policy permit-all-to-otherlsys then permit
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
  ls-product-design-trust policy permit-all-from-otherlsys match source-address otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
  ls-product-design-trust policy permit-all-from-otherlsys match destination-address
  product-designers
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
  ls-product-design-trust policy permit-all-from-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
  ls-product-d esign-trust policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPv6 security policies for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

   ```
   [edit logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
     ls-product-design-untrust]
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
     source-address product-designers
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
     destination-address otherlsys
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
     application any
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys then
     permit
   ```

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

   ```
   [edit logical-systems lsys1 security policies from-zone ls-product-design-untrust
     to-zone ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
     source-address otherlsys
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
     destination-address product-designers
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
     application any
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys then
     permit
   ```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy permit-all-from-otherlsys {
        match {
            source-address otherlsys;
            destination-address product-designers;
            application any;
        }
        then {
            permit;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation

- Understanding Logical System Security Policies on page 3306
- User Logical System Configuration Overview on page 3225
- Troubleshooting Security Policies on page 1162
- Example: Configuring IPv6 Zones for a User Logical System on page 3449
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 3441

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring IPv6 Dual-Stack Lite for a User Logical System

This example shows how to configure a softwire concentrator for a user logical system.

- Requirements on page 3456
- Overview on page 3456
- Configuration on page 3456
- Verification on page 3457

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 3225.
- Use the **show system security-profile dslite-softwire-initiator** command to see the number softwire initiators that can be connected to a softwire concentrator in the logical system.

Overview

This example shows how to configure a softwire concentrator to decapsulate IPv4-in-IPv6 packets in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242. The IPv6 address of the softwire concentrator is 3000::1 and the name of the softwire configuration is sc_1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security softwires softwire-name sc_1 softwire-concentrator 3000::1 softwire-type
    IPv4-in-IPv6
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an IPv6 DS-Lite softwire concentrator:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Specify the address of the softwire concentrator and the softwire type.

   ```
   [edit security]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
lsdesignadmin1@host:ls-product-design# set softwires softwire-name sc_1
  softwire-concentrator 3000::1 softwire-type IPv4-in-IPv6
```

Results From configuration mode, confirm your configuration by entering the **show security softwires** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsdesignadmin1@host:ls-product-design# show security softwires
softwire-name sc_1 {
   softwire-concentrator 3000::1;
   softwire-type IPv4-in-IPv6;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the DS-Lite Configuration

Purpose Verify that the softwire initiators can connect to the softwire concentrator configured in the user logical system.

Action From operational mode, enter the **show security softwires** command.

If a softwire initiator is not connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Active         0
```

If a softwire initiator is connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  sc_1                      3000::1         Connected      1
```

Related Documentation

- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 3440
- User Logical System Configuration Overview on page 3225

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 50

Configuring System Resources Allocation

- System Resources Allocation (Master Administrators Only) on page 3461

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 153

System Resources Allocation (Master Administrators Only)

- Understanding CPU Allocation and Control on page 3461
- Example: Configuring CPU Utilization (Master Administrators Only) on page 3465
- Example: Deleting an SRX Series Services Gateway Logical System (Master Administrators Only) on page 3468

Understanding CPU Allocation and Control

When device CPU utilization is low, logical systems can acquire and use CPU resources above their allocated reserve quotas as long as the system-wide utilization remains within a stable range. CPU utilization on a device should never reach 100 percent because a device running at 100 percent CPU utilization might be slow to respond to management or system events or be unable to handle traffic bursts.

CPU resources are used on a first-come first-served basis. Without controls, logical systems can compete for CPU resources and drive CPU utilization up to 100 percent. You cannot rely on the configuration of static resources, such as security policies and zones, to directly control CPU usage because a logical system with small numbers of static resources allocated could still consume a large amount of CPU. Instead, the master administrator can enable CPU resource control and configure CPU utilization parameters for logical systems.

NOTE: **Only the master administrator can enable CPU control and configure CPU utilization parameters. User logical system administrators can use the show system security-profile cpu command to view CPU utilization for their logical systems.**

This topic includes the following sections:

- CPU Control on page 3462
- Reserved CPU Utilization Quota for Logical Systems on page 3462
- CPU Control Target on page 3463

Copyright © 2016, Juniper Networks, Inc.

- Shared CPU Resources and CPU Quotas on page 3463
- Monitoring CPU Utilization on page 3465

CPU Control

The master administrator enables CPU control with the **cpu-control** configuration statement at the [**edit system security-profile resources**] hierarchy level.

NOTE: The resources security profile is a special security profile that contains global settings that apply to all logical systems in the device. Other security profiles configured by the master administrator are bound to specific logical systems.

When CPU control is enabled, the master administrator can then configure the following CPU utilization parameters:

- A reserved CPU quota is the percentage of CPU utilization that is guaranteed for a logical system.
- The CPU control target is the upper limit, in percent, for system-wide CPU utilization on the device under normal operating conditions.

Reserved CPU Utilization Quota for Logical Systems

A configured reserved CPU quota guarantees that a specified percentage of CPU is always available to a logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The reserved CPU quota is used to calculate the amount of CPU each logical system can use based on the runtime utilization.

The master administrator specifies the reserved CPU quota in a logical system security profile with the **cpu reserved** configuration statement at the [**edit system security-profile** *profile-name*] hierarchy level. The security profile is bound to one or more logical systems. Unlike other resources that are allocated to a logical system in a security profile, no maximum allowed quota can be configured for CPU utilization.

The Junos OS software checks to ensure that the sum of reserved CPU quotas for all logical systems on the device is less than 90 percent of the CPU control target value. If CPU control is enabled and reserved CPU quotas are not configured, the default reserved CPU quota for the master logical system is 1 percent and the default reserved CPU quota for user logical systems is 0 percent. The master administrator can configure reserved CPU quotas even if CPU control is not enabled. The master administrator can enable or disable CPU control without changing security profiles.

CAUTION: The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota because traffic loss could occur.

Copyright © 2016, Juniper Networks, Inc.

CPU Control Target

CPU control target is the upper limit, in percent, for CPU utilization on the device under normal operating conditions. If CPU utilization on the device surpasses the configured target value, the Junos OS software initiates controls to bring CPU utilization between the target value and 90 percent of the target value. For example, if the CPU control target value is 80 and CPU utilization on the device surpasses 80 percent, then controls are initiated to bring CPU utilization within the range of 72 (90 percent of 80) and 80 percent.

During runtime, CPU utilization by each logical system is measured every two seconds. Dropping packets reduces the CPU usage for a logical system. If the CPU usage of a logical system exceeds its quota, CPU utilization control drops the packets received on that logical system. The packet drop rate is calculated every two seconds based on CPU utilization of all logical systems.

The master administrator configures the CPU control target with the **cpu-control-target** configuration statement at the [**edit system security-profile resources**] hierarchy level. A stable level of CPU utilization should be relatively close to 100 percent but allow for bursts in CPU utilization. The master administrator should configure the CPU control target level based on an understanding of the usage pattern of the logical system's deployment on the device.

CPU control must be enabled for the Junos OS software to control CPU usage. If the master administrator enables CPU control without specifying a CPU control target value, the default CPU control target is 80 percent.

Shared CPU Resources and CPU Quotas

The sum of the reserved CPU quotas for all logical systems on the device must be less than 90 percent of the CPU control target; the difference is called the shared CPU resource. The shared CPU resource is dynamically allocated among the logical systems that need additional CPU. This means that a logical system can use more CPU than its reserved CPU quota.

The CPU quota for a logical system is the sum of its reserved CPU quota and its portion of the shared CPU resource. If multiple logical systems need more CPU resources, they split the shared CPU resource based on the relative weights of their reserved CPU quotas. Logical systems with larger reserved CPU quotas receive larger portions of the shared CPU resource. The goal for CPU control is to keep the actual CPU utilization of a logical system at its CPU quota. If a logical system's CPU needs are greater than its CPU quota, packets are dropped for that logical system.

The following scenarios illustrate CPU control for logical systems. In each scenario, the CPU control target value is 80, which means that CPU controls will keep the maximum system-wide CPU utilization between 72 and 80 percent. The reserved CPU quotas for the logical systems are configured as follows: master and lsys1 logical systems are 10 percent each and the lsys2 logical system is 5 percent.

Copyright © 2016, Juniper Networks, Inc.

CPU Utilization Scenario 1

In this scenario, each of the three logical systems needs 40 percent of CPU. Table 349 shows the CPU quotas for each logical system. Because the CPU needed by each logical system is greater than its CPU quota, packets are dropped for each logical system.

Table 349: CPU Utilization Scenario 1

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
master	40%	28.8%	Yes
lsys1	40%	28.8%	Yes
lsys2	40%	14.4%	Yes

CPU Utilization Scenario 2

In this scenario, the master logical system needs 25 percent of CPU while the two user logical systems need 40 percent. Table 350 shows the CPU quota for the master logical system is equal to the CPU it needs, so no packets are dropped for the master logical system and CPU control monitors the CPU utilization of the master logical system. Packets are dropped for lsys1 and lsys2.

Table 350: CPU Utilization Scenario 2

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
master	25%	25%	No
lsys1	40%	31.3%	Yes
lsys2	40%	15.6%	Yes

CPU Utilization Scenario 3

In this scenario, the master and lsys2 logical systems need 5 percent and 3 percent of CPU, respectively, while lsys1 needs 40 percent. Table 351 shows system-wide CPU utilization is 48 percent, which is less than 72 percent (90 percent of the CPU control target), so no packets are dropped and CPU control monitors all logical systems.

Table 351: CPU Utilization Scenario 3

Logical System	Needed CPU	CPU Quota	Packets Dropped?
master	5%	5%	No
lsys1	40%	40%	No
lsys2	3%	3%	No

Copyright © 2016, Juniper Networks, Inc.

Monitoring CPU Utilization

CPU utilization can be monitored by either the master administrator or the user logical system administrators. The master administrator can monitor CPU utilization for the master logical system, a specified user logical system, or all logical systems. User logical system administrators can only monitor CPU utilization for their logical system.

The **show system security-profile cpu** command shows the usage and drop rate in addition to the reserved CPU quota configured for the logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The usage and drop rates displayed are the values at the interval prior to when the **show** command is run. If the **detail** option is not specified, the utilization of the central point (CP) and the average utilization of all services processing units (SPUs) is shown. The **detail** option displays the CPU utilization on each SPU.

The CPU utilization log file **lsys-cpu-utilization-log** contains utilization data for all logical systems on the device. Only the master administrator can view the log file with the **show log lsys-cpu-utilization-log** command.

Related Documentation

- Example: Configuring CPU Utilization (Master Administrators Only) on page 3465
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253

Example: Configuring CPU Utilization (Master Administrators Only)

The master administrator can enable CPU control and configure CPU utilization parameters. This example shows how to enable CPU utilization control and configure CPU utilization quotas and a control target.

- Requirements on page 3465
- Overview on page 3465
- Configuration on page 3466
- Verification on page 3467

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 3221.
- Bind security profiles to the master logical system and user logical systems configured on the device. See "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 3258.

Overview

In this example, you enable CPU control and set the CPU control target to be 85 percent. You allocate reserved CPU quotas to the logical systems shown in "Example: Creating

Copyright © 2016, Juniper Networks, Inc.

User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 3242. The logical systems are bound to the security profiles shown in Table 352 and are assigned the reserved CPU quotas in the security profiles.

Table 352: Logical Systems, Security Profiles, and Reserved CPU Quotas

Logical System	Security Profile	Reserved CPU Quotas
root-logical-system (master)	master-profile	2 percent
ls-product-design	ls-design-profile	2 percent
ls-marketing-dept, ls-accounting-dept	ls-accnt-mrkt-profile	1 percent

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system security-profile resources cpu-control
set system security-profile resources cpu-control-target 85
set system security-profile master-profile cpu reserved 2
set system security-profile ls-design-profile cpu reserved 2
set system security-profile ls-accnt-mrkt-profile cpu reserved 1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure CPU utilization control parameters:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   [edit]
   admin@host> configure
   admin@host#
   ```

2. Enable CPU control.

   ```
   [edit system security-profile resources]
   admin@host# set cpu-control
   ```

3. Configure the CPU control target.

   ```
   [edit system security-profile resources]
   admin@host# set cpu-control-target 85
   ```

4. Configure the reserved CPU quotas in the security profiles.

   ```
   [edit system]
   admin@host# set security-profile security-profile master-profile cpu reserved 2
   admin@host# set security-profile security-profile ls-design-profile cpu reserved 2
   ```


Copyright © 2016, Juniper Networks, Inc.

```
admin@host# set security-profile security-profile ls-accnt-mrkt-profile cpu reserved
  1
```

Results From configuration mode, confirm your configuration by entering the **show system security-profile** command. If the output does not display the intended configuration, repeat the \ instructions in this example to correct the configuration.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show system security-profile
  resources {
    cpu-control;
    cpu-control-target 85;
  }
    ls-accnt-mrkt-profile {
      ...
      cpu {
        reserved 1;
      }
      logical-system [ ls-marketing-dept ls-accounting-dept ];
    }
    ls-design-profile {
      ...
      cpu {
        reserved 2;
      }
      logical-system ls-product-design;
    }
    master-profile {
      ...
      cpu {
        reserved 2;
      }
      logical-system root-logical-system;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying CPU Utilization on page 3467

Verifying CPU Utilization

Purpose Display the configured reserved CPU quota, the actual CPU usage, and the drop rate.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show system security-profile cpu logical-system all** command.

```
admin@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 85.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    master-profile CP              0.10%          2.00%
     0.00%
root-logical-system    master-Profile SPU             0.25%          2.00%
     0.00%
ls-product-design      ls-design-profile CP           0.53%          2.00%
     0.00%
ls-product-design      ls-design-profile SPU          0.26%          2.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP        0.10%          1.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU       0.15%          1.00%
     0.00%
ls-accounting-dept     ls-acct-mrkt-profile CP        0.23%          1.00%
     0.00%
ls-accounting-dept     ls-acct-mrkt-profile SPU       0.34%          1.00%
     0.00%
```

Related Documentation

- Understanding CPU Allocation and Control on page 3461
- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253

Example: Deleting an SRX Series Services Gateway Logical System (Master Administrators Only)

This example shows how to delete a logical system configured for an SRX Series Services Gateway device running logical systems. Only the master administrator can delete a logical system.

- Requirements on page 3468
- Overview on page 3468
- Configuration on page 3469
- Verification on page 3471

Requirements

The example uses an SRX5600 device running Junos OS with Logical Systems.

Alternatively, follow those instructions substituting your own configuration values.

Overview

This example shows how to delete a logical system, which you can do at any time. However, if you have configured the device to include the maximum number of logical systems that are supported you must first delete an existing logical system before you can add another one.

Copyright © 2016, Juniper Networks, Inc.

Deletion of a logical system is a simple procedure that includes these tasks:

- Remove from the logical system the security profile that is bound to it.

  Note that in this step you are not deleting the security profile—it might be used for other logical systems—but simply detaching it from the logical system that you intend to delete.

- Detach from the logical system any login classes that are associated with it.

  Removing them from the logical system does not delete the login classes.

- Delete the logical system.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
delete system security-profile ls-design-profile logical-system ls-product-design
delete system login class ls-design-admin logical-system ls-product-design
delete system login class ls-design-user logical-system ls-product-design
delete logical-system ls-product-design
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To delete a logical system:

1. Determine that the logical system that you want to delete exists.

   ```
   [edit]
   user@host# show logical-systems ?
   interconnect-logical-system Logical system name
   ls-accounting-dept Logical system name
   ls-marketing-dept Logical system name
   ls-product-design Logical system name
   ```

2. Delete the security profile.

   a. Verify that security profile that you intend to detach from the logical system is bound to it.

      ```
      [edit]
      user@host# show system security-profile ls-design-profile
      logical-system [ ls-product-design ];
      ```

   b. Detach the security profile from the logical system.

      ```
      [edit]
      ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# delete system security-profile ls-design-profile logical-system
  ls-product-design
```

4. Delete the login classes.

   a. Display the login class and login user configurations for the user logical system administrator.

```
user@host> show configuration system login class ls-design-admin
logical-system ls-product-design;
permissions all;
user@host> show configuration system login user lsdesignadmin1
full-name lsdesignadmin1;
uid 2006;
class ls-design-admin;
authentication {
  encrypted-password "$ABC123"; ## SECRET-DATA
  }
```

   b. Detach the login class for the administrator from the logical system.

```
[edit]
user@host# delete system login class ls-design-admin logical-system
  ls-product-design
```

   c. Display the login class and login user configurations for the user.

```
user@host> show configuration system login class ls-design-user
logical-system ls-product-design;
permissions view;
user@host> show configuration system login user lsdesignuser1
full-name lsdesignuser1
uid 2007;
class ls-design-user;
authentication {
  encrypted-password "$ABC123"; ## SECRET-DATA
  }
```

   d. Detach the login class for the user from the logical system.

```
user@host# delete system login class ls-design-user logical-system
  ls-product-design
```

5. Delete the logical system.

```
[edit]
user@host# delete logical-system ls-product-design
```

Results From configuration mode, confirm your configuration by entering the **show logical-systems** command. In this case, the logical system that you deleted should not be included in displayed list of logical systems configured for the device. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show logical-systems
interconnect-logical-system Logical system name
ls-accounting-dept Logical system name
```


Copyright © 2016, Juniper Networks, Inc.

interconnect-logical-system Logical system name
ls-marketing-dept Logical system name

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted on page 3471

Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted

Purpose Verify if the logical system has been deleted using the show command described previously.

Related Documentation

- Understanding User Logical Systems and the User Logical System Administrator Role on page 3227
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 51

Troubleshooting

- Troubleshooting Logical Systems (Master Administrators Only) on page 3475

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 154

Troubleshooting Logical Systems (Master Administrators Only)

- Understanding Security Logs and Logical Systems on page 3475
- Understanding Data Path Debugging for Logical Systems on page 3476
- Performing Tracing for Logical Systems (Master Administrators Only) on page 3477
- Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only) on page 3481

Understanding Security Logs and Logical Systems

Security logs are system log messages that include security events. If a device is configured for logical systems, security logs generated within the context of a logical system use the name ***logname*_LS** (for example, **IDP_ATTACK_LOG_EVENT_LS**). The logical system version of a log has the same set of attributes as the log for devices that are not configured for logical systems, but it also includes logical-system-name as the first attribute.

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT log for a device that is *not* configured for logical systems:

```
IDP_ATTACK_LOG_EVENT {
  help "IDP attack log";
  description "IDP Attack log generated for attack";
  type event;
  args timestamp message-type source-address source-port destination-address
destination-port protocol-name service-name application-name rule-name
rulebase-name policy-name repeat-count action threat-severity attack-name
nat-source-address nat-source-port nat-destination-address nat-destination-port
elapsed-time inbound-bytes outbound-bytes inbound-packets outbound-packets
source-zone-name source-interface-name destination-zone-name
destination-interface-name packet-log-id message;
  severity LOG_INFO;
  flag auditable;
  edit "2010/10/01 mvr created";
}
```

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT_LS log for a device that is configured for logical systems (note that logical-system-name is the first attribute):

Copyright © 2016, Juniper Networks, Inc.

```
IDP_ATTACK_LOG_EVENT_LS {
  help "IDP attack log";
  description "IDP Attack log generated for attack";
  type event;
  args logical-system-name timestamp message-type source-address source-port
destination-address destination-port protocol-name service-name application-name
rule-name rulebase-name policy-name repeat-count action threat-severity attack-name
 nat-source-address nat-source-port nat-destination-address nat-destination-port
elapsed-time inbound-bytes outbound-bytes inbound-packets outbound-packets
source-zone-name source-interface-name destination-zone-name
destination-interface-name packet-log-id message;
  severity LOG_INFO;
  flag auditable;
  edit "2010/10/01 mvr created";
}
```

If a device is configured for logical systems, log parsing scripts might need to be modified because the log name includes the **_LS** suffix and the logical-system-name attribute can be used to segregate logs by logical system.

If a device is not configured for logical systems, the security logs remain unchanged and scripts built to parse logs do not need any modification.

NOTE: Only the master administrator can configure logging at the [**edit security log**] hierarchy level. User logical system administrators cannot configure logging for their logical systems.

Related Documentation

- *Understanding System Logging for Security Devices*

Understanding Data Path Debugging for Logical Systems

Data path debugging provides tracing and debugging at multiple processing units along the packet-processing path. Data path debugging can also be performed on traffic between logical systems.

NOTE: Only the master administrator can configure data path debugging for logical systems at the [**edit security datapath-debug**] level. User logical system administrators cannot configure data path debugging for their logical systems.

End-to-end event tracing traces the path of a packet from when it enters the device to when it leaves the device. When the master administrator configures end-to-end event tracing, the trace output contains logical system information.

The master administrator can also configure tracing for traffic between logical systems. The trace output shows traffic entering and leaving the logical tunnel between logical systems. When the **preserve-trace-order** option is configured, the trace message is sorted chronologically. In addition to the trace action, other actions such as packet-dump and packet-summary may be configured for traffic between logical systems.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Performing Tracing for Logical Systems (Master Administrators Only) on page 3477
- *Understanding Data Path Debugging for SRX Series Devices*

Performing Tracing for Logical Systems (Master Administrators Only)

NOTE: Only the master administrator can configure data path debugging for logical systems at the root level.

To configure an action profile for a trace or packet capture:

1. Specify event types and trace actions. You can specify any combination of event types and trace actions. For example, the following statements configure multiple trace actions for each event type:

```
[edit security datapath-debug]
user@host# set action-profile p1 event lbt trace
user@host# set action-profile p1 event lbt count
user@host# set action-profile p1 event lbt packet-summary
user@host# set action-profile p1 event lbt packet-dump
user@host# set action-profile p1 event pot trace
user@host# set action-profile p1 event pot count
user@host# set action-profile p1 event pot packet-summary
user@host# set action-profile p1 event pot packet-dump
user@host# set action-profile p1 event np-ingress trace
user@host# set action-profile p1 event np-ingress count
user@host# set action-profile p1 event np-ingress packet-summary
user@host# set action-profile p1 event np-ingress packet-dump
user@host# set action-profile p1 event np-egress trace
user@host# set action-profile p1 event np-egress count
user@host# set action-profile p1 event np-egress packet-summary
user@host# set action-profile p1 event np-egress packet-dump
user@host# set action-profile p1 event jexec trace
user@host# set action-profile p1 event jexec count
user@host# set action-profile p1 event jexec packet-summary
user@host# set action-profile p1 event jexec packet-dump
user@host# set action-profile p1 event lt-enter trace
user@host# set action-profile p1 event lt-enter count
user@host# set action-profile p1 event lt-enter packet-summary
user@host# set action-profile p1 event lt-enter packet-dump
user@host# set action-profile p1 event lt-leave trace
user@host# set action-profile p1 event lt-leave count
user@host# set action-profile p1 event lt-leave packet-summary
user@host# set action-profile p1 event lt-leave packet-dump
```

2. Specify action profile options.

```
[edit security datapath-debug]
user@host# set action-profile p1 record-pic-history
user@host# set action-profile p1 preserve-trace-order
```

3. Configure packet filter options.

```
[edit security datapath-debug]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set packet-filter 1 action-profile p1
user@host# set packet-filter 1 protocol udp
```

To capture trace messages for logical systems:

1. Configure the trace capture file.

```
[edit security datapath-debug]
user@host# set traceoptions file e2e.trace
user@host# set traceoptions file size 10m
```

2. Display the captured trace in operational mode.

```
user@host> show log e2e.trace
  Jul  7 09:49:56
  09:49:56.417578:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:0
  PIC History: ->C0/F1/P0
  NP ingress channel 0 packet
  Meta: Src: F1/P0 Dst: F0/P0
  IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

  Jul  7 09:49:56
  09:49:55.1414031:CID-00:FPC-00:PIC-00:THREAD_ID-04:FINDEX:0:IIF:75:SEQ:0:TC:1
  PIC History: ->C0/F1/P0->C0/F0/P0
  LBT pkt, payload: DATA
  Meta: Src: F1/P0 Dst: F0/P0
  IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

  ...
  (Some trace information omitted)
  ...

  .Jul  7 09:49:56
  09:49:55.1415649:CID-00:FPC-00:PIC-00:THREAD_ID-05:FINDEX:0:IIF:75:SEQ:0:TC:16
  PIC History: ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0
  POT pkt, action: POT_SEND payload: DATA
  Meta: Src: F0/P0 Dst: F1/P0
  IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

  Jul  7 09:49:56
  09:49:56.419274:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:17
  PIC History:
  ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F1/P0
  NP egress channel 0 packet
  Meta: Src: F0/P0 Dst: F1/P0
  IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500
```

3. Clear the log.

```
user@host> clear log e2e.trace
```

To perform packet capture for logical systems:

1. Configure the packet capture file.

```
[edit security datapath-debug]
user@host# set capture-file e2e.pcap
user@host# set capture-file format pcap
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set capture-file size 10m
user@host# set capture-file world-readable
user@host# set capture-file maximum-capture-size 1500
```

2. Enter operational mode to start and then stop the packet capture.

```
user@host> request security datapath-debug capture start
user@host> request security datapath-debug capture stop
```

NOTE: Packet capture files can be opened and analyzed offline with tcpdump or any packet analyzer that recognizes the libpcap format. You can also use FTP or the Session Control Protocol (SCP) to transfer the packet capture files to an external device.

3. Disable packet capture from configuration mode.

NOTE: Disable packet capture before opening the file for analysis or transferring the file to an external device with FTP or SCP. Disabling packet capture ensures that the internal file buffer is flushed and all the captured packets are written to the file.

```
[edit forwarding-options]
user@host# set packet-capture disable
```

4. Display the packet capture.

   - To display the packet capture with the tcpdump utility:

```
user@host# tcpdump -nr /var/log/e2e.pcap
  09:49:55.1413990 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1414154 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
  S 0:460(460) win 0
  09:49:55.1415062 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415184 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
  S 0:460(460) win 0
  09:49:55.1414093 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1414638 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415011 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415129 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415511 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415649 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415249 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
  09:49:55.1415558 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
   S 0:460(460) win 0
```

Copyright © 2016, Juniper Networks, Inc.

```
09:49:55.1414226 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
 S 0:460(460) win 0
09:49:55.1414696 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 > 30.1.1.2.12345:
 S 0:460(460) win 0
09:49:55.1414828 C0/F0/P0 event:16(lt-enter) SEQ:0 IP 10.1.1.2.23451 >
30.1.1.2.12345: S 0:460(460) win 0
09:49:55.1414919 C0/F0/P0 event:15(lt-leave) SEQ:0 IP 10.1.1.2.23451 >
30.1.1.2.12345: S 0:460(460) win 0
09:49:56.417560 C0/F1/P0 event:1(np-ingress) SEQ:0 IP 10.1.1.2.23451 >
30.1.1.2.12345: S 0:460(460) win 0
09:49:56.419263 C0/F1/P0 event:2(np-egress) SEQ:0 IP 10.1.1.2.23451 >
30.1.1.2.12345: S 0:460(460) win 0
```

- To display the packet capture from CLI operational mode:

```
user@host> show security datapath-debug capture
  Packet 1, len 568: (C0/F0/P0/SEQ:0:lbt)
  00 00 00 00 00 00 50 c5 8d 0c 99 4a 00 00 0a 01
  01 02 08 00 45 60 01 f4 00 00 00 00 40 06 4e 9f
  0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
  00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 00 04
  00 00 00 00 b3 e3 15 4e 66 93 15 00 04 22 38 02
  38 02 00 00 00 01 00 03 0b 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08
  Packet 2, len 624: (C0/F0/P0/SEQ:0:lbt)
  aa 35 00 00 00 00 00 00 00 00 00 00 03 00 00
  00 0a 00 00 00 00 00 00 05 bd 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 50 c5
  8d 0c 99 4a 00 00 0a 01 01 02 08 00 45 60 01 f4
  00 00 00 00 40 06 4e 9f 0a 01 01 02 ac 7a 00 04
  00 00 00 00 b3 e3 15 4e 0a 94 15 00 04 5a 70 02
  70 02 00 00 00 03 00 03 0b 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08

  ...
  (Packets 3 through 17 omitted)
  ...

  Packet 18, len 568: (C0/F1/P0/SEQ:0:np-egress)
  00 00 00 04 00 00 00 00 1e 01 01 02 50 c5 8d 0c
  99 4b 08 00 45 60 01 f4 00 00 00 00 3e 06 50 9f
  0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
  00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 04 00
  00 00 00 00 b4 e3 15 4e bf 65 06 00 04 22 38 02
  38 02 00 00 00 11 00 03 02 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08

user@host> show security datapath-debug counters
  Datapath debug counters
  Packet Filter 1:
  lt-enter
  Chassis 0 FPC 0 PIC 1: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
lt-enter
Chassis 0 FPC 0 PIC 0: 1
lt-leave
Chassis 0 FPC 0 PIC 1: 0
lt-leave
Chassis 0 FPC 0 PIC 0: 1
np-egress
Chassis 0 FPC 1 PIC 3: 0
np-egress
Chassis 0 FPC 1 PIC 1: 0
np-egress
Chassis 0 FPC 1 PIC 2: 0
np-egress
Chassis 0 FPC 1 PIC 0: 1
pot
Chassis 0 FPC 0 PIC 1: 0
pot
Chassis 0 FPC 0 PIC 0: 6
np-ingress
Chassis 0 FPC 1 PIC 3: 0
np-ingress
Chassis 0 FPC 1 PIC 1: 0
np-ingress
Chassis 0 FPC 1 PIC 2: 0
np-ingress
Chassis 0 FPC 1 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 4
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 4
```

Related Documentation

- Understanding Data Path Debugging for Logical Systems on page 3476
- *Debugging the Data Path (CLI Procedure)*

Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only)

Problem **Description:** The address of a hostname in an address book entry that is used in a security policy might fail to resolve correctly.

Cause Normally, address book entries that contain dynamic hostnames refresh automatically for SRX Series devices. The TTL field associated with a DNS entry indicates the time after which the entry should be refreshed in the policy cache. Once the TTL value expires, the SRX Series device automatically refreshes the DNS entry for an address book entry.

However, if the SRX Series device is unable to obtain a response from the DNS server (for example, the DNS request or response packet is lost in the network or the DNS server cannot send a response), the address of a hostname in an address book entry might fail

Copyright © 2016, Juniper Networks, Inc.

to resolve correctly. This can cause traffic to drop as no security policy or session match is found.

Solution The master administrator can use the **show security dns-cache** command to display DNS cache information on the SRX Series device. If the DNS cache information needs to be refreshed, the master administrator can use the **clear security dns-cache** command.

NOTE: These commands are only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Related Documentation

- Understanding Logical System Security Policies on page 3306

Copyright © 2016, Juniper Networks, Inc.

PART 52

Configuration Statements and Operational Commands

- Configuration Statements on page 3485
- Operational Commands on page 3637

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 155

Configuration Statements

- Chassis Configuration Statement Hierarchy on page 3487
- Logical-Systems Configuration Statement Hierarchy on page 3490
- Security Configuration Statement Hierarchy on page 3492
- System Configuration Statement Hierarchy on page 3493
- [edit security address-book] Hierarchy Level on page 3524
- [edit security application-firewall] Hierarchy Level on page 3524
- [edit security application-tracking] Hierarchy Level on page 3525
- [edit security datapath-debug] Hierarchy Level on page 3526
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 3527
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 3541
- [edit security nat] Hierarchy Level on page 3542
- [edit security policies] Hierarchy Level on page 3545
- [edit security screen] Hierarchy Level on page 3550
- [edit security softwires] Hierarchy Level on page 3552
- [edit security zones] Hierarchy Level on page 3553
- address-book on page 3555
- address-book (System) on page 3556
- appfw-profile (System) on page 3557
- appfw-rule on page 3558
- appfw-rule-set on page 3559
- application-firewall on page 3560
- application-tracking on page 3561
- auth-entry on page 3562
- cluster (Chassis) on page 3563

Copyright © 2016, Juniper Networks, Inc.

- cpu on page 3565
- datapath-debug on page 3566
- dslite-softwire-initiator on page 3568
- file (System Logging) on page 3569
- firewall-authentication (Security) on page 3571
- flow (Security Flow) on page 3572
- flow-gate on page 3574
- flow-session on page 3575
- idp (Security) on page 3577
- idp-policy on page 3585
- ike (Security) on page 3586
- ipsec (Security) on page 3588
- log (Security) on page 3591
- logical-system (System Security Profile) on page 3593
- logical-systems (All) on page 3594
- nat on page 3595
- nat-cone-binding on page 3599
- nat-destination-pool on page 3600
- nat-destination-rule on page 3601
- nat-interface-port-ol (System) on page 3602
- nat-nopat-address on page 3603
- nat-pat-address on page 3604
- nat-pat-portnum on page 3605
- nat-port-ol-ipnumber on page 3606
- nat-rule-referenced-prefix (System) on page 3607
- nat-source-pool on page 3608
- nat-source-rule on page 3609
- nat-static-rule on page 3610
- policies on page 3611
- policy (System Security Profile) on page 3616
- policy-with-count on page 3617
- profile (Access) on page 3618
- purging on page 3621
- root-authentication on page 3622
- root-logical-system on page 3623
- scheduler (System Security Profile) on page 3624
- screen (Security) on page 3625

Copyright © 2016, Juniper Networks, Inc.

- security-profile on page 3628
- security-profile-resources on page 3631
- softwires on page 3632
- zone (System Security Profile) on page 3633
- zones on page 3634

Chassis Configuration Statement Hierarchy

Use the statements in the **chassis** configuration hierarchy to configure alarms, aggregated devices, clusters, the Routing Engine, and other chassis properties.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority number;
      }
    }
    sonet {
      device-count number;
    }
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
    ethernet {
      link-down (ignore | red | yellow);
    }
    integrated-services {
      failure (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
    serial {
      cts-absent (ignore | red | yellow);
      dcd-absent (ignore | red | yellow);
      dsr-absent (ignore | red | yellow);
      loss-of-rx-clock (ignore | red | yellow);
      loss-of-tx-clock (ignore | red | yellow);
    }
    services {
      hw-down (ignore | red | yellow);
      linkdown (ignore | red | yellow);
      pic-hold-reset (ignore | red | yellow);
      pic-reset (ignore | red | yellow);
      rx-errors (ignore | red | yellow);
```

Copyright © 2016, Juniper Networks, Inc.

```
    sw-down (ignore | red | yellow);
    tx-errors (ignore | red | yellow);
  }
  t3 {
    ais (ignore | red | yellow);
    exz (ignore | red | yellow);
    ferf (ignore | red | yellow);
    idle (ignore | red | yellow);
    lcv (ignore | red | yellow);
    lof (ignore | red | yellow);
    los (ignore | red | yellow);
    pll (ignore | red | yellow);
    ylw (ignore | red | yellow);
  }
}
cluster {
  configuration-synchronize {
    no-secondary-bootup-auto;
  }
  control-link-recovery;
  heartbeat-interval milliseconds;
  heartbeat-threshold number;
  network-management {
    cluster-master;
  }
  redundancy-group  group-number {
    gratuitous-arp-count number;
    hold-down-interval number;
    interface-monitor interface-name {
      weight number;
    }
    ip-monitoring {
      family {
        inet {
          ipv4-address {
            interface {
              logical-interface-name;
              secondary-ip-address ip-address;
            }
            weight number;
          }
        }
      }
      global-threshold number;
      global-weight number;
      retry-count number;
      retry-interval seconds;
    }
    node (0 | 1 ) {
      priority number;
    }
    preempt;
  }
  reth-count number;
  traceoptions {
    file {
```


Copyright © 2016, Juniper Networks, Inc.

```
        filename;
        files number;
        match regular-expression;
        (world-readable | no-world-readable);
        size maximum-file-size;
      }
      flag flag;
      level {
        (alert | all | critical | debug | emergency | error | info | notice | warning);
      }
      no-remote-trace;
    }
  }
  config-button {
    no-clear;
    no-rescue;
  }
  craft-lockout;
  fpc slot-number {
    offline;
    pic slot-number {
      aggregate-ports;
      framing {
        (e1 | e3 | sdh | sonet | t1 | t3);
      }
      ingress-policer-overhead bytes
      max-queues-per-interface (4 | 8);
      mlfr-uni-nni-bundles number;
      no-multi-rate;
      np-cache;
      port slot-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
        speed (oc12-stm4 | oc3-stm1 | oc48-stm16);
      }
      q-pic-large-buffer (large-scale | small-scale);
      services-offload {
        low-latency;
        per-session-statistics;
      }
      shdsl {
        pic-mode (1-port-atm | 2-port-atm | 4-port-atm | efm);
      }
      sparse-dlcis;
      traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode (egress-only | ingress-and-egress);
      }
      tunnel-queuing;
    }
    services-offload;
  }
  ioc-npc-connectivity {
    ioc slot-number {
      npc (npc-slot-number | none);
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    maximum-ecmp (16 | 32 | 64);
    network-services (ethernet | IP);
    routing-engine {
      bios {
        no-auto-upgrade;
      }
      on-disk-failure {
        disk-failure-action (halt | reboot);
      }
      usb-wwan {
        port 1;
      }
    }
    usb {
      storage {
        disable;
      }
    }
  }
```

Related Documentation

- cluster (Chassis) on page 3563
- *ip-monitoring*

Logical-Systems Configuration Statement Hierarchy

To configure a logical system, use the configuration statements in the **logical-systems** configuration hierarchy. As indicated in the following hierarchy, you can include at this hierarchy level several of the hierarchies that can be included at the [**edit**] hierarchy level. However, some statements in a subhierarchy are not valid for logical systems. To learn which statements can be included on your device, issue the **set ?** command at the hierarchy level of interest.

NOTE: The **logical-systems** configuration hierarchy can only be used by the master administrator. Configuration statements in the other hierarchies can also be used by a user logical system administrator at the [**edit**] hierarchy level to configure their user logical system.

```
logical-systems {
  logical-system-name  {
    access{
      ...most statements as in [edit access] Hierarchy Level...
    }
    access-profile profile-name;
    applications {
      ...same statements as in [edit applications] Hierarchy Level...
    }
    firewall {
      ...same statements as in [edit firewall] Hierarchy Level...
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
forwarding-options {
   ...most statements as in [edit forwarding-options] Hierarchy Level...
}
interfaces {
   ...most statements as in [edit interfaces] Hierarchy Level...
}
policy-options {
   ...same statements as in [edit policy-options] Hierarchy Level...
}
protocols {
   ...same statements as in [edit protocols] Hierarchy Level...
}
routing-instances {
   ...most statements in [edit routing-instances] Hierarchy Level...
}
routing-options {
   ...most statements in [edit routing-options] Hierarchy Level...
}
schedulers {
   ...same statements as in the [edit schedulers] Hierarchy Level...
}
security {
   address-book {
      ... same statements as in [edit security address-book] Hierarchy Level...
   }
   application-firewall {
      ... same statements as in [edit security application-firewall] Hierarchy Level...
   }
   application-tracking {
      ... same statements as in [edit security application-tracking] Hierarchy Level...
   }
   firewall-authentication {
      ... same statements as in [edit security firewall-authentication] Hierarchy Level...
   }
   flow {
      ... same statements as in [edit security flow] Hierarchy Level...
   }
   ipsec {
      ... same statements as in [edit security ipsec] Hierarchy Level...
   }
   nat {
      ... same statements as in [edit security nat] Hierarchy Level...
   }
   policies {
      ... same statements as in [edit security policies] Hierarchy Level...
   }
   screen {
      ... same statements as in [edit security screen] Hierarchy Level...
   }
   softwires {
      ... same statements as in [edit security softwires] Hierarchy Level...
   }
   utm {
      ... same statements as in [edit security utm] Hierarchy Level...
   }
   zones {
```

Copyright © 2016, Juniper Networks, Inc.

```
            ... same statements as in [edit security zones] Hierarchy Level...
          }
        }
        system {
          ...most statements as in the [edit system] Hierarchy Level...
        }
      }
    }
```

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537

Copyright © 2016, Juniper Networks, Inc.

- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

System Configuration Statement Hierarchy

Use the statements in the **system** configuration hierarchy to configure system management functions including addresses of the Domain Name System (DNS) servers; device's hostname, address, and domain name; health monitoring; interface filtering; properties of the device's auxiliary and console ports; security profiles for logical systems; time zones and Network Time Protocol (NTP) properties; trace options; and user login accounts, including user authentication and the root-level user account. Statement descriptions that are exclusive to the SRX Series devices running Junos OS are described in this section.

```
system {
    accounting {
        destination {
            radius {
                server server-address {
                    accounting-port port-number;
                    max-outstanding-requests number;
                    port number;
                    retry number;
                    secret password;
                    source-address address;
                    timeout seconds;
                }
            }
            tacplus {
                server server-address {
```

Copyright © 2016, Juniper Networks, Inc.

```
            port port-number;
            secret password;
            single-connection;
            source-address source-address;
            timeout seconds;
          }
        }
      }
      events [change-log interactive-commands login];
      traceoptions {
        file {
          filename;
          files number;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
    allow-v4mapped-packets;
    archival {
      configuration {
        archive-sites url {
          password password;
        }
        transfer-interval interval;
        transfer-on-commit;
      }
    }
    arp {
      aging-timer minutes;
      gratuitous-arp-delay seconds;
      gratuitous-arp-on-ifup;
      interfaces {
        interface name {
          aging-timer minutes;
        }
      }
      passive-learning;
      purging;
    }
    authentication-order [password radius tacplus];
    auto-configuration {
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
auto-snapshot;
autoinstallation {
  configuration-servers {
    url {
      password password;
    }
  }
  interfaces {
    interface-name {
      bootp;
      rarp;
    }
  }
  usb {
    disable;
  }
}
auto-snapshot;
backup-router {
  address;
  destination [network];
}
commit {
  server {
    commit-interval seconds;
    days-to-keep-error-logs days;
    maximum-aggregate-pool number;
    maximum entries number;
    traceoptions {
      file {
        filename;
        files number;
        microsecond-stamp;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  synchronize;
}
compress-configuration-files;
default-address-selection;
diag-port-authentication {
  encrypted-password passsword;
  plain-text-password;
}
domain-name domain-name;
domain-search [domain-list];
donot-disable-ip6op-ondad;
dump-device (boot-device | compact-flash | usb);
dynamic-profile-options {
  versioning;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
encrypt-configuration-files;
extensions {
  providers {
    provider-id {
      license-type license deployment-scope [deployments];
    }
  }
  resource-limits {
    package package-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size mbytes;
          locked-in mbytes;
          resident-set-size mbytes;
          socket-buffers mbytes;
          stack-size mbytes;
        }
      }
    }
    process process-ui-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size mbytes;
          locked-in mbytes;
          resident-set-size mbytes;
          socket-buffers mbytes;
          stack-size mbytes;
        }
      }
    }
  }
}
fips {
  level (0 | 1 | 2 | 3 | 4);
}
host-name hostname;
inet6-backup-router {
  address;
```


Copyright © 2016, Juniper Networks, Inc.

```
    destination destination;
  }
  internet-options {
    icmpv4-rate-limit {
      bucket-size seconds;
      packet-rate packets-per-second;
    }
    icmpv6-rate-limit {
      bucket-size seconds;
      packet-rate packets-per-second;
    }
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    ipv6-duplicate-addr-detection-transmits number;
    (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
    ipv6-path-mtu-discovery-timeout minutes;
    no-tcp-reset (drop-all-tcp | drop-tcp-with-syn-only);
    no-tcp-rfc1323;
    no-tcp-rfc1323-paws;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port upper-limit upper-limit;
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
    tcp-mss bytes;
  }
  kernel-replication;
  license {
    autoupdate {
    url url;
    password password;
    }
    renew {
      before-expiration number;
      interval interval-hours;
    }
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
```

```
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
  }
  login {
    announcement text;
    class class-name {
      access-end hh:mm;
      access-start hh:mm;
      allow-commands regular-expression;
      allow-configuration regular-expression;
      allow-configuration-regexps [regular-expression];
      allowed-days [day];
      deny-commands regular-expression;
      deny-configuration regular-expression;
      deny-configuration-regexps [regular-expression];
      idle-timeout minutes;
      logical-system logical-system;
      login-alarms;
      login-script script;
      login-tip;
      permissions [permissions ];
      security-role (audit-administrator | crypto-administrator | ids-administrator |
        security-administrator);
    }
    deny-sources {
      address [address-or-hostname];
    }
    message text;
    }
    password {
      change-type (character-set | set-transitions);
      format (des | md5 | sha1);
      maximum-length length;
      minimum-changes number;
      minimum-length length;
    }
    retry-options {
      backoff-factor seconds;
      backoff-threshold number;
      lockout-period time;
      maximum-time seconds;
      minimum-time seconds;
      tries-before-disconnect number;
    }
    user username {
      authentication {
        encrypted-password password;
        load-key-file url;
        plain-text-password;
        ssh-dsa public-key;
        ssh-rsa public-key;
      }
      class class-name;
      full-name complete-name;
      uid uid-value;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  log-vital {
    interval minutes;
    files days;
    storage-limit percentage;
    file-size Mbytes;
    add oid{
      comment comment;
    }
    group {
      operating;
      idp;
      storage;
      cluster-counter;
      screen zone-name;
      spu spu-name;
      }
    }
  max-configuration-rollbacks number;
  max-configurations-on-flash number;
  mirror-flash-on-disk;
  name-server ip-address;
  nd-maxmcast-solicit value;
  nd-retrasmit-timer value;
  no-compress-configuration-files;
  no-debugger-on-alt-break;
  no-multicast-echo;
  no-neighbor-learn;
  no-ping-record-route;
  no-ping-time-stamp;
  no-redirects;
  no-saved-core-context;
  ntp {
    authentication-key key-number {
      type md5;
      value password;
    }
    boot-server address;
    broadcast broadcast-address  {
      key key;
      ttl value;
      version version;
    }
    broadcast-client;
    multicast-client {
      address;
    }
    peer peer-address {
      key key;
      prefer;
      version version;
    }
    server server-address {
      key key;
      prefer;
```

```
        version version;
    }
    source-address source-address;
    trusted-key [key-number];
}
pic-console-authentication {
    encrypted-password password;
    plain-text-password;
}
ports {
    auxiliary {
        disable;
        insecure;
        type (ansi | small-xterm | vt100 | xterm);
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type (ansi | small-xterm | vt100 | xterm);
    }
}
processes {
    802.1x-protocol-daemon {
        command binary-file-path;
        disable;
    }
    adaptive-services {
        command binary-file-path;
        disable;
        failover (alternate-media | other-routing-engine);
    }
    alarm-control {
        command binary-file-path;
        disable;
        failover (alternate-media | other-routing-engine);
    }
    application-identification {
        command binary-file-path;
        disable;
        failover (alternate-media | other-routing-engine);
    }
    application-security {
        command binary-file-path;
        disable;
        failover (alternate-media | other-routing-engine);
    }
    audit-process {
        command binary-file-path;
        disable;
    }
    auto-configuration {
        command binary-file-path;
        disable;
        failover (alternate-media | other-routing-engine);
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
bootp {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
chassis-control {
  disable;
  failover alternate-media;
}
class-of-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
craft-control {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
database-replication {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
datapath-trace-service {
  disable;
  traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
dhcp {
  command binary-file-path;
  disable;
}
dhcp-service {
  disable;
  failover (alternate-media | other-routing-engine);
  interface-traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
```

Copyright © 2016, Juniper Networks, Inc.

```
    no-remote-trace;
  }
  traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
dialer-services {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
diameter-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
disk-monitoring {
  command binary-file-path;
  disable;
}
dynamic-flow-capture {
  command binary-file-path;
  disable;
}
ecc-error-logging {
  command binary-file-path;
  disable;
}
```


Copyright © 2016, Juniper Networks, Inc.

```
ethernet-connectivity-fault-management {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
ethernet-link-fault-management {
  command binary-file-path;
  disable;
}
ethernet-switching {
  command binary-file-path;
  disable;
}
event-processing {
  command binary-file-path;
  disable;
}
fipsd {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
firewall {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
firewall-authentication-service {
  disable;
}
forwarding {
  command binary-file-path;
  disable;
}
general-authentication-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
gprs-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
group-key-member {
  disable;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
group-key-server {
   disable;
}
idp-policy {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ilmi {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
inet-process {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
init {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
interface-control {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ipmi {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
ipsec-key-management {
   (disable | enable);
}
jsrp-service {
   disable;
}
jtasktest {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
kernel-replication {
   command binary-file-path;
   disable;
}
l2-learning {
   command binary-file-path;
   disable;
}
l2cpd-service {
   command binary-file-path;
   disable;
}
```


Copyright © 2016, Juniper Networks, Inc.

```
lacp {
  command binary-file-path;
  disable;
}
lldpd-service {
  command binary-file-path;
  disable;
}
logical-system-mux {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
logical-system-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
mib-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mobile-ip {
  command binary-file-path;
  disable;
}
mountd-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mspd {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
multicast-snooping {
  command binary-file-path;
  disable;
}
named-service {
  disable;
  failover (alternate-media | other-routing-engine);
}
neighbor-liveness {
  command binary-file-path;
```

Copyright © 2016, Juniper Networks, Inc.

```
    disable;
    failover (alternate-media | other-routing-engine);
}
network-security {
    disable;
}
network-security-trace {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
nfsd-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
ntp {
    disable;
    failover (alternate-media | other-routing-engine);
}
ntpd-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
peer-selection-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
periodic-packet-services {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pgcp-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pgm {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
pic-services-logging {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
ppp {
    command binary-file-path;
    disable;
}
pppoe {
    command binary-file-path;
```


Copyright © 2016, Juniper Networks, Inc.

```
    disable;
}
process-monitor {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error |info |notice | verbose | warning);
        no-remote-trace;
    }
}
profilerd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
r2cp {
    command binary-file-path;
    disable;
}
redundancy-interface-process {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
remote-operations {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
resource-cleanup {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error |info |notice | verbose | warning);
        no-remote-trace;
    }
}
routing {
    disable;
    failover (alternate-media | other-routing-engine);
}
sampling {
```

```
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  sbc-configuration-process {
    disable;
    failover (alternate-media | other-routing-engine);
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  sdk-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error |info |notice | verbose | warning);
      no-remote-trace;
    }
  }
  secure-neighbor-discovery {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  security-log {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  send {
    disable;
  }
  service-deployment {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  shm-rtsdbd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
```


Copyright © 2016, Juniper Networks, Inc.

```
}
simple-mail-client-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
smtpd-service {
   disable;
}
snmp {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
static-subscribers {
   disable;
}
statistics-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
subscriber-management {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
subscriber-management-helper {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
system-health-management {
   disable;
}
system-log-vital {
   disable;
}
tunnel-oamd {
   command binary-file-path;
   disable;
}
uac-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
usb-control {
   command binary-file-path;
   disable;
}
virtualization-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
```

Copyright © 2016, Juniper Networks, Inc.

```
vrrp {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
wan-acceleration {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
watchdog {
  enable;
  disable;
  timeout value;
}
web-management {
  disable;
  failover (alternate media | other-routing-engine);
}
wireless-lan-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
wireless-wan-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

```
proxy {
    password password;
    port port-number;
    server url;
    username user-name;
}
radius-options {
    attributes {
        nas-ip-address nas-ip-address;
        }
    password-protocol mschap-v2;
}
radius-server server-address {
    accounting-port number;
    max-outstanding-requests number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
root-authentication {
    encrypted-password password;
    load-key-file url;
    plain-text-password;
    ssh-dsa public-key {
        <from pattern-list>;
    }
    ssh-rsa public-key {
        <from pattern-list>;
    }
}
saved-core-context;
saved-core-files number;
scripts {
    commit {
        allow-transients;
        direct-access;
        file filename {
            checksum (md5 | sha-256 | sha1);
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        refresh;
        refresh-from url;
        traceoptions {
            file {
                filename;
                files number;
                size maximum-file-size;
                (world-readable | no-world-readable);
            }
            flag flag;
            no-remote-trace;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    load-scripts-from-flash;
    op {
        file filename {
            arguments name {
                description text;
            }
            checksum (md5 | sha-256 | sha1);
            command filename-alias;
            description cli-help-text;
            refresh;
            refresh-from url;
            source url;
        }
        no-allow-url;
        refresh;
        refresh-from url;
        traceoptions {
            file {
                filename;
                files number;
                size maximum-file-size;
                (world-readable | no-world-readable);
            }
            flag flag;
            no-remote-trace;
        }
}
security-profile security-profile-name {
    address-book {
        maximum amount;
        reserved amount;
    }
    appfw-profile {
        maximum amount;
        reserved amount;
    }
    appfw-rule {
        maximum amount;
        reserved amount;
    }
    appfw-rule-set {
        maximum amount;
        reserved amount;
    }
    auth-entry {
        maximum amount;
        reserved amount;
    }
    cpu {
        reserved percent;
    }
    dslite-softwire-initiator {
        maximum amount;
        reserved amount;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
flow-gate {
    maximum amount;
    reserved amount;
}
flow-session {
    maximum amount;
    reserved amount;
}
idp-policy idp-policy-name;
logical-system logical-system-name;
nat-cone-binding {
    maximum amount;
    reserved amount;
}
nat-destination-pool {
    maximum amount;
    reserved amount;
}
nat-destination-rule {
    maximum amount;
    reserved amount;
}
nat-interface-port-ol {
    maximum amount;
    reserved amount;
}
nat-nopat-address {
    maximum amount;
    reserved amount;
}
nat-pat-address {
    maximum amount;
    reserved amount;
}
nat-pat-portnum {
    maximum amount
    reserved amount
}
nat-port-ol-ipnumber {
    maximum amount;
    reserved amount;
}
nat-rule-referenced-prefix {
    maximum amount;
    reserved amount;
}
nat-source-pool {
    maximum amount;
    reserved amount;
}
nat-source-rule {
    maximum amount;
    reserved amount;
}
nat-static-rule {
```

Copyright © 2016, Juniper Networks, Inc.

```
            maximum amount;
            reserved amount;
        }
        policy {
            maximum amount;
            reserved amount;
        }
        policy-with-count {
            maximum amount;
            reserved amount;
        }
        root-logical-system;
        scheduler {
            maximum amount;
            reserved amount;
        }
        zone {
            maximum amount;
            reserved amount;
        }
    }
    security-profile-resources {
      cpu-control;
      cpu-control-target percent;
    }
    services {
      database-replication {
        traceoptions {
          file {
            filename ;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
        }
      }
      dhcp {
        boot-file filename;
        boot-server (address | hostname);
        default-lease-time (infinite | seconds);
        domain-name domain-name;
        domain-search dns-search-suffix;
        maximum-lease-time (infinite | seconds);
        name-server ip-address;
        next-server ip-address;
        option option-identifier-code array type-name [ type-values ] | byte 8-bit-value | flag
          (false | off | on | true) | integer signed-32-bit-value | ip-address address | short
          signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
          unsigned-short 16-bit-value);
        pool subnet-ip-address/mask {
          address-range {
            high address;
            low address;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        boot-file filename;
        boot-server (address | hostname);
        default-lease-time (infinite | seconds);
        domain-name domain-name;
        domain-search dns-search-suffix;
        exclude-address ip-address;
        maximum-lease-time (infinite | seconds);
        name-server ip-address;
        next-server ip-address;
        option option-identifier-code array type-name [ type-values ] | byte 8-bit-value |
          flag (false | off | on | true) | integer signed-32-bit-value | ip-address address |
          short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
          unsigned-short 16-bit-value);
        propagate-ppp-settings interface-name;
        propagate-settings interface-name;
        router ip-address;
        server-identifier dhcp-server;
        sip-server {
          address ip-address;
          name sip-server-name;
        }
        wins-server ip-address;
      }
      propagate-ppp-settings interface-name;
      propagate-settings interface-name;
      router ip-address;
      server-identifier dhcp-server;
      sip-server {
        address ip-address;
        name sip-server-name;
      }
      static-binding mac-address;
      traceoptions {
        file {
          filename ;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
      wins-server ip-address;
    }
    dhcp-local-server {
      dhcpv6 {
        authentication {
          password password;
          username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name;
```

```
      interface-name;
      logical-system-name;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix;
    }
  }
  dynamic-profile {
    profile-name;
    aggregate-clients {
      merge;
      replace;
    }
    junos-default-profile;
    use-primary dynamic-profile-name;
  }
  group group-name {
    authentication {
      password password;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix;
      }
    }
    dynamic-profile {
      profile-name;
      aggregate-clients {
        merge;
        replace;
      }
      junos-default-profile;
      use-primary dynamic-profile;
    }
    interface interface-name {
      dynamic-profile {
        profile-name;
        aggregate-clients {
          merge;
          replace;
        }
        junos-default-profile;
        use-primary dynamic-profile-name;
      }
      exclude;
      overrides {
```


Copyright © 2016, Juniper Networks, Inc.

```
            delegated-pool pool-name;
            interface-client-limit number;
            process-inform {
              pool pool-name;
            }
            rapid-commit ;
          }
          service-profile service-profile-name
          trace ;
          upto interface-name;
        }
        liveness-detection {
          failure-action {
            clear-binding;
            clear-binding-if-interface-up;
            log-only;
          }
          method {
            bfd {
              detection-time {
                threshold milliseconds;
              }
              holddown-interval interval;
              minimum-interval milliseconds;
              minimum-receive-interval milliseconds;
              multiplier number;
              no-adaptation;
              session-mode (automatic | multihop | single-hop);
              transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
              }
            version (0 | 1 | automatic);
          }
        }
        overrides {
          delegated-pool pool-name;
          interface-client-limit number;
          process-inform {
            pool pool-name;
          }
          rapid-commit ;
        }
        reconfigure {
          attempts number;
          clear-on-abort;
          strict;
          timeout number;
          token token-name;
          trigger {
            radius-disconnect;
          }
        }
        service-profile service-profile-name;
      }
      liveness-detection {
```

```
            failure-action {
              clear-binding;
              clear-binding-if-interface-up;
              log-only;
            }
            method {
              bfd {
                detection-time {
                  threshold milliseconds;
                }
                holddown-interval interval;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                session-mode (automatic | multihop | single-hop);
                transmit-interval {
                  minimum-interval milliseconds;
                  threshold milliseconds;
                }
              version (0 | 1 | automatic);
            }
          }
          overrides {
            delegated-pool pool-name;
            interface-client-limit number;
            process-inform {
              pool pool-name;
            }
            rapid-commit ;
          }
          reconfigure {
            attempts number;
            clear-on-abort;
            strict;
            timeout number;
            token token-name;
            trigger {
              radius-disconnect;
            }
          }
          service-profile service-profile-name;
        }
        group group-name {
          interface interface-name {
            exclude;
            upto upto-interface-name;
          }
        }
      }
      dns {
          dns-proxy {
            cache hostname inet ip-address;
            default-domain domain-name {
              forwarders ip-address;
            }
```


Copyright © 2016, Juniper Networks, Inc.

```
        interface interface-name;
        propogate-setting (enable | disable);
        view view-name {
          domain domain-name {
            forward-only;
            forwarders ip-address;
          }
          match-clients subnet-address;
        }
      }
    }
    dnssec {
      disable;
      dlv {
        domain-name domain-name trusted-anchor  trusted-anchor;
        }
    secure-domains domain-name;
      trusted-keys (key dns-key | load-key-file url);
    forwarders {
      ip-address;
    }
    max-cache-ttl seconds;
    max-ncache-ttl seconds;
    traceoptions {
      category {
        category-type;
      }
      debug-level level;
      file {
        filename;
        files number;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  dynamic-dns {
    client hostname {
      agent agent-name;
      interface interface-name;
      password server-password;
      server server-name;
      username user-name;
    }
  }
  finger {
    connection-limit number;
    rate-limit number;
  }
  ftp {
    connection-limit number;
    rate-limit number;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
netconf {
  ssh {
    connection-limit number;
    port port-number;
    rate-limit number;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    on-demand;
  }
}
outbound-ssh {
  client client-id {
    address {
      port port-number;
      retry number;
      timeout value;
    }
    device-id device-id;
    keep-alive {
      retry number;
      time-out value;
    }
    reconnect-strategy (in-order |sticky);
    secret secret;
    services {
      netconf;
    }
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
service-deployment {
  local-certificate certificate-name;
  servers server-address {
    port port-number;
    security-options {
      ssl3;
      tls;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        user user-name;
    }
    source-address source-address;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
ssh {
    ciphers [cipher];
    client-alive-count-max number;
    client-alive-interval seconds;
    connection-limit number;
    hostkey-algorithm {
        (ssh-dss | no-ssh-dss);
        (ssh-ecdsa |no-ssh-ecdsa);
        (ssh-rsa | no-ssh-rsa);
    }
    key-exchange [algorithm];
    macs [algorithm];
    max-sessions-per-connection number;
    protocol-version {
        v1;
        v2;
    }
    rate-limit number;
    root-login (allow | deny | deny-password);
    (tcp-forwarding | no-tcp-forwarding);
}
subscriber-management {
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    maintain-subscriber interface-delete;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
subscriber-management-helper {
    traceoptions {
        file {
```

```
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
telnet {
    connection-limit number;
    rate-limit number;
}
web-management {
    control {
        max-threads number;
    }
    http {
        interface [interface-name];
        port port-number;
    }
    https {
        interface [interface-name];
        local-certificate name;
        pki-local-certificate name;
        port port-number;
        system-generated-certificate;
    }
    management-url url;
    session {
        idle-timeout minutes;
        session-limit number;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
xnm-clear-text {
    connection-limit number;
    rate-limit number;
}
xnm-ssl {
    connection-limit number;
    local-certificate name;
    rate-limit number;
}
```


Copyright © 2016, Juniper Networks, Inc.

```
}
static-host-mapping hostname {
  alias [host-name-alias];
  inet [ip- address];
  inet6 [ipv6- address];
  sysid system-identifier;
}
syslog {
  allow-duplicates;
  archive {
    binary-data;
    files number;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  console {
    (any | facility) severity;
  }
  file filename {
    allow-duplicates;
    archive {
      archive-sites url {
        password password;
      }
      (binary-data| no-binary-data);
      files number;
      size maximum-file-size;
      start-time "YYYY-MM-DD.hh:mm";
      transfer-interval minutes;
      (world-readable | no-world-readable);
    }
    structure-data {
      brief;
    }
    (any | facility) severity;
  }
  host (hostname | other-routing-engine) {
    (any | facility) severity;
  }
  log-rotate-frequency minutes;
  source-address source-address;
  time-format {
    millisecond;
    year;
  }
  user (username | *) {
    (any | facility) severity;
  }
}
tacplus-options {
  (exclude-cmd-attribute | no-cmd-attribute-value);
  service-name service-name;
}
tacplus-server server-address {
  port port-number;
  secret password;
```

```
    single-connection;
    source-address source-address;
    timeout seconds;
  }
  time-zone (GMThour-offset | time-zone);
  tracing {
    destination-override {
      syslog {
        host address;
        }
    }
  }
  use-imported-time-zones;
}
```

Related Documentation
- Security Configuration Statement Hierarchy on page 603

[edit security address-book] Hierarchy Level

```
security {
  address-book (book-name | global) {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
    attach {
      zone zone-name;
    }
    description text;
  }
}
```

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- Understanding Address Books on page 1069

[edit security application-firewall] Hierarchy Level

```
security {
  application-firewall {
```


Copyright © 2016, Juniper Networks, Inc.

```
    profile profile-name {
      block-message type {
        custom-text content custom-html-text;
        custom-redirect-url content custom-redirect-url;
      }
    }
    rule-sets rule-set-name {
      default-rule {
        (deny [block-message] | permit | reject [block-message]);
      }
      profile profile-name;
      rule rule-name {
        match {
          dynamic-application [system-application];
          dynamic-application-group [system-application-group];
          ssl-encryption (any | yes | no);
        }
        then {
          (deny [block-message] | permit | reject [block-message]);
        }
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Application Firewall Overview on page 555

[edit security application-tracking] Hierarchy Level

```
security {
  application-tracking {
    disable;
    (first-update | first-update-interval first-update-interval);
    session-update-interval session-update-interval;
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

[edit security datapath-debug] Hierarchy Level

```
security {
  datapath-debug {
    action-profile profile-name {
      event (jexec | lbt | lt-enter | lt-leave | mac-egress | mac-ingress | np-egress |
        np-ingress | pot) {
        count;
        packet-dump;
        packet-summary;
        trace;
      }
      module {
        flow {
          flag {
            all;
          }
        }
      }
      preserve-trace-order;
      record-pic-history;
    }
    capture-file {
      filename;
      files files-number;
      format pacp-format;
      (no-world-readable | world-readable);
      size maximum-file-size;
    }
    maximum-capture-size value;
    packet-filter packet-filter-name {
      action-profile (profile-name | default);
      destination-port (port-range | protocol-name);
      destination-prefix destination-prefix;
      interface logical-interface-name;
      protocol (protocol-number | protocol-name;
      source-port (port-range | protocol- name);
      source-prefix source-prefix;
      }
    trace-options {
      file {
        filename;
        files files-number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

[edit security firewall-authentication] Hierarchy Level

```
security {
  firewall-authentication {
    traceoptions {
      flag flag;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security flow] Hierarchy Level

```
security {
  flow {
    aging {
      early-ageout seconds;
      high-watermark percent;
      low-watermark percent;
    }
    allow-dns-reply;
    allow-embedded-icmp;
    ethernet-switching {
      block-non-ip-all;
      bpdu-vlan-flooding;
      bypass-non-ip-unicast;
      no-packet-flooding {
        no-trace-route;
      }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
    load distribution {
      session-affinity ipsec;
    }
    pending-sess-queue-length (high | moderate | normal);
    route-change-timeout seconds;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    tcp-mss {
      all-tcp mss value;
      gre-in {
        mss value;
      }
      gre-out {
        mss value;
      }
      ipsec-vpn {
        mss value;
      }
    }
    tcp-session {
      fin-invalidate-session;
```

Copyright © 2016, Juniper Networks, Inc.

```
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
          (session-ageout | session-timeout seconds);
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
        packet-filter filter-name {
          destination-port port-identifier;
          destination-prefix address;
          interface interface-name;
          protocol protocol-identifier;
          source-port port-identifier;
          source-prefix address;
        }
        rate-limit messages-per-second;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Juniper Networks Devices Processing Overview on page 1665
- *Network Monitoring and Troubleshooting Guide for Security Devices*

[edit security idp] Hierarchy Level

```
security {
  idp {
    active-policy policy-name;
    custom-attack attack-name {
      attack-type {
        anomaly {
          direction (any | client-to-server | server-to-client);
          service service-name;
          shellcode (all | intel | no-shellcode | sparc);
          test test-condition;
        }
        chain {
          expression boolean-expression;
```


Copyright © 2016, Juniper Networks, Inc.

```
    member member-name {
      attack-type {
        (anomaly ...same statements as in [edit security idp custom-attack
          attack-name attack-type anomaly] hierarchy level | signature ...same
          statements as in [edit security idp custom-attack attack-name attack-type
          signature] hierarchy level);
      }
    }
    order;
    protocol-binding {
      application application-name;
      icmp;
      icmpv6;
      ip {
        protocol-number transport-layer-protocol-number;
      }
      ipv6 {
        protocol-number transport-layer-protocol-number;
      }
      rpc {
        program-number rpc-program-number;
      }
      tcp {
        minimum-port port-number <maximum-port port-number>;
      }
      udp {
        minimum-port port-number <maximum-port port-number>;
      }
    }
    reset;
    scope (session | transaction);
  }
  signature {
    context context-name;
    direction (any | client-to-server | server-to-client);
    negate;
    pattern signature-pattern;
    pattern-pcre signature-pattern-pcre;
    protocol {
      icmp {
        code {
          match (equal | greater-than | less-than | not-equal);
          value code-value;
        }
        data-length {
          match (equal | greater-than | less-than | not-equal);
          value data-length;
        }
        identification {
          match (equal | greater-than | less-than | not-equal);
          value identification-value;
        }
        sequence-number {
          match (equal | greater-than | less-than | not-equal);
          value sequence-number;
        }
```

```
    type {
      match (equal | greater-than | less-than | not-equal);
      value type-value;
    }
  }
  icmpv6 {
    code {
      match (equal | greater-than | less-than | not-equal);
      value code-value;
    }
    data-length {
      match (equal | greater-than | less-than | not-equal);
      value data-length;
    }
    identification {
      match (equal | greater-than | less-than | not-equal);
      value identification-value;
    }
    sequence-number {
      match (equal | greater-than | less-than | not-equal);
      value sequence-number;
    }
    type {
      match (equal | greater-than | less-than | not-equal);
      value type-value;
    }
  }
  ipv4 {
    destination {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    identification {
      match (equal | greater-than | less-than | not-equal);
      value identification-value;
    }
    ip-flags {
      (df | no-df);
      (mf | no-mf);
      (rb | no-rb);
    }
    protocol {
      match (equal | greater-than | less-than | not-equal);
      value transport-layer-protocol-id;
    }
    source {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    tos {
      match (equal | greater-than | less-than | not-equal);
      value type-of-service-in-decimal;
    }
    total-length {
      match (equal | greater-than | less-than | not-equal);
      value total-length-of-ip-datagram;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    ttl {
      match (equal | greater-than | less-than | not-equal);
      value time-to-live;
    }
  }
  ipv6 {
    destination {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    flow-label {
      match (equal | greater-than | less-than | not-equal);
      value flow-label-value;
    }
    hop-limit {
      match (equal | greater-than | less-than | not-equal);
      value hop-limit-value;
    }
    next-header {
      match (equal | greater-than | less-than | not-equal);
      value next-header-value;
    }
    payload-length {
      match (equal | greater-than | less-than | not-equal);
      value payload-length-value;
    }
    source {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    traffic-class {
      match (equal | greater-than | less-than | not-equal);
      value traffic-class-value;
    }
  tcp {
    ack-number {
      match (equal | greater-than | less-than | not-equal);
      value acknowledgement-number;
    }
    data-length {
      match (equal | greater-than | less-than | not-equal);
      value tcp-data-length;
    }
    destination-port {
      match (equal | greater-than | less-than | not-equal);
      value destination-port;
    }
    header-length {
      match (equal | greater-than | less-than | not-equal);
      value header-length;
    }
    mss {
      match (equal | greater-than | less-than | not-equal);
      value maximum-segment-size;
    }
```

```
        option {
          match (equal | greater-than | less-than | not-equal);
          value tcp-option;
        }
        sequence-number {
          match (equal | greater-than | less-than | not-equal);
          value sequence-number;
        }
        source-port {
          match (equal | greater-than | less-than | not-equal);
          value source-port;
        }
        tcp-flags {
          (ack | no-ack);
          (fin | no-fin);
          (psh | no-psh);
          (r1 | no-r1);
          (r2 | no-r2);
          (rst | no-rst);
          (syn | no-syn);
          (urg | no-urg);
        }
        urgent-pointer {
          match (equal | greater-than | less-than | not-equal);
          value urgent-pointer;
        }
        window-scale {
          match (equal | greater-than | less-than | not-equal);
          value window-scale-factor;
        }
        window-size {
          match (equal | greater-than | less-than | not-equal);
          value window-size;
        }
      }
      udp {
        data-length {
          match (equal | greater-than | less-than | not-equal);
          value data-length;
        }
        destination-port {
          match (equal | greater-than | less-than | not-equal);
          value destination-port;
        }
        source-port {
          match (equal | greater-than | less-than | not-equal);
          value source-port;
        }
      }
    }
    protocol-binding {
      application application-name;
      icmp;
      icmpv6;
      ip {
        protocol-number transport-layer-protocol-number;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            ipv6 {
              protocol-number transport-layer-protocol-number;
            }
            rpc {
              program-number rpc-program-number;
            }
            tcp {
              minimum-port port-number <maximum-port port-number>;
            }
            udp {
              minimum-port port-number <maximum-port port-number>;
            }
          }
          regexp regular-expression;
          shellcode (all | intel | no-shellcode | sparc);
        }
      }
      recommended-action (close | close-client | close-server | drop | drop-packet | ignore
        | none);
      severity (critical | info | major | minor | warning);
      time-binding {
        count count-value;
        scope (destination | peer | source);
      }
    }
    custom-attack-group custom-attack-group-name {
      group-members [attack-or-attack-group-name];
    }
    dynamic-attack-group dynamic-attack-group-name {
      filters {
        category {
          values [category-value];
        }
        direction {
          expression (and | or);
          values [any client-to-server exclude-any exclude-client-to-server
            exclude-server-to-client server-to-client];
        }
        false-positives {
          values [frequently occasionally rarely unknown];
        }
        performance {
          values [fast normal slow unknown];
        }
        products {
          values [product-value];
        }
        recommended;
        no-recommended;
        service {
          values [service-value];
        }
        severity {
          values [critical info major minor warning];
        }
```

```
    type {
      values [anomaly signature];
    }
  }
}
idp-policy policy-name {
  rulebase-exempt {
    rule rule-name {
      description text;
      match {
        attacks {
          custom-attack-groups [attack-group-name];
          custom-attacks [attack-name];
          dynamic-attack-groups [attack-group-name];
          predefined-attack-groups [attack-group-name];
          predefined-attacks [attack-name];
        }
        destination-address ([address-name] | any | any-ipv4 | any-ipv6);
        destination-except [address-name];
        from-zone (zone-name | any );
        source-address ([address-name] | any | any-ipv4 | any-ipv6);
        source-except [address-name];
        to-zone (zone-name | any);
      }
    }
  }
  rulebase-ips {
    rule rule-name {
      description text;
      match {
        application (application-name | any | default);
        attacks {
          custom-attack-groups [attack-group-name];
          custom-attacks [attack-name];
          dynamic-attack-groups [attack-group-name];
          predefined-attack-groups [attack-group-name];
          predefined-attacks [attack-name];
        }
        destination-address ([address-name] | any | any-ipv4 | any-ipv6);
        destination-except [address-name];
        from-zone (zone-name | any );
        source-address ([address-name] | any | any-ipv4 | any-ipv6);
        source-except [address-name];
        to-zone (zone-name | any);
      }
      terminal;
      then {
        action {
          class-of-service {
            dscp-code-point number;
            forwarding-class forwarding-class;
          }
          (close-client | close-client-and-server | close-server |drop-connection |
            drop-packet | ignore-connection | mark-diffserv value | no-action |
            recommended);
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
            ip-action {
              (ip-block | ip-close | ip-notify);
              log;
              log-create;
              refresh-timeout;
              target (destination-address | service | source-address | source-zone |
                source-zone-address | zone-service);
              timeout seconds;
            }
            notification {
              log-attacks {
                alert;
              }
              packet-log {
                post-attack number;
                post-attack-timeout seconds;
                pre-attack number;
              }
            }
            severity (critical | info | major | minor | warning);
          }
        }
      }
    }
    security-package {
      automatic {
        download-timeout minutes;
        enable;
        interval hours;
        start-time start-time;
      }
      install {
        ignore-version-check;
      }
      source-address address;
      url url-name;
    }
    sensor-configuration {
      application-identification {
        max-packet-memory-ratio percemtage-value;
        max-reass-packet-memory-ratio percemtage-value;
        max-tcp-session-packet-memory value;
        max-udp-session-packet-memory value;
      }
      detector {
        protocol-name protocol-name {
          tunable-name  tunable-name {
            tunable-value protocol-value;
          }
        }
      }
      flow {
        (allow-icmp-without-flow | no-allow-icmp-without-flow);
        drop-if-no-policy-loaded;
        drop-on-failover;
        drop-on-limit;
```

```
    fifo-max-size value;
    hash-table-size value;
    (log-errors | no-log-errors);
    max-sessions-offset value;
    max-timers-poll-ticks value;
    min-objcache-limit-lt lower-threshold-value;
    min-objcache-limit-ut upper-threshold-value;
    reject-timeout value;
    (reset-on-policy | no-reset-on-policy);
    udp-anticipated-timeout value;
}
global {
    (enable-all-qmodules | no-enable-all-qmodules);
    (enable-packet-pool | no-enable-packet-pool);
    memory-limit-percent value;
    (policy-lookup-cache | no-policy-lookup-cache);
}
high-availability {
    no-policy-cold-synchronization;
}
ips {
    content-decompression-max-memory-kb value;
    content-decompression-max-ratio value;
    (detect-shellcode | no-detect-shellcode);
    fifo-max-size value;
    (ignore-regular-expression | no-ignore-regular-expression);
    log-supercede-min minimum-value;
    pre-filter-shellcode;
    (process-ignore-s2c | no-process-ignore-s2c);
    (process-override | no-process-override);
    process-port port-number;
}
log {
    cache-size size;
    suppression {
        disable;
        (include-destination-address | no-include-destination-address);
        max-logs-operate value;
        max-time-report value;
        start-log value;
    }
}
packet-log {
    host ip-address <port number>;
    max-sessions percentage;
    source-address ip-address;
    total-memory percentage;
}
re-assembler {
    action-on-reassembly-failure (drop | drop-session | ignore);
    (force-tcp-window-checks | no-force-tcp-window-checks);
    (ignore-memory-overflow | no-ignore-memory-overflow);
    (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
    ignore-reassembly-overflow;
    max-flow-mem value;
    max-packet-mem-ratio percetnage-value;
```


Copyright © 2016, Juniper Networks, Inc.

```
            max-synacks-queued value;
            (tcp-error-logging | no-tcp-error-logging);
          }
          ssl-inspection {
            cache-prune-chunk-size number;
            key-protection;
            maximum-cache-size number;
            session-id-cache-timeout seconds;
            sessions number;
          }
        }
        traceoptions {
          file {
            filename;
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag all;
          level (all | error | info | notice | verbose | warning);
          no-remote-trace;
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- *Understanding Intrusion Detection and Prevention for SRX Series*

[edit security ike] Hierarchy Level

```
security {
  ike {
    gateway gateway-name {
      address [ip-address-or-hostname];
      dead-peer-detection {
        (always-send | optimized | probe-idle-tunnel);
        interval seconds;
        threshold number;
      }
      dynamic {
        connections-limit number;
        (distinguished-name <container container-string> <wildcard wildcard-string> |
          hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
          e-mail-address);
        ike-user-type (group-ike-id | shared-ike-id);
      }
      external-interface external-interface-name;
      general-ikeid;
      ike-policy policy-name;
      local-address (ipv4-address | ipv6-address);
      local-identity {
        (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address
          | user-at-hostname e-mail-address);
```

```
        }
        nat-keepalive seconds;
        no-nat-traversal;
        remote-identity {
          (distinguished-name <container container-string> <wildcard wildcard-string> |
            hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
            e-mail-address);
        }
        version (v1-only | v2-only);
        xauth {
          access-profile profile-name;
        }
      }
      policy policy-name {
        certificate {
          local-certificate certificate-id;
          peer-certificate-type (pkcs7 | x509-signature);
        }
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
        proposals [proposal-name];
      }
      proposal proposal-name {
        authentication-algorithm (md5 | sha-256 | sha-384| sha1);
        authentication-method (dsa-signatures | ecdsa-signatures-256 |
          ecdsa-signatures-384 | pre-shared-keys | rsa-signatures);
        description description;
        dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
        encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
        lifetime-seconds seconds;
      }
      respond-bad-spi <max-responses>;
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
        rate-limit messages-per-second;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

[edit security ipsec] Hierarchy Level

```
security {
  ipsec {
    internal {
      security-association {
        manual encryption {
          iked_encryption enabled;
          algorithm 3des-cbc;
          key ascii-text key;
          }
      }
    }
    policy policy-name {
      description description;
      perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24
        | group5);
      proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
      proposals [proposal-name];
    }
    proposal proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
      description description;
      encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc |
        aes-192-gcm | aes-256-cbc | aes-256-gcm | des-cbc);
      lifetime-kilobytes kilobytes;
      lifetime-seconds seconds;
      protocol (ah | esp);
    }
    security-association sa-name {
      manual {
        direction bidirectional {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key {
              ascii-text key;
              hexadecimal key;
            }
          }
          auxiliary-spi auxiliary-spi-value;
          encryption {
            algorithm (3des-cbc | des-cbc | null);
            key {
              ascii-text key;
              hexadecimal key;
            }
          }
          protocol (ah | esp);
          spi spi-value;
        }
      }
      mode transport;
    }
    traceoptions {
```

Copyright © 2016, Juniper Networks, Inc.

```
    flag flag;
  }
  vpn vpn-name {
    bind-interface interface-name;
    copy-outer-dscp;
    df-bit (clear | copy | set);
    establish-tunnels (immediately | on-traffic);
    ike {
      gateway gateway-name;
      idle-time seconds;
      install-interval seconds;
      ipsec-policy ipsec-policy-name;
      no-anti-replay;
      proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
      }
    }
    manual {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      encryption {
        algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
        key (ascii-text key | hexadecimal key);
      }
      external-interface external-interface-name;
      gateway ip-address;
      protocol (ah | esp);
      spi spi-value;
    }
    traffic-selector traffic-selector-name {
      local-ip ip-address/netmask;
      remote-ip ip-address/netmask;
    }
    vpn-monitor {
      destination-ip ip-address;
      optimized;
      source-interface interface-name;
    }
  }
  vpn-monitor-options {
    interval seconds;
    threshold number;
  }
 }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

[edit security log] Hierarchy Level

```
security {
  log {
    cache {
      exclude exclude-name {
        destination-address destination-address;
        destination-port destination-port;
        event-id event-id;
        failure;
        interface-name interface-name;
        policy-name policy-name;
        process process-name;
        protocol protocol;
        source-address source-address;
        source-port source-port;
        success;
        user-name user-name;
      }
      limit value;
    }
    disable;
    event-rate rate;
    facility-override;
    file {
      files max-file-number;
      name file-name;
      path binary-log-file-path;
      size maximum-file-size;
    }
    format (binary | sd-syslog | syslog);
    mode (event | stream);
    source-address source-address | source-interface interface-name;
    stream stream-name {
      category(all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec
        | idp | rtlog | pst-ds-lite | appqos | secintel);
      filter {
        threat-attack;
      }
      format (binary | sd-syslog | syslog | welf);
      host {
        ip-address;
        port port-number;
      }
      severity (alert | critical | debug | emergency | error | info | notice | warning);
    }
    traceoptions {
      file {
        file-name;
        files max-file-number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
        flag flag;
        no-remote-trace;
      }
      transport {
        protocol (udp | tcp | tls);
        tls-profile tls-profile-name;
        tcp-connections tcp-connections;
      }
      utc-time-stamp;
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- SSL Proxy Overview on page 531

[edit security nat] Hierarchy Level

```
security {
  nat {
    destination {
      pool pool-name {
        address <ip-address> {
          (port port-number | to ip-address);
        }
        description text;
        routing-instance (routing-instance-name | default);
      }
      rule-set rule-set-name {
        description text;
        from {
          interface [interface-name];
          routing-instance [routing-instance-name];
          zone [zone-name];
        }
        rule rule-name {
          description text;
          match {
            application {
              [application];
              any;
            }
            (destination-address ip-address | destination-address-name address-name);
            destination-port (port-or-low <to high>);
            protocol [protocol-name-or-number];
            source-address [ip-address];
            source-address-name [address-name];
          }
          then {
            destination-nat (off | pool pool-name | rule-session-count-alarm
              (clear-threshold value | raise-threshold value));
          }
        }
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
proxy-arp interface interface-name address ip-address;
  to ip-address;
}
proxy-ndp interface interface-name address ip-address;
  to ip-address;
}
source {
  address-persistent;
  interface (port-overloading off | port-overloading-factor number);
  pool pool-name {
    address ip-address {
      to ip-address;
    }
    address-persistent subscriber ipv6-prefix-length prefix-length;
    address-pooling (paired | no-paired);
    address-shared;
    description text;
    host-address-base ip-address;
    overflow-pool (pool-name | interface);
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port {
      block-allocation {
        active-block-timeouttimeout-interval;
        block-size block-size;
        log disable;
        maximum-blocks-per-host maximum-block-number
      }
      deterministic {
        block-size block-size;
        host {
          address ip-address;
          address-name address-name;
      }
        no-translation;
        port-overloading-factor number;
        range {
          port-low <to port-high>;
          to port-high;
          twin-port port-low <to port-high>;
        }
    }
    routing-instance routing-instance-name;
  }
  pool-default-port-range lower-port-range to upper-port-range;
  pool-default-twin-port-range lower-port-range to upper-port-range;
  pool-utilization-alarm (clear-threshold value | raise-threshold value);
  port-randomization disable;
  port-round-robin disable;
  rule-set rule-set-name {
    description text;
    from {
      interface [interface-name];
      routing-instance [routing-instance-name];
      zone [zone-name];
    }
    rule rule-name {
```

```
        description text;
        match {
          application {
            [application];
            any;
          }
          (destination-address <ip-address> | destination-address-name
            <address-name>);
          destination-port (port-or-low <to high>);
          protocol [protocol-name-or-number];
          source-address [ip-address];
          source-address-name [address-name];
          source-port (port-or-low <to high>);
        }
        then source-nat;
          interface {
            persistent-nat {
              address-mapping;
              inactivity-timeout seconds;
              max-session-number value;
              permit (any-remote-host | target-host | target-host-port);
            }
          off;
          pool <pool-name>
            persistent-nat
              address-mapping;
              inactivity-timeout seconds;
              max-session-number number;
              permit (any-remote-host | target-host | target-host-port);
            }
          rule-session-count-alarm (clear-threshold value | raise-threshold value);
        }
      }
      to {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
    }
}
static rule-set rule-set-name;
      description text;
      from {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
      rule rule-name {
        description text;
        match {
          (destination-address <ip-address> | destination-address-name
            <address-name>);
          destination-port (port-or-low  | <to high>);
          source-address [ip-address];
          source-address-name [address-name];
          source-port (port-or-low  <to high>);
```


Copyright © 2016, Juniper Networks, Inc.

```
                }
                then static-nat;
                  inet {
                    routing-instance (routing-instance-name | default);
                  }
                  nptv6-prefix {
                    address-prefix;
                      routing-instance routing-instance-name;
                  nptv6-prefix-name {
                    address-prefix-name;
                      routing-instance routing-instance-name;
                  [
                  prefix {
                    address-prefix;
                    mapped-port lower-port-range to upper-port-range;
                    routing-instance (routing-instance-name| default);
                  }
                  prefix-name {
                    address-prefix-name;
                    mapped-port lower-port-range to upper-port-range;
                    routing-instance (routing-instance-name | default);
                  }
                  rule-session-count-alarm (clear-threshold value | raise-threshold value);
                  }
                }
              }
            }
          }
          traceoptions {
            file {
              filename;
              files number;
              match regular-expression;
              (world-readable | no-world-readable);
              size maximum-file-size;
            }
            flag flag;
            no-remote-trace;
          }
        }
      }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Introduction to NAT on page 4839

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
```

```
description description;
match {
  application {
    [application];
    any;
  }
  destination-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
    }
  destination-address-excluded;
  source-address {
    [address];
    any;
    any-ipv4;
    any-ipv6;
    }
  source-address-excluded;
  source-identity {
    [role-name];
    any;
    authenticated-user;
    unauthenticated-user;
    unknown-user;
  }
}
scheduler-name scheduler-name;
then {
  count {
    alarm {
      per-minute-threshold number;
      per-second-threshold number;
    }
  }
  deny;
  log {
    session-close;
    session-init;
  }
  permit {
    application-services {
      application-firewall {
        rule-set rule-set-name;
      }
      application-traffic-control {
        rule-set rule-set-name;
      }
      gprs-gtp-profile profile-name;
      gprs-sctp-profile profile-name;
      idp;
      redirect-wx | reverse-redirect-wx;
      ssl-proxy {
        profile-name profile-name;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
  global {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        from-zone {
```

```
      [zone-name];
      any;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
    to-zone {
      [zone-name];
      any;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
          captive-portal captive-portal;
        }
        utm-policy policy-name;
      }
      destination-address {
        drop-translated;
```


Copyright © 2016, Juniper Networks, Inc.

```
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
        }
        reject;
      }
    }
  }
  policy-rematch;
  policy-stats {
    system-wide (disable | enable);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
      size maximum-file-size;
    }
    flag flag;
    no-remote-trace;
  }
 }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

Copyright © 2016, Juniper Networks, Inc.

[edit security screen] Hierarchy Level

```
security {
  screen {
    ids-option screen-name {
        alarm-without-drop;
        description text;
        icmp {
          flood {
            threshold number;
          }
          fragment;
          icmpv6-malformed;
          ip-sweep {
            threshold number;
          }
          large;
          ping-death;
        }
        ip {
          bad-option;
          block-frag;
          ipv6-extension-header {
            AH-header;
            ESP-header;
            HIP-header;
            destination-header {
              ILNP-nonce-option;
              home-address-option;
              line-identification-option;
              tunnel-encapsulation-limit-option;
              user-defined-option-type <type-low> to <type-high>;
            }
            fragment-header;
            hop-by-hop-header {
              CALIPSO-option;
              RPL-option;
              SFM-DPD-option;
              jumbo-payload-option;
              quick-start-option;
              router-alert-option;
              user-defined-option-type <type-low> to <type-high>;
            }
            mobility-header;
            no-next-header;
            routing-header;
            shim6-header
            user-defined-option-type <type-low> to <type-high>;
          }
          ipv6-extension-header-limit limit;
          ipv6-malformed-header;
          loose-source-route-option;
          record-route-option;
          security-option;
```


Copyright © 2016, Juniper Networks, Inc.

```
    source-route-option;
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
      gre {
        gre-4in4;
        gre-4in6;
        gre-6in4;
        gre-6in6;
      }
      ip-in-udp {
        teredo;
      }
      ipip {
        ipip-4in4;
        ipip-4in6;
        ipip-6in4;
        ipip-6in6;
        ipip-6over4;
        ipip-6to4relay;
        isatap;
        dslite;
      }
      bad-inner-header;
    }
  }
  limit-session {
    destination-ip-based number;
    source-ip-based number;
  }
  tcp {
    fin-no-ack;
    land;
    port-scan {
      threshold number;
    }
    syn-ack-ack-proxy {
      threshold number;
    }
    syn-fin;
    syn-flood {
      alarm-threshold number;
      attack-threshold number;
      destination-threshold number;
      source-threshold number;
      timeout seconds;
      white-list name {
        destination-address destination-address;
        source-address source-address;
      }
    }
    syn-frag;
```

```
            tcp-no-flag;
            tcp-sweep {
              threshold threshold number;
            }
            winnuke;
          }
          udp {
            flood {
              threshold number;
            }
            udp-sweep {
              threshold threshold number;
            }
          }
        }
      }
      traceoptions {
        file filename {
          files number;
          match regular-expression;
          (no-world-readable | world-readable);
          size maximum-file-size;
        }
        flag flag;
        no-remote-trace;
      }
      trap {
        interval trap interval;
      }
    }
  }
```

Related Documentation

- Attack Detection and Prevention Overview on page 835
- Example: Configuring Multiple Screening Options on page 849
- Security Configuration Statement Hierarchy on page 603

[edit security softwires] Hierarchy Level

```
security {
  softwires {
    softwire-name name {
      softwire-concentrator ipv6-address;
      softwire-type IPv4-in-IPv6;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag (all | configuration | flow);
```


Copyright © 2016, Juniper Networks, Inc.

```
        no-remote-trace;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 3440

[edit security zones] Hierarchy Level

```
security {
  zones {
    functional-zone {
      management {
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
        interfaces interface-name  {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
      }
    }
    security-zone zone-name {
      address-book {
        address address-name {
          ip-prefix {
            description text;
          }
          description text;
          dns-name domain-name {
            ipv4-only;
            ipv6-only;
          }
          range-address lower-limit to upper-limit;
          wildcard-address ipv4-address/wildcard-mask;
        }
        address-set address-set-name {
          address address-name;
          address-set address-set-name;
          description text;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      application-tracking;
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
      interfaces interface-name {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
      tcp-rst;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Security Zones and Interfaces Overview on page 1049

Copyright © 2016, Juniper Networks, Inc.

address-book

Syntax

```
address-book (book-name | global) {
  address address-name {
    ip-prefix {
      description text;
    }
    description text;
    dns-name domain-name {
      ipv4-only;
      ipv6-only;
    }
    range-address lower-limit to upper-limit;
    wildcard-address ipv4-address/wildcard-mask;
  }
  address-set address-set-name {
    address address-name;
    address-set address-set-name;
    description text;
  }
  attach {
    zone zone-name;
  }
  description text;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OSRelease 11.1. Statement moved under the security hierarchy in Junos OS Release 11.2. Support for address range added in Junos OS Release 12.1. The **description** option added in Junos OS Release 12.1.

Description Define entries in the address book. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use them when configuring different features, such as security policies and NAT.

NOTE: IPv6 wildcard address configuration is not supported in this release.

Options

- **address-book *book-name*—**Name of the address book.
- **global—**An address book that is available by default. You can add any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, DNS names, or address range to the global address book. You do not need to attach the global address book to a security zone; entries in the global address book are available to all security zones that are not attached to address books.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Address Books on page 1069
- Understanding Address Sets on page 1071

address-book (System)

Syntax

```
address-book {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level
[edit system security-profile *security-profile-name*]

Release Information
Statement introduced in Junos OS Release 11.2.

Description
Specify the number of address books that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Understanding Address Books on page 1069

Copyright © 2016, Juniper Networks, Inc.

appfw-profile (System)

Syntax

```
appfw-profile {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the application firewall profile quota of a logical system.

Options

- **maximum *amount***—Specify the maximum allowed quota value.

  Range: 0 through 1024

- **reserved *amount***—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

appfw-rule

Syntax

```
appfw-rule {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the number of application firewall rule configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to the master logical system and the user logical systems
- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

 Copyright © 2016, Juniper Networks, Inc.

appfw-rule-set

Syntax

```
appfw-rule-set {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the number of application firewall rule set configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to the master logical system and the user logical systems
- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-firewall

Syntax

```
application-firewall {
  rule-sets rule-set-name {
    default-rule {
      (deny | permit);
    }
    rule rule-name {
      match {
        dynamic-application [system-application];
        dynamic-application-group [system-application-group];
      }
      then {
        (deny | permit);
      }
    }
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure application firewall rule sets with rules defining match criteria and the action to be performed.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

application-tracking

Syntax

```
application-tracking {
  disable;
  (first-update | first-update-interval first-update-interval);
  session-update-interval session-update-interval;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 10.2. Support for **disable** added in Junos OS Release 11.4.

Description AppTrack, an application tracking tool, is a form of statistical profiling. Enabling this feature for a zone logs flow statistics (the byte count, packet count, and start and end times for a session) at session end. You can modify the logging time and log frequency with command options. Periodically, a network management tool, such as STRM, collects the logged statistics sent by each network device for bandwidth usage analysis of the network.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring AppTrack on page 574

Copyright © 2016, Juniper Networks, Inc.

auth-entry

Syntax

```
auth-entry {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of firewall authentication entries that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253

 Copyright © 2016, Juniper Networks, Inc.

cluster (Chassis)

Syntax

```
cluster {
  configuration-synchronize {
    no-secondary-bootup-auto;
  }
  control-link-recovery;
  heartbeat-interval milliseconds;
  heartbeat-threshold number;
  network-management {
    cluster-master;
  }
  redundancy-group  group-number {
    gratuitous-arp-count number;
    hold-down-interval number;
    interface-monitor interface-name {
      weight number;
    }
    ip-monitoring {
      family {
        inet {
          ipv4-address {
            interface {
              logical-interface-name;
              secondary-ip-address ip-address;
            }
            weight number;
          }
        }
      }
      global-threshold number;
      global-weight number;
      retry-count number;
      retry-interval seconds;
    }
    node (0 | 1 ) {
      priority number;
    }
    preempt;
  }
  reth-count number;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (world-readable | no-world-readable);
      size maximum-file-size;
    }
    flag flag;
    level {
      (alert | all | critical | debug | emergency | error | info | notice | warning);
    }
    no-remote-trace;
```

Copyright © 2016, Juniper Networks, Inc.

```
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure a chassis cluster.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *ip-monitoring*

Copyright © 2016, Juniper Networks, Inc.

cpu

Syntax

```
cpu {
  reserved percent;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the percentage of CPU utilization that is always available to a logical system. This value is configured in a security profile that is bound to a logical system.

Only the master administrator can create security profiles and bind them to logical systems.

NOTE: The cpu-control option at the [edit system security-profile resources] hierarchy level must be specified for the reserved value to take effect.

Options

reserved *percent*—A reserved quota that guarantees that the percentage of CPU specified is always available to the logical system.

Range: 0 through 100 percent (decimal point allowed).

Default: 1 percent for the master logical system and 0 percent for user logical systems.

CAUTION: The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota as traffic loss could occur.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253

Copyright © 2016, Juniper Networks, Inc.

datapath-debug

Syntax

```
datapath-debug {
  action-profile profile-name {
    event (jexec | lbt | lt-enter | lt-leave | mac-egress | mac-ingress | np-egress | np-ingress
      | pot) {
      count;
      packet-dump;
      packet-summary;
      trace;
    }
    module {
      flow {
        flag {
          all;
        }
      }
    }
    preserve-trace-order;
    record-pic-history;
  }
  capture-file {
    filename;
    files number;
    format pacp-format;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  maximum-capture-size value;
  packet-filter packet-filter-name {
    action-profile (profile-name | default);
    destination-port (port-range | protocol-name);
    destination-prefix destination-prefix;
    interface logical-interface-name;
    protocol (protocol-number | protocol-name;
    source-port (port-range | protocol- name);
    source-prefix source-prefix;
    }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Release Information Command introduced in Junos OS Release 10.0.

Copyright © 2016, Juniper Networks, Inc.

Description Configure the data path debugging options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Data Path Debugging for Logical Systems on page 3476
- *Packet Capture Overview*
- *Understanding Data Path Debugging for SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

dslite-softwire-initiator

Syntax

```
dslite-softwire-initiator {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the number of IPv6 dual-stack lite (DS-Lite) softwire initiators that can connect to the softwire concentrator configured in either a user logical system or the master logical system. This statement is configured in the security profile that is bound to the logical system.

Only the master administrator can create security profiles and bind them to logical systems. The master administrator:

- Uses security profiles to provision logical systems with resources
- Binds security profiles to user logical systems and the master logical system
- Configures more than one security profile, specifying different amounts of resource allocations in various profiles

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources. The default is the system maximum.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system. The default is 0.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 3440

Copyright © 2016, Juniper Networks, Inc.

file (System Logging)

Syntax

```
file filename {
  allow-duplicates;
  any (alert | any | critical | emergency | error | info | none | notice | warning);
  archive {
    archive-sites {
      url password;
    }
    (binary-data | no-binary-data);
    files number;
    size size;
    start-time start-time;
    transfer-interval transfer-interval;
    (world-readable | no-world-readable);
  }
  authorization (alert | any | critical | emergency | error | info | none | notice | warning);
  change-log (alert | any | critical | emergency | error | info | none | notice | warning);
  conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
  daemon (alert | any | critical | emergency | error | info | none | notice | warning);
  dfc (alert | any | critical | emergency | error | info | none | notice | warning);
  explicit-priority;
  external (alert | any | critical | emergency | error | info | none | notice | warning);
  firewall (alert | any | critical | emergency | error | info | none | notice | warning);
  ftp (alert | any | critical | emergency | error | info | none | notice | warning);
  interactive-commands (alert | any | critical | emergency | error | info | none | notice | warning);
  kernel (alert | any | critical | emergency | error | info | none | notice | warning);
  match "regular-expression";
  ntp (alert | any | critical | emergency | error | info | none | notice | warning);
  pfe (alert | any | critical | emergency | error | info | none | notice | warning);
  security (alert | any | critical | emergency | error | info | none | notice | warning);
  structured-data {
    brief;
  }
  user (alert | any | critical | emergency | error | info | none | notice | warning);
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 12.1X47 for SRX Series.

Description Specify the file in which to log data.

Options

- *filename*—Specify the name of the file in which to log data.
- *allow-duplicates*—Do not suppress the repeated messages.
- *any*—Specify all facilities information.
  - *alert*—Specify the conditions that should be corrected immediately.
  - *critical*—Specify the critical conditions.
  - *emergency*—Specify the conditions that cause security functions to stop.
  - *error*—Specify the general error conditions.

Copyright © 2016, Juniper Networks, Inc.

    - *info*—Specify the information about normal security operations.
    - *none*—Do not specify any messages.
    - *notice*—Specify the conditions that should be handled specifically.
    - *warning*—Specify the general warning conditions.
- *archive*—Specify the archive file information.
    - *archive-sites*—Specify a list of destination URLs for the archived log files.
        - *url*—Specify the primary and failover URLs to receive archive files.
    - *binary-data*—Mark file such that it contains binary data.
    - *no-binary-data*—Do not mark the file such that it contains binary data.
    - *files*—Specify the number of files to be archived. Range: 1 through 1000 files.
    - *size*—Specify the size of files to be archived. Range: 65,536 through 1,073,741,824 bytes.
    - *world-readable*—Allow any user to read the log file.
    - *no-world-readable*—Do not allow any user to read the log file.
    - *start-time*—Specify the start time for file transmission. Enter the start time in the yyyy-mm-dd.hh:mm format.
    - *transfer-interval*—Specify the frequency at which to transfer the files to archive sites.
- *authorization*—Specify the authorization system.
- *change-log*—Specify the configuration change log.
- *conflict-log*—Specify the configuration conflict log.
- *daemon*—Specify the various system processes.
- *dfc*—Specify the dynamic flow capture.
- *explicit-priority*—Include the priority and facility in messages.
- *external*—Specify the local external applications.
- *firewall*—Specify the firewall filtering system.
- *ftp*—Specify the FTP process.
- *interactive-commands*—Specify the commands executed by the UI.
- *kernel*—Specify the kernel information.
- *match*—Specify the regular expression for lines to be logged.
- *ntp*—Specify the NTP process.
- *pfe*—Specify the Packet Forwarding Engine.
- *security*—Specify the security-related information.

Copyright © 2016, Juniper Networks, Inc.

- *structured-data*—Log the messages in structured log format.
  - *brief*—Omit English language text from the end of the logged message.
- *user*—Specify the user processes.
  - *info*—Specify the informational messages.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Junos OS System Log Overview*
- *syslog (System)*

firewall-authentication (Security)

Syntax

```
firewall-authentication {
  traceoptions {
    flag flag;
  }
}
```

Hierarchy Level
[edit security]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Define data-plane firewall authentication tracing options.

Options
- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple flag statements.
  - **all**—Enable all tracing operations.
  - **authentication**—Trace data-plane firewall authentication events.
  - **proxy**—Trace data-plane firewall authentication proxy events.
- **detail**—Display moderate amount of data.
- **extensive**—Display extensive amount of data.
- **terse**—Display minimum amount of data.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Understanding Logical System Firewall Authentication on page 3267

Copyright © 2016, Juniper Networks, Inc.

flow (Security Flow)

Syntax

```
flow {
    aging {
        early-ageout seconds;
        high-watermark percent;
        low-watermark percent;
    }
    allow-dns-reply;
    ethernet-switching {
        block-non-ip-all;
        bpdu-vlan-flooding;
        bypass-non-ip-unicast;
        no-packet-flooding {
            no-trace-route;
        }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
    load distribution {
        session-affinity ipsec;
    }
    pending-sess-queue-length (high | moderate | normal);
    route-change-timeout seconds;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    tcp-mss {
        all-tcp mss value;
        gre-in {
            mss value;
        }
        gre-out {
            mss value;
        }
        ipsec-vpn {
            mss value;
        }
    }
    tcp-session {
        fin-invalidate-session;
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
            (session-ageout | session-timeout seconds);
        }
    }
    traceoptions {
        file {
            filename;
            files number;
```


Copyright © 2016, Juniper Networks, Inc.

```
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
      packet-filter filter-name {
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
      }
      rate-limit messages-per-second;
    }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.5.

Description Determine how the device manages packet flow. The device can regulate packet flow in the following ways:

- Enable or disable DNS replies when there is no matching DNS request.
- Set the initial session-timeout values.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Understanding Session Characteristics for SRX Series Services Gateways on page 1689
- Understanding Flow in Logical Systems for SRX Series Devices on page 3211

Copyright © 2016, Juniper Networks, Inc.

flow-gate

Syntax

```
flow-gate {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of flow gates, also known as **pinholes**, that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253

Copyright © 2016, Juniper Networks, Inc.

flow-session

Syntax

```
flow-session {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Release 11.2 of Junos OS.

Description Specify the number of flow sessions that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

NOTE: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

Copyright © 2016, Juniper Networks, Inc.

idp (Security)

Syntax

```
idp {
  active-policy policy-name;
  custom-attack attack-name {
    attack-type {
      anomaly {
        direction (any | client-to-server | server-to-client);
        service service-name;
        shellcode (all | intel | no-shellcode | sparc);
        test test-condition;
      }
      chain {
        expression boolean-expression;
        member member-name {
          attack-type {
            (anomaly ...same statements as in [edit security idp custom-attack attack-name
              attack-type anomaly] hierarchy level | signature ...same statements as in [edit
              security idp custom-attack attack-name attack-type signature] hierarchy
              level);
          }
        }
        order;
        protocol-binding {
          application application-name;
          icmp;
          icmpv6;
          ip {
            protocol-number transport-layer-protocol-number;
          }
          ipv6 {
            protocol-number transport-layer-protocol-number;
          }
          rpc {
            program-number rpc-program-number;
          }
          tcp {
            minimum-port port-number <maximum-port port-number>;
          }
          udp {
            minimum-port port-number <maximum-port port-number>;
          }
        }
        reset;
        scope (session | transaction);
      }
      signature {
        context context-name;
        direction (any | client-to-server | server-to-client);
        negate;
        pattern signature-pattern;
        protocol {
          icmp {
            code {
```

Copyright © 2016, Juniper Networks, Inc.

```
        match (equal | greater-than | less-than | not-equal);
        value code-value;
      }
      data-length {
        match (equal | greater-than | less-than | not-equal);
        value data-length;
      }
      identification {
        match (equal | greater-than | less-than | not-equal);
        value identification-value;
      }
      sequence-number {
        match (equal | greater-than | less-than | not-equal);
        value sequence-number;
      }
      type {
        match (equal | greater-than | less-than | not-equal);
        value type-value;
      }
    }
    ipv4 {
      destination {
        match (equal | greater-than | less-than | not-equal);
        value ip-address-or-hostname;
      }
      identification {
        match (equal | greater-than | less-than | not-equal);
        value identification-value;
      }
      ip-flags {
        (df | no-df);
        (mf | no-mf);
        (rb | no-rb);
      }
      protocol {
        match (equal | greater-than | less-than | not-equal);
        value transport-layer-protocol-id;
      }
      source {
        match (equal | greater-than | less-than | not-equal);
        value ip-address-or-hostname;
      }
      tos {
        match (equal | greater-than | less-than | not-equal);
        value type-of-service-in-decimal;
      }
      total-length {
        match (equal | greater-than | less-than | not-equal);
        value total-length-of-ip-datagram;
      }
      ttl {
        match (equal | greater-than | less-than | not-equal);
        value time-to-live;
      }
    }
    ipv6 {
```


Copyright © 2016, Juniper Networks, Inc.

```
    destination {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    flow-label {
      match (equal | greater-than | less-than | not-equal);
      value flow-label-value;
    }
    hop-limit {
      match (equal | greater-than | less-than | not-equal);
      value hop-limit-value;
    }
    next-header {
      match (equal | greater-than | less-than | not-equal);
      value next-header-value;
    }
    payload-length {
      match (equal | greater-than | less-than | not-equal);
      value payload-length-value;
    }
    source {
      match (equal | greater-than | less-than | not-equal);
      value ip-address-or-hostname;
    }
    traffic-class {
      match (equal | greater-than | less-than | not-equal);
      value traffic-class-value;
    }
  tcp {
    ack-number {
      match (equal | greater-than | less-than | not-equal);
      value acknowledgement-number;
    }
    data-length {
      match (equal | greater-than | less-than | not-equal);
      value tcp-data-length;
    }
    destination-port {
      match (equal | greater-than | less-than | not-equal);
      value destination-port;
    }
    header-length {
      match (equal | greater-than | less-than | not-equal);
      value header-length;
    }
    mss {
      match (equal | greater-than | less-than | not-equal);
      value maximum-segment-size;
    }
    option {
      match (equal | greater-than | less-than | not-equal);
      value tcp-option;
    }
    sequence-number {
      match (equal | greater-than | less-than | not-equal);
      value sequence-number;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        source-port {
            match (equal | greater-than | less-than | not-equal);
            value source-port;
        }
        tcp-flags {
            (ack | no-ack);
            (fin | no-fin);
            (psh | no-psh);
            (r1 | no-r1);
            (r2 | no-r2);
            (rst | no-rst);
            (syn | no-syn);
            (urg | no-urg);
        }
        urgent-pointer {
            match (equal | greater-than | less-than | not-equal);
            value urgent-pointer;
        }
        window-scale {
            match (equal | greater-than | less-than | not-equal);
            value window-scale-factor;
        }
        window-size {
            match (equal | greater-than | less-than | not-equal);
            value window-size;
        }
    }
    udp {
        data-length {
            match (equal | greater-than | less-than | not-equal);
            value data-length;
        }
        destination-port {
            match (equal | greater-than | less-than | not-equal);
            value destination-port;
        }
        source-port {
            match (equal | greater-than | less-than | not-equal);
            value source-port;
        }
    }
}
protocol-binding {
    application application-name;
    icmp;
    icmpv6;
    ip {
        protocol-number transport-layer-protocol-number;
    }
    ipv6 {
        protocol-number transport-layer-protocol-number;
    }
    rpc {
        program-number rpc-program-number;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
            tcp {
              minimum-port port-number <maximum-port port-number>;
            }
            udp {
              minimum-port port-number <maximum-port port-number>;
            }
          }
          regexp regular-expression;
          shellcode (all | intel | no-shellcode | sparc);
        }
      }
      recommended-action (close | close-client | close-server | drop | drop-packet | ignore |
        none);
      severity (critical | info | major | minor | warning);
      time-binding {
        count count-value;
        scope (destination | peer | source);
      }
    }
    custom-attack-group custom-attack-group-name {
      group-members [attack-or-attack-group-name];
    }
    dynamic-attack-group dynamic-attack-group-name {
      filters {
        category {
          values [category-value];
        }
        direction {
          expression (and | or);
          values [any client-to-server exclude-any exclude-client-to-server
            exclude-server-to-client server-to-client];
        }
        false-positives {
          values [frequently occasionally rarely unknown];
        }
        performance {
          values [fast normal slow unknown];
        }
        products {
          values [product-value];
        }
        recommended;
        service {
          values [service-value];
        }
        severity {
          values [critical info major minor warning];
        }
        type {
          values [anomaly signature];
        }
      }
    }
    idp-policy policy-name {
      rulebase-exempt {
        rule rule-name {
```

```
    description text;
    match {
      attacks {
        custom-attack-groups [attack-group-name];
        custom-attacks [attack-name];
        dynamic-attack-groups [attack-group-name];
        predefined-attack-groups [attack-group-name];
        predefined-attacks [attack-name];
      }
      destination-address ([address-name] | any | any-ipv4 | any-ipv6);
      destination-except [address-name];
      from-zone (zone-name | any );
      source-address ([address-name] | any | any-ipv4 | any-ipv6);
      source-except [address-name];
      to-zone (zone-name | any);
    }
  }
}
rulebase-ips {
  rule rule-name {
    description text;
    match {
      application (application-name | any | default);
      attacks {
        custom-attack-groups [attack-group-name];
        custom-attacks [attack-name];
        dynamic-attack-groups [attack-group-name];
        predefined-attack-groups [attack-group-name];
        predefined-attacks [attack-name];
      }
      destination-address ([address-name] | any | any-ipv4 | any-ipv6);
      destination-except [address-name];
      from-zone (zone-name | any );
      source-address ([address-name] | any | any-ipv4 | any-ipv6);
      source-except [address-name];
      to-zone (zone-name | any);
    }
    terminal;
    then {
      action {
        class-of-service {
          dscp-code-point number;
          forwarding-class forwarding-class;
        }
        (close-client | close-client-and-server | close-server |drop-connection |
          drop-packet | ignore-connection | mark-diffserv value | no-action |
          recommended);
      }
      ip-action {
        (ip-block | ip-close | ip-notify);
        log;
        log-create;
        refresh-timeout;
        target (destination-address | service | source-address | source-zone |
          source-zone-address | zone-service);
        timeout seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            notification {
              log-attacks {
                alert;
              }
              packet-log {
                post-attack number;
                post-attack-timeout seconds;
                pre-attack number;
              }
            }
            severity (critical | info | major | minor | warning);
          }
        }
      }
    }
    security-package {
      automatic {
        download-timeout minutes;
        enable;
        interval hours;
        start-time start-time;
      }
      install {
        ignore-version-check;
      }
      source-address address;
      url url-name;
    }
    sensor-configuration {
      application-identification {
        max-packet-memory value;
        max-tcp-session-packet-memory value;
        max-udp-session-packet-memory value;
      }
      detector {
        protocol-name protocol-name {
          tunable-name  tunable-name {
            tunable-value protocol-value;
          }
        }
      }
      flow {
        (allow-icmp-without-flow | no-allow-icmp-without-flow);
        fifo-max-size value;
        hash-table-size value;
        (log-errors | no-log-errors);
        max-session-offset value;
        max-timers-poll-ticks value;
        reject-timeout value;
        (reset-on-policy | no-reset-on-policy);
        udp-anticipated-timeout value;
      }
      global {
        (enable-all-qmodules | no-enable-all-qmodules);
        (enable-packet-pool | no-enable-packet-pool);
```

```
    gtp (decapsulation | no-decapsulation);
    memory-limit-percent value;
    (policy-lookup-cache | no-policy-lookup-cache);
  }
  high-availability {
    no-policy-cold-synchronization;
  }
  ips {
    content-decompression-max-memory-kb value;
    content-decompression-max-ratio value;
    (detect-shellcode | no-detect-shellcode);
    fifo-max-size value;
    (ignore-regular-expression | no-ignore-regular-expression);
    log-supercede-min minimum-value;
    pre-filter-shellcode;
    (process-ignore-s2c | no-process-ignore-s2c);
    (process-override | no-process-override);
    process-port port-number;
  }
  log {
    cache-size size;
    suppression {
      disable;
      (include-destination-address | no-include-destination-address);
      max-logs-operate value;
      max-time-report value;
      start-log value;
    }
  }
  packet-log {
    host ip-address <port number>;
    max-sessions percentage;
    source-address ip-address;
    total-memory percentage;
  }
  re-assembler {
    action-on-reassembly-failure (drop | drop-session | ignore);
    (force-tcp-window-checks | no-force-tcp-window-checks);
    (ignore-memory-overflow | no-ignore-memory-overflow);
    (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
    ignore-reassembly-overflow;
    max-flow-mem value;
    max-packet-mem value;
    (tcp-error-logging | no-tcp-error-logging);
  }
  ssl-inspection {
    cache-prune-chunk-size number;
    key-protection;
    maximum-cache-size number;
    session-id-cache-timeout seconds;
    sessions number;
  }
}
traceoptions {
  file {
    filename;
```


Copyright © 2016, Juniper Networks, Inc.

```
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag all;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.3. The **expression** option added in Junos OS Release 11.4.

Description Configure Intrusion Detection and Prevention (IDP) to selectively enforce various IDP attack detection and prevention techniques on the network.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- *Understanding Intrusion Detection and Prevention for SRX Series*

idp-policy

Syntax idp-policy *idp-policy-name*;

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the IDP policy for the security profile.

Options *idp-policy-name*—Name of the IDP policy.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Understanding Intrusion Detection and Prevention for SRX Series*

Copyright © 2016, Juniper Networks, Inc.

ike (Security)

Syntax

```
ike {
  gateway gateway-name {
    address [ip-address-or-hostname];
    advpn {
      suggester {
        disable;
      }
      partner {
        connection-limit <number>;
        idle-threshold <packets/sec>;
        idle-time <seconds>;
        disable;
      }
    }
    dead-peer-detection {
      (always-send | optimized | probe-idle-tunnel);
      interval seconds;
      threshold number;
    }
    dynamic {
      connections-limit number;
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
      ike-user-type (group-ike-id | shared-ike-id);
    }
    external-interface external-interface-name;
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
      (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address |
        user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    remote-identity {
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    version (v1-only | v2-only);
    xauth {
      access-profile profile-name;
    }
  }
  policy policy-name {
    certificate {
      local-certificate certificate-id;
      peer-certificate-type (pkcs7 | x509-signature);
      policy-oids [ oid ];
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1);
    authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384
      | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
  }
  respond-bad-spi <max-responses>;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    rate-limit messages-per-second;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The **inet6** option added in Junos OS Release 11.1.

Description Define Internet Key Exchange (IKE) configuration.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- ALG Overview on page 3
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

ipsec (Security)

Syntax

```
ipsec {
  policy policy-name {
    description description;
    perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 |
      group5);
    proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
    description description;
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm
      | aes-256-cbc | aes-256-gcm | des-cbc);
    lifetime-kilobytes kilobytes;
    lifetime-seconds seconds;
    protocol (ah | esp);
  }
  security-association sa-name {
    manual {
      direction bidirectional {
        authentication {
          algorithm (hmac-md5-96 | hmac-sha1-96);
          key {
            ascii-text key;
            hexadecimal key;
          }
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
          algorithm (3des-cbc | des-cbc | null);
          key {
            ascii-text key;
            hexadecimal key;
          }
        }
        protocol (ah | esp);
        spi spi-value;
      }
    }
    mode transport;
  }
  traceoptions {
    flag flag;
  }
  vpn vpn-name {
    bind-interface interface-name;
    copy-outer-dscp;
    ike {
    gateway gateway-name;
    ipsec-policy ipsec-phase2-policy;
    }
    establish-tunnels (immediately | on-traffic);
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    ike {
      gateway gateway-name;
      idle-time seconds;
      install-interval seconds;
      ipsec-policy ipsec-policy-name;
      no-anti-replay;
      proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
      }
    }
    manual {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      encryption {
        algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
        key (ascii-text key | hexadecimal key);
      }
      external-interface external-interface-name;
      gateway ip-address;
      protocol (ah | esp);
      spi spi-value;
    }
    traffic-selector traffic-selector-name {
      local-ip ip-address/netmask;
      remote-ip ip-address/netmask;
    }
    vpn-monitor {
      destination-ip ip-address;
      optimized;
      source-interface interface-name;
    }
  }
  vpn-monitor-options {
    interval seconds;
    threshold number;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5.

Description Define IPsec configuration.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

log (Security)

Syntax

```
log {
  cache {
    exclude exclude-name {
      destination-address destination-address;
      destination-port destination-port;
      event-id event-id;
      failure;
      interface-name interface-name;
      policy-name policy-name;
      process process-name;
      protocol protocol;
      source-address source-address;
      source-port source-port;
      success;
      user-name user-name;
    }
    limit value;
  }
  disable;
  event-rate rate;
  facility-override;
  file {
    files max-file-number;
    name file-name;
    path binary-log-file-path;
    size maximum-file-size;
  }
  format (binary | sd-syslog | syslog);
  mode (event | stream);
  rate-cap rate-cap-value;
  (source-address source-address | source-interface interface-name);
  stream stream-name {
    category (all | content-security);
    format (binary | sd-syslog | syslog | welf);
    host {
      ip-address;
      port port-number;
    }
    severity (alert | critical | debug | emergency | error | info | notice | warning);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
  transport {
```

Copyright © 2016, Juniper Networks, Inc.

```
    protocol (udp | tcp | tls);
    tls-profile tls-profile-name;
    tcp-connections tcp-connections;
  }
  utc-time-stamp;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.2.

Description You can set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

- **disable**—Disable the security logging for the device.
- **event-rate** *rate*—Limits the rate (0 through 1500) at which logs will be streamed per second.
- **rate-cap** *rate-cap-value*—Works with event mode only. Limits the rate (0 through 5000) at which data plane logs will be generated per second.
- **source-address** *source-address*—Specify a source IP address or IP address used when exporting security logs.
- **source-interface** *interface-name*—Specify a source interface name, which is mandatory to configure **stream**.

NOTE: The **source-address** and **source-interface** are alternate values. Using one of the options is mandatory.

- **utc-time-stamp**—Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

logical-system (System Security Profile)

Syntax logical-system *logical-system-name*;

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the user logical system to bind the security profile to.

The master administrator uses security profiles to provision logical systems with resources. You can bind security profiles to user logical systems and the master logical system. The master administrator can configure more than one security profile allocating different amounts of a resource in various ones.

Only the master administrator can create security profiles and bind them to logical systems.

Options ***logical-system-name***—Name of the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

logical-systems (All)

Syntax

```
logical-systems logical-system-name {
    ...
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure a logical system. Only the master administrator can configure a logical system at this hierarchy level.

At this hierarchy level, you can include several of the hierarchies that can be included at the [**edit**] hierarchy level. For descriptions of the applicable statements, see the appropriate hierarchies.

NOTE: The logical-systems configuration statement can only be used by the master administrator.

Options *logical-system-name*—Name of the logical system.

Required Privilege Level
all—To view this statement in the configuration.
all—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

nat

Syntax

```
nat {
    destination {
        pool pool-name {
            address ip-address {
                (port port-number | to ip-address);
            }
            description text;
            routing-instance routing-instance-name;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
                    source-address [ip-address];
                    source-address-name [address-name];
                }
                then {
                    destination-nat (off | pool pool-name);
                }
            }
        }
    }
    proxy-arp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    proxy-ndp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    source {
        address-persistent;
        interface {
            port-overloading {
                off;
            }
        }
```

```
pool pool-name {
  address ip-address {
    to ip-address;
  }
  description text;
  host-address-base ip-address;
  overflow-pool (interface | pool-name);
  port {
    (no-translation | port-overloading-factor number | range port-low <to port-high>);
  }
  routing-instance routing-instance-name;
}
pool-default-port-range lower-port-range to upper-port-range;
pool-utilization-alarm {
  clear-threshold value;
  raise-threshold value;
}
port-randomization {
  disable;
}
port-round-robin {
  disable;
}
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      (destination-address <ip-address> | destination-address-name <address-name>);
      destination-port port-number;
      protocol [protocol-name-or-number];
      source-address [ip-address];
      source-address-name [address-name];
    }
    then {
      source-nat {
        interface {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number value;
            permit (any-remote-host | target-host | target-host-port);
          }
        }
        off;
        pool {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number number;
            permit (any-remote-host | target-host | target-host-port);
```


Copyright © 2016, Juniper Networks, Inc.

```
                    }
                    pool-name;
                  }
                }
              }
            }
            to {
              interface [interface-name];
              routing-instance [routing-instance-name];
              zone [zone-name];
            }
          }
        }
        static {
          rule-set rule-set-name {
            description text;
            from {
              interface [interface-name];
              routing-instance [routing-instance-name];
              zone [zone-name];
            }
            rule rule-name {
              description text;
              match {
                (destination-address ip-address | destination-address-name address-name);
              }
              then {
                static-nat {
                  inet {
                    routing-instance (default | routing-instance-name);
                  }
                  prefix {
                    address-prefix;
                    routing-instance (default | routing-instance-name);
                  }
                  prefix-name {
                    address-prefix-name;
                    routing-instance (default | routing-instance-name);
                  }
                }
              }
            }
          }
        }
        traceoptions {
          file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
        }
      }
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1.

Description Configure Network Address Translation (NAT) for SRX Series devices.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Introduction to NAT on page 4839
- Understanding Logical System Network Address Translation on page 3355

Copyright © 2016, Juniper Networks, Inc.

nat-cone-binding

Syntax

```
nat-cone-binding {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT cone binding configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-destination-pool

Syntax

```
nat-destination-pool {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT destination pool configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-destination-rule

Syntax

```
nat-destination-rule {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT destination rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-interface-port-ol (System)

Syntax

```
nat-interface-port-ol {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the security NAT interface port overloading the quota of a logical system.

Options

- **maximum *amount***—Specify the maximum allowed quota value.

  Range: 0 through 64

- **reserved *amount***—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-nopat-address

Syntax

```
nat-nopat-address {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT without port address translation configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-pat-address

Syntax

```
nat-pat-address {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT with port address translation (PAT) configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Network Address Translation on page 3355
- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

nat-pat-portnum

Syntax

```
nat-pat-portnum {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the maximum quantity and the reserved quantity of ports for the logical system as part of its security profile. The total number of PAT pools must not exceed the configured maximum ports for the logical system.

Options

maximum *amount*—Specify the maximum number of ports allowed for a logical system. The maximum quantity is not guaranteed and is shared among multiple logical systems.

reserved *amount*—Specify the number of resources guaranteed for a logical system.

Range: For SRX5600 and SRX5800 devices, up to 402,653,184 ports are supported. Pool specifications for logical systems can be viewed using the **show security nat source summary logical-system all** command.

Required Privilege Level

system—To view this statement in the configuration.
system—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

nat-port-ol-ipnumber

Syntax

```
nat-port-ol-ipnumber {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT port overloading IP number configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

 Copyright © 2016, Juniper Networks, Inc.

nat-rule-referenced-prefix (System)

Syntax

```
nat-rule-referenced-prefix {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the security NAT rule referenced IP prefix quota of a logical system.

Options

- **maximum *amount*** —Specify the maximum allowed quota value.

  Range: 0 through 1,048,576

- **reserved *amount*** —Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

nat-source-pool

Syntax

```
nat-source-pool {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the NAT source pool configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

nat-source-rule

Syntax

```
nat-source-rule {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the NAT source rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

nat-static-rule

Syntax

```
nat-static-rule {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT static rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

policies

Syntax

```
policies {
  default-policy (deny-all | permit-all);
  from-zone zone-name to-zone zone-name {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
        deny;
        log {
          session-close;
          session-init;
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
}
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
```


Copyright © 2016, Juniper Networks, Inc.

```
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    from-zone {
      [zone-name];
      any;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
    to-zone {
      [zone-name];
      any;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
```

```
                captive-portal captive-portal;
              }
              utm-policy policy-name;
            }
            destination-address {
              drop-translated;
              drop-untranslated;
            }
            firewall-authentication {
              pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
              }
              web-authentication {
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable) ;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Copyright © 2016, Juniper Networks, Inc.

Release Information

Statement introduced in Junos OS Release 8.5.
Support for the **services-offload** option added in Junos OS Release 11.4.
Support for the **source-identitiy** option added in Junos OS Release 12.1.
Support for the **description** option added in Junos OS Release 12.1.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.
Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.
Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description

Configure network security policies.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

policy (System Security Profile)

Syntax

```
policy {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of security policies that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

policy-with-count

Syntax

```
policy-with-count {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of security policies with a count that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order [accounting-method];
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order [accounting-method];
  address-assignment pool pool-name;
  authentication-order [ldap | none | password | radius | securid];
  authorization-order [jsrc];
  client client-name {
    chap-secret chap-secret;
    client-group [ group-names ];
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
```


Copyright © 2016, Juniper Networks, Inc.

```
radius {
  accounting-server [server];
  attributes {
    exclude {
      acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
      acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
      acc-loop-cir-id [access-request | accounting-start | accounting-stop];
      accounting-authentic [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-delay-time [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-session-id [access-request];
      accounting-terminate-cause [accounting-off];
      act-data-rate-dn [access-request | accounting-start | accounting-stop];
      act-data-rate-up [access-request | accounting-start | accounting-stop];
      act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      act-interlv-delay-up [access-request | accounting-start | accounting-stop];
      att-data-rate-dn [access-request | accounting-start | accounting-stop];
      att-data-rate-up [access-request | accounting-start | accounting-stop];
      called-station-id [access-request | accounting-start | accounting-stop];
      calling-station-id [access-request | accounting-start | accounting-stop];
      class [access-request | accounting-start | accounting-stop];
      delegated-ipv6-prefix [accounting-start | accounting-stop];
      dhcp-gi-address [access-request | accounting-start | accounting-stop];
      dhcp-mac-address [access-request | accounting-start | accounting-stop];
      dhcp-options [access-request | accounting-start | accounting-stop];
      downstream-calculated-qos-rate [access-request | accounting-start |
        accounting-stop];
      dsl-forum-attributes [access-request | accounting-start | accounting-stop];
      dsl-line-state [access-request | accounting-start | accounting-stop];
      dsl-type [access-request | accounting-start | accounting-stop];
      dynamic-iflset-name [accounting-start | accounting-stop];
      event-time-stamp [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      framed-interface-id [access-request | accounting-start | accounting-stop];
      framed-ip-address [access-request | accounting-start | accounting-stop];
      framed-ip-netmask [access-request | accounting-start | accounting-stop];
      framed-ip-route [access-request | accounting-start | accounting-stop];
      framed-ipv6-pool [accounting-start | accounting-stop];
      framed-ipv6-prefix [accounting-start | accounting-stop];
      framed-ipv6-route [accounting-start | accounting-stop];
      framed-pool [accounting-start | accounting-stop];
      input-filter [accounting-start | accounting-stop];
      input-gigapackets [accounting-stop];
      input-gigawords [accounting-stop];
      input-ipv6-gigawords [accounting-stop];
      input-ipv6-octets [accounting-stop];
      input-ipv6-packets [accounting-stop];
      interface-description [access-request | accounting-start | accounting-stop];
      l2c-downstream-data [access-request | accounting-start | accounting-stop];
      l2c-upstream-data [access-request | accounting-start | accounting-stop];
      max-data-rate-dn [access-request | accounting-start | accounting-stop];
      max-data-rate-up [access-request | accounting-start | accounting-stop];
      max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      max-interlv-delay-up [access-request | accounting-start | accounting-stop];
      min-data-rate-dn [access-request | accounting-start | accounting-stop];
```

Copyright © 2016, Juniper Networks, Inc.

```
            min-data-rate-up [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
            nas-identifier [access-request | accounting-start | accounting-stop];
            nas-port [access-request | accounting-off | accounting-on | accounting-start |
              accounting-stop];
            nas-port-id [access-request | accounting-start | accounting-stop];
            nas-port-type [access-request | accounting-start | accounting-stop];
            output-filter [accounting-start | accounting-stop];
            output-gigapackets [accounting-stop];
            output-gigawords [accounting-stop];
            output-ipv6-gigawords [accounting-stop];
            output-ipv6-octets [accounting-stop];
            output-ipv6-packets [accounting-stop];
            upstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];
          }
          ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
          }
        }
      authentication-server [server];
      radius-options {
        request-rate number;
        revert-interval seconds;
      }
      radius-server server-address {
        accounting-port port-number
        max-outstanding-requests number-of--outstanding-requests;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
      }
      service {
        accounting-order {
          activation-protocol;
          radius;
        }
      }
      session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
      }
    }
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

 Copyright © 2016, Juniper Networks, Inc.

Description Create a profile containing a set of attributes that define device management access.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441
- Understanding User Authentication for Security Devices on page 5173
- Layer 2 Transparent and Switching Overview on page 2893

purging

Syntax purging;

Hierarchy Level [edit system arp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Purge obsolete ARP entries from the cache when an interface or link goes offline.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

root-authentication

Syntax

```
root-authentication {
    encrypted-password password;
    load-key-file URL;
    plain-text-password;
    ssh-dsa public-key {
    <from pattern-list>;
    }
    ssh-rsa public-key {
    <from pattern-list>;
    }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify authentication information for the root login.

Options

- **encrypted-password *password***—Specify the encrypted authentication password. You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.
- **plain-text-password**—The CLI prompts you for a password encrypts it, and stores the encrypted version in its user database.
- **load-key-file*URL***—File URL containing one or more SSH keys.
- **ssh-dsa *public-key***—SSH DSA public key string.
  - **from *pattern-list***—Pattern list of allowed hosts.
- **ssh-rsa *public-key***—SSH RSA public key string.
  - **from *pattern-list***—Pattern list of allowed hosts.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- System Configuration Statement Hierarchy on page 612

Copyright © 2016, Juniper Networks, Inc.

root-logical-system

Syntax root-logical-system;

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify root-logical-system to bind the security profile to the master logical system.

The master administrator uses security profiles to provision logical systems with resources. The security profile containing this statement must be bound to root-logical-system only.

Only the master administrator can create security profiles and bind them to logical systems.

Options none

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

scheduler (System Security Profile)

Syntax

```
scheduler {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of schedulers that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

screen (Security)

Syntax

```
screen {
    ids-option screen-name {
        alarm-without-drop;
        description text;
        icmp {
            flood {
                threshold number;
            }
            fragment;
            icmpv6-malformed;
            ip-sweep {
                threshold number;
            }
            large;
            ping-death;
        }
        ip {
            bad-option;
            block-frag;
            ipv6-extension-header {
                AH-header;
                ESP-header;
                HIP-header;
                }
                destination-header {
                    ILNP-nonce-option;
                    home-address-option;
                    line-identification-option;
                    tunnel-encapsulation-limit-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                fragment-header;
                hop-by-hop-header {
                    CALIPSO-option;
                    RPL-option;
                    SFM-DPD-option;
                    jumbo-payload-option;
                    quick-start-option;
                    router-alert-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                mobility-header;
                no-next-header;
                routing-header;
                shim6-header
                user-defined-option-type <type-low> to <type-high>;
            }
            ipv6-extension-header-limit limit;
            ipv6-malformed-header;
            loose-source-route-option;
            record-route-option;
            security-option;
```

```
    source-route-option;
    spoofing;
    stream-option;
    strict-source-route-option;
    tear-drop;
    timestamp-option;
    unknown-protocol;
    tunnel {
      gre {
        gre-4in4;
        gre-4in6;
        gre-6in4;
        gre-6in6;
      }
      ip-in-udp {
        teredo;
      }
      ipip {
        ipip-4in4;
        ipip-4in6;
        ipip-6in4;
        ipip-6in6;
        ipip-6over4;
        ipip-6to4relay;
        isatap;
        dslite;
      }
      bad-inner-header;
    }
  }
  limit-session {
    destination-ip-based number;
    source-ip-based number;
  }
  tcp {
    fin-no-ack;
    land;
    port-scan {
      threshold number;
    }
    syn-ack-ack-proxy {
      threshold number;
    }
    syn-fin;
    syn-flood {
      alarm-threshold number;
      attack-threshold number;
      destination-threshold number;
      source-threshold number;
      timeout seconds;
      white-list name {
        destination-address destination-address;
        source-address source-address;
      }
    }
    syn-frag;
```


Copyright © 2016, Juniper Networks, Inc.

```
            tcp-no-flag;
            tcp-sweep {
              threshold threshold number;
            }
            winnuke;
          }
          udp {
            flood {
              threshold number;
            }
            port-scan {
              threshold number;
            }
            udp-sweep {
              threshold threshold number;
            }
          }
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. The **description** option added in Junos OS Release 12.1.

Description Configure security screen options.

Options ***screen-name***—Name of the screen configured at the **security screen ids-options** level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Attack Detection and Prevention Overview on page 835

Copyright © 2016, Juniper Networks, Inc.

security-profile

Syntax

```
security-profile security-profile-name {
    address-book {
      maximum amount;
      reserved amount;
    }
    appfw-profile {
      maximum amount;
      reserved amount;
    }
    appfw-rule {
      maximum amount;
      reserved amount;
    }
    appfw-rule-set {
      maximum amount;
      reserved amount;
    }
    auth-entry {
      maximum amount;
      reserved amount;
    }
    cpu {
      reserved percent;
    }
    dslite-softwire-initiator {
      maximum amount;
      reserved amount;
    }
    flow-gate {
      maximum amount;
      reserved amount;
    }
    flow-session {
      maximum amount;
      reserved amount;
    }
    idp-policy idp-policy-name;
    logical-system [logical-system-name];
    nat-cone-binding {
      maximum amount;
      reserved amount;
    }
    nat-destination-pool {
      maximum amount;
      reserved amount;
    }
    nat-destination-rule {
      maximum amount;
      reserved amount;
    }
    nat-interface-port-ol {
      maximum amount;
```


Copyright © 2016, Juniper Networks, Inc.

```
    reserved amount;
}
nat-nopat-address {
    maximum amount;
    reserved amount;
}
nat-pat-address {
    maximum amount;
    reserved amount;
}
nat-pat-portnum {
    maximum amount
    reserved amount
}
nat-port-ol-ipnumber {
    maximum amount;
    reserved amount;
}
nat-rule-referenced-prefix {
    maximum amount;
    reserved amount;
}
nat-source-pool {
    maximum amount;
    reserved amount;
}
nat-source-rule {
    maximum amount;
    reserved amount;
}
nat-static-rule {
    maximum amount;
    reserved amount;
}
policy {
    maximum amount;
    reserved amount;
}
policy-with-count {
    maximum amount;
    reserved amount;
}
root-logical-system;
scheduler {
    maximum amount;
    reserved amount;
}
zone {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.2. The **dslite-softwire-initiator** option introduced in Junos OS Release 12.1.

Description Create a security profile and specify the kinds and amounts of resources to allocate to a logical system to which the security profile is bound.

As a master administrator, you can create a security profile and bind it to more than one logical system if you want to allocate the same kinds and amounts of resources to them. For details on how many security profiles you can create, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 3253. When you reach the limit, you must delete a security profile and commit the configuration before you can create and commit the configuration for another security profile.

Only the master administrator can create security profiles.

Options

- ***security-profile-name***—Name of the security profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

 Copyright © 2016, Juniper Networks, Inc.

security-profile-resources

Syntax

```
security-profile-resources {
    cpu-control;
    cpu-control-target percent;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure global settings that apply to all logical systems in the device.

Options

cpu-control—Enable CPU utilization control.

cpu-control-target ***percent***—Specify the upper limit for CPU utilization on the device under normal operating conditions.

Range: 0 through 100 percent (decimal point allowed).

Default: 80 percent.

NOTE: **The cpu-control option must be specified for the cpu-control-target value to take effect.**

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

Copyright © 2016, Juniper Networks, Inc.

softwires

Syntax

```
softwires {
  softwire-name name {
    softwire-concentrator ipv6-address;
    softwire-type IPv4-in-IPv6;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag (all | configuration | flow);
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before Junos OS Release 12.1.

Description Configure softwires for IPv6 dual-stack lite (DS-Lite). DS-Lite allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content.

Options

- **softwire-name *name***—Name of the softwire configuration.
  - **softwire-concentrator *ipv6-address***—IPv6 address of the concentrator.
  - **softwire-type**—Must be IPv4-in-IPv6.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [edit security softwires] Hierarchy Level on page 3552

 Copyright © 2016, Juniper Networks, Inc.

zone (System Security Profile)

Syntax

```
zone {
  maximum amount;
  reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the zones that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum *amount***—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
- **reserved *amount***—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 3253
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 3258

Copyright © 2016, Juniper Networks, Inc.

zones

Syntax

```
zones {
  functional-zone {
    management {
      description text;
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
      interfaces interface-name  {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
    }
  }
  security-zone zone-name {
    address-book {
      address address-name {
        ip-prefix {
          description text;
        }
        description text;
        dns-name domain-name {
          ipv4-only;
          ipv6-only;
        }
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
      }
      address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
      }
    }
    application-tracking;
    description text;
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
```


Copyright © 2016, Juniper Networks, Inc.

```
            except;
          }
        }
        interfaces interface-name {
          host-inbound-traffic {
            protocols protocol-name {
              except;
            }
            system-services service-name {
              except;
            }
          }
        }
        screen screen-name;
        tcp-rst;
      }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description A zone is a collection of interfaces for security purposes. All interfaces in a zone are equivalent from a security point of view. Configure the following zones:

- Functional zone—Special-purpose zone, such as a management zone that can host dedicated management interfaces.
- Security zone—Most common type of zone that is used as a building block in policies.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 156

Operational Commands

- clear security application-firewall rule-set statistics logical-system
- clear security dns-cache
- request security datapath-debug capture start
- request security datapath-debug capture stop
- set chassis cluster cluster-id node reboot
- show chassis cluster status
- show log
- show security application-firewall rule-set
- show security application-firewall rule-set logical-system
- show security application-tracking counters
- show security datapath-debug capture
- show security datapath-debug counter
- show security dns-cache
- show security firewall-authentication history
- show security firewall-authentication users
- show security flow session
- show security idp logical-system policy-association
- show security ike security-associations
- show security ipsec security-associations
- show security match-policies
- show security nat destination rule
- show security nat destination summary
- show security nat source rule
- show security nat source summary
- show security nat static rule
- show security policies
- show security screen statistics
- show system security-profile

Copyright © 2016, Juniper Networks, Inc.

- show security softwires
- show security zones

Copyright © 2016, Juniper Networks, Inc.

clear security application-firewall rule-set statistics logical-system

Syntax The master, or root, administrator can issue the following statements:

clear security application-firewall rule-set statistics [logical-system *logical-system-name* | all | root-logical-system]

The user logical system administrator can issue the following statement:

clear security application-firewall rule-set statistics all

Release Information Command introduced in Junos OS Release 11.4.

Description Clear all security application firewall rule set statistics.

NOTE: User logical system administrators can clear statistics only for the logical systems they can access. For information about master and user administrator roles in logical systems, see "Understanding the Master Logical System and the Master Administrator Role" on page 3221.

Options ***logical-system-name***—Name of a specific logical system.

all—(default) Clear all rule set statistics for a specific logical system or all logical systems.

root-logical-system—Clear application firewall rule set statistics on the root logical system (master administrator only).

Required Privilege Level clear

Related Documentation

- show security application-firewall rule-set logical-system on page 777

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security dns-cache

Syntax clear security dns-cache <dns-name *dns-name*>

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Reset DNS cache information.

NOTE: This command is only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options

- **dns-name**—Clear DNS cache information for the specified name.

Required Privilege Level clear

Related Documentation

- show security dns-cache on page 3660
- Understanding the Master Logical System and the Master Administrator Role on page 3221

Copyright © 2016, Juniper Networks, Inc.

request security datapath-debug capture start

Syntax request security datapath-debug capture start

Release Information Command introduced in Junos OS Release 10.0.

Description Start the data path debugging capture.

Required Privilege Level maintenance

Related Documentation
- Understanding Data Path Debugging for Logical Systems on page 3476

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture start

```
user@host> request security datapath-debug capture start
datapath-debug capture started on file
```

Copyright © 2016, Juniper Networks, Inc.

request security datapath-debug capture stop

Syntax request security datapath-debug capture stop

Release Information Command introduced in Junos OS Release 10.0.

Description Stop the data path debugging capture.

Required Privilege Level maintenance

Related Documentation
- Understanding Data Path Debugging for Logical Systems on page 3476

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture stop

```
user@host> request security datapath-debug capture stop
datapath-debug capture successfully stopped, use show security datapath-debug
capture to view
```


Copyright © 2016, Juniper Networks, Inc.

set chassis cluster cluster-id node reboot

Syntax set chassis cluster cluster-id *cluster-id* node *node* reboot

Release Information Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X45-D10.

Description This operational mode command sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. The system uses the chassis cluster ID and chassis cluster node ID to apply the correct configuration for each node (for example, when you use the **apply-groups** command to configure the chassis cluster management interface). The chassis cluster ID and node ID statements are written to the EPROM, and the statements take effect when the system is rebooted.

Setting a cluster ID to 0 is equivalent to disabling a cluster. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back.

NOTE: If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X45-D10 or later and re-create a cluster with cluster IDs greater than 16. If for any reason you decide to revert to the previous version of Junos OS that did not support extended cluster IDs, the system comes up with standalone devices after you reboot. If the cluster ID set is less than 16 and you roll back to a previous release, the system comes back with the previous setup.

Options **cluster-id *cluster-id*** —Identifies the cluster within the Layer 2 domain.

Range: 0 through 255

node *node* —Identifies a node within a cluster.

Range: 0 to 1

Required Privilege Level maintenance

Related Documentation

- *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*
- *Example: Setting the Chassis Cluster Node ID and Cluster ID for High-End SRX Series Devices*
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 3210
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 3370

Output Fields When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

show chassis cluster status

Syntax **show chassis cluster status**
<redundancy-group *group-number* >

Release Information Command modified in Junos OS Release 9.2. Support for dual control ports added in Junos OS Release 10.0. Support for monitoring failures added in Junos OS Release 12.1X47-D10.

Description Display the failover status of a chassis cluster.

Options
- none—Display the status of all redundancy groups in the chassis cluster.
- **redundancy-group *group-number*** —(Optional) Display the status of the specified redundancy group.

Required Privilege Level view

Related Documentation
- *redundancy-group (Chassis Cluster)*
- *clear chassis cluster failover-count*
- *request chassis cluster failover node*
- *request chassis cluster failover reset*

List of Sample Output
show chassis cluster status on page 3645
show chassis cluster status redundancy-group 1 on page 3646

Output Fields Table 353 lists the output fields for the **show chassis cluster status** command. Output fields are listed in the approximate order in which they appear.

Table 353: show chassis cluster status Output Fields

Field Name	Field Description
Cluster ID	ID number (1-15) of a cluster is applicable for releases upto 12.1X45-D10. ID number (1-255) is applicable for releases 12.1X45-D10 and later. Setting a cluster ID to 0 is equivalent to disabling a cluster.
Redundancy-Group	ID number (1-128) of a redundancy group in the chassis cluster.
Node name	Node (device) in the chassis cluster (**node0** or **node1**).
Priority	Assigned priority for the redundancy group on that node.

Copyright © 2016, Juniper Networks, Inc.

Table 353: show chassis cluster status Output Fields *(continued)*

Field Name	Field Description
Status	State of the redundancy group (**Primary**, **Secondary**, **Lost**, or **Unavailable**). • **Primary**—Redundancy group is active and passing traffic. • **Secondary**—Redundancy group is passive and not passing traffic. • **Lost**—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and due to control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • **Unavailable**—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
Preempt	• **Yes**: Mastership can be preempted based on priority. • **No**: Change in priority will not preempt the mastership.
Manual failover	• **Yes**: If the Mastership is set manually through the CLI with the **request chassis cluster failover node** or **request chassis cluster failover redundancy-group** command. This overrides **Priority** and **Preempt**. • **No**: Mastership is not set manually through the CLI.
Monitor-failures	• **None**: Cluster working properly. • **Monitor Failure code**: Cluster is not working properly and the respective failure code is displayed.

Sample Output

Displays chassis cluster status with all redundancy groups.

show chassis cluster status

```
user@host> show chassis cluster status

Monitor Failure codes:
  CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 0 , Failover count: 1
node0  200      primary        no      no       None
node1  1        secondary      no      no       None

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

Displays chassis cluster status with redundancy group 1 only.

show chassis cluster status redundancy-group 1

```
user@host> show chassis cluster status redundancy-group 1

Monitor Failure codes:
    CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node   Priority Status         Preempt Manual   Monitor-failures

Redundancy group: 1 , Failover count: 1
node0  101      primary        no      no       None
node1  1        secondary      no      no       None
```


Copyright © 2016, Juniper Networks, Inc.

show log

List of Syntax

Syntax on page 3647
Syntax (QFX Series and OCX Series) on page 3647
Syntax (TX Matrix Router) on page 3647

Syntax

show log
<*filename* | user <*username*>>

Syntax (QFX Series and OCX Series)

show log *filename*
<*device-type (device-id* | *device-alias*)>

Syntax (TX Matrix Router)

show log
<all-lcc | lcc *number* | scc>
<*filename* | user <*username*>>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Option *device-type (device-id* | *device-alias*) is introduced in Junos OS Release 13.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace ***number*** with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

Copyright © 2016, Juniper Networks, Inc.

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.

NOTE: If you specify the device-type optional parameter, you must also specify either the device-id or device-alias optional parameter.

(***device-id*** **|** ***device-alias***)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

user <*username*>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include ***username***, display logging information about the specified user.

Required Privilege Level trace

Related Documentation
- *syslog (System)*

List of Sample Output
show log on page 3648
show log filename on page 3649
show log filename (QFabric System) on page 3649
show log user on page 3650

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
```


Copyright © 2016, Juniper Networks, Inc.

```
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
```

Copyright © 2016, Juniper Networks, Inc.

```
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user
darius   mg2546                    Thu Oct  1 19:37   still logged in
darius   mg2529                    Thu Oct  1 19:08 - 19:36  (00:28)
darius   mg2518                    Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                    Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    jun.site.per    Wed Sep 30 18:39 - 18:41  (00:02)
alex     ttyp1    192.168.1.2      Wed Sep 30 01:03 - 01:22  (00:19)
```


Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set

Syntax show security application-firewall rule-set (*<rule-set-name>* | all)

Release Information Command introduced in Junos OS Release 11.1. Updated in Junos OS Release 12.1X44-D10 with output format changes. Updated in Junos OS Release 12.1X45-D10 with redirection counters.

Description Display information about the specified rule set defined in the application firewall.

Options ***rule-set-name***—Name of the rule set.

all—Display information about all the application firewall rule sets.

Required Privilege Level view

Related Documentation
- clear security application-firewall rule-set statistics on page 745

List of Sample Output
show security application-firewall rule-set my_ruleset1 on page 3652
show security application-firewall rule-set all on page 3652

Output Fields Table 44 lists the output fields for the **show security application-firewall rule-set** command. Output fields are listed in the approximate order in which they appear.

Table 354: show security application-firewall rule-set Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system of the rule set.
Profile	The redirect profile to be used for rules requiring redirection for reject or deny actions.

Rule	Name of the rule
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **SSL-Encryption**—Setting for SSL traffic.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	  ◦ **reject**
	  ◦ **redirect**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.
	• **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Copyright © 2016, Juniper Networks, Inc.

Table 354: show security application-firewall rule-set Output Fields *(continued)*

Field Name	Field Description
Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule. • **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.
Number of sessions with appid pending	Number of sessions that are pending application identification processing

Sample Output

show security application-firewall rule-set my_ruleset1

```
user@host>show security application-firewall rule-set my_ruleset1
Rule-set: my_ruleset1
     Rule: rule1
         Dynamic Applications: junos:facebook, junos:messenger
         Dynamic Application Groups: junos:web, junos:chat
         SSL-Encryption: any
         Action: deny or redirect
         Number of sessions matched: 10
         Number of sessions redirected: 10
     Default rule: permit
         Number of sessions matched: 200
         Number of sessions redirected: 0
     Number of sessions with appid pending: 2
```

Sample Output

show security application-firewall rule-set all

```
user@host> show security application-firewall rule-set all
Rule-set: appfw
    Logical system: root-logical-system
    Profile: lsy2_pf555
    Rule: 2
        Dynamic Applications: junos:HTTP
        SSL-Encryption: any
        Action:deny or redirect
        Number of sessions matched: 2
        Number of sessions redirected: 2
    Rule: 1
        Dynamic Applications: junos:FTP
        SSL-Encryption: any
        Action:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Default rule:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Number of sessions with appid pending: 0
```


Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set logical-system

Syntax

The master, or root, administrator can issue the following statements:

show security application-firewall rule-set all
show security application-firewall rule-set *rule-set-name* | all | logical-system *logical-system-name* | all | root-logical-system [*logical-system-name* | all]

The user logical system administrator can issue the following statement:

show security application-firewall rule-set all

Release Information

Command introduced in Junos OS Release 11.4.

Description

Display information about application firewall rule set(s) associated with a specific logical system, all logical systems, or the root logical system configured on a device.

NOTE: The master administrator can configure and view application firewall rule sets for the root logical system and all user logical systems configured on the device. User logical system administrators can configure and view application firewall rule set information only for the user logical systems for which they have access. For information about master and user administrator roles in logical systems, see "Understanding Logical Systems for SRX Series Services Gateways" on page 3205.

Options

rule-set-name—Name of a specific rule set.

logical-system-name—Name of a specific logical system.

all—(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.

root-logical-system—Display application firewall rule set information for the root logical system (master administrator only).

Required Privilege Level

view

Related Documentation

- clear security application-firewall rule-set statistics logical-system on page 746

List of Sample Output

show security application-firewall rule-set logical-system all on page 3655
show security application-firewall rule-set all on page 3656

Output Fields

Table 45 lists the output fields for the **show security application-firewall rule-set logical-system** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 355: show security application-firewall rule-set logical-system Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system.

Rule	Name of the rule.
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.

Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule.
Number of sessions with appid pending	Number of sessions that are pending with the application ID processing.

Sample Output

show security application-firewall rule-set logical-system all

```
root@host> show security application-firewall rule-set logical-system all

Rule-set: root_rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FTP
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 10
```

Copyright © 2016, Juniper Networks, Inc.

show security application-firewall rule-set all

```
root@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```


Copyright © 2016, Juniper Networks, Inc.

show security application-tracking counters

Syntax show security application-tracking counters

Release Information Command introduced in Junos OS Release 10.2.

Description Display the status of AppTrack counters.

Required Privilege Level view

Related Documentation
- Example: Configuring AppTrack on page 574

Output Fields Table 46 lists the output fields for the **show security application-tracking counters** command. Output fields are listed in the approximate order in which they appear.

Table 356: show security application-tracking counters

Field Name	Field Description
Session create messages	The number of log messages generated when a session was created.
Session close messages	The number of log messages generated when a session was closed.
Session volume updates	The number of log messages generated when an update interval was exceeded.
Failed messages	The number of messages that were not generated due to memory or session constraints.

Sample Output

show security application-tracking counters

```
user@host> show security application-tracking counters
AVT counters:                       Value
 Session create messages              0
 Session close messages               0
 Session volume updates               0
 Failed messages                      0
```

Copyright © 2016, Juniper Networks, Inc.

show security datapath-debug capture

Syntax show security datapath-debug capture

Release Information Command introduced in Junos OS Release 10.0.

Description Display details of the data path debugging capture file.

Required Privilege Level view

Related Documentation
- show security datapath-debug counter on page 3659
- Understanding Data Path Debugging for Logical Systems on page 3476

List of Sample Output show security datapath—debug capture on page 3658

Output Fields Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath—debug capture

```
user@host> show security datapath-debug capture
Packet 1, len 120: (C0/F0/P0/SEQ:71:lbt)
91 00 00 47 11 00 10 00 9a 14 00 19 03 00 00 00
00 00 00 00 00 01 00 47 10 00 00 00 00 00 00 00
00 1f 12 f8 dd 29 00 21 59 84 f4 01 81 00 02 1e
08 00 45 60 01 f4 00 00 00 00 3f 06 73 9f 01 01
01 02 03 01 01 02 d4 31 d4 31 00 00 00 00 00 00
00 00 50 02 00 00 ff ad 00 00 00 00
Packet 2, len 120: (C0/F0/P0/SEQ:71:lbt)
90 00 00 47 04 00 00 00 00 00 00 00 02 02 00 47
10 00 00 00 00 00 00 00 50 00 a6 1c 00 00 00 00
00 00 00 0a 00 00 00 00 00 00 09 d9 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 1f 12 f8
dd 29 00 21 59 84 f4 01 81 00 02 1e
```


Copyright © 2016, Juniper Networks, Inc.

show security datapath-debug counter

Syntax show security datapath-debug counter

Release Information Command introduced in Junos OS Release 10.0.

Description Display details of the data path debugging counter.

Required Privilege Level view

Related Documentation
- show security datapath-debug capture on page 3658
- Understanding Data Path Debugging for Logical Systems on page 3476

List of Sample Output show security datapath-debug counter on page 3659

Output Fields Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath-debug counter

```
user@host> show security datapath-debug counter
Datapath debug counters
Packet Filter 1:
np-ingress
Chassis 0 FPC 4 : 1
np-ingress
Chassis 0 FPC 3 : 0
np-egress
Chassis 0 FPC 4 : 1
np-egress
Chassis 0 FPC 3 : 0
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 2
pot
Chassis 0 FPC 0 PIC 1: 0
pot
```

Copyright © 2016, Juniper Networks, Inc.

show security dns-cache

Syntax show security dns-cache <dns-name *dns-name*>

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display DNS cache information.

NOTE: This command is only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options
- **dns-name**—Display DNS cache information for the specified name.

Required Privilege Level view

Related Documentation
- clear security dns-cache on page 3640

List of Sample Output
show security dns-cache on page 3660
show security dns-cache dns-name dns2.test.com on page 3661

Output Fields Table 357 lists the output fields for the **show security dns-cache** command. Output fields are listed in the approximate order in which they appear.

Table 357: show security dns-cache Output Fields

Field Name	Field Description
DNS Name	DNS name.
Address Family	IPv4 or IPv6.
TTL	Time-to-live value.
IP Address	IP address for the DNS name.

Sample Output

show security dns-cache

```
user@host> show security dns-cache
DNS Name: dns1.test.com:
  Address Family: IPv4, TTL: 10
    IP Address: 1.1.1.1
  Address Family: IPv6: TTL = 15
    IP Address: 2001:1.1.1.1
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
```


Copyright © 2016, Juniper Networks, Inc.

```
    IP Address: 2.2.2.2
    IP Address: 2.2.2.3
```

Sample Output

show security dns-cache dns-name dns2.test.com

```
user@host> show security dns-cache dns-name dns2.test.com
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
    IP Address: 2.2.2.2
    IP Address: 2.2.2.3
```

show security firewall-authentication history

Syntax **show security firewall-authentication history**
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display security firewall authentication history information.

Options
- none—Display history of firewall authentication information.
- **node**—(Optional) For chassis cluster configurations, display all firewall authentication history on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- Understanding Logical System Firewall Authentication on page 3267
- Firewall User Authentication Overview on page 5179

List of Sample Output
show security firewall-authentication history on page 3663
show security firewall-authentication history node all on page 3663

Output Fields Table 358 lists the output fields for the **show security firewall-authentication history** command. Output fields are listed in the approximate order in which they appear.

Table 358: show security firewall-authentication history Output Fields

Field Name	Field Description
Authentications	Number of authentications.
Id	Identification number.
Source IP	IP address of the authentication source.
Date	Authentication date.
Time	Authentication time.
Duration	Authentication duration.
Status	Authentication status success or failure.

Copyright © 2016, Juniper Networks, Inc.

Table 358: show security firewall-authentication history Output Fields *(continued)*

Field Name	Field Description
User	Name of the user.

Sample Output

show security firewall-authentication history

```
user@host> show security firewall-authentication history
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           1 211.0.0.6       2007-04-03 11:43:06 00:00:45   Success  hello
```

Sample Output

show security firewall-authentication history node all

```
user@host> show security firewall-authentication history node all
node0:
--------------------------------------------------------------------------
History of firewall authentication data:
Authentications: 2
Id Source Ip       Date       Time     Duration   Status   User
1 100.0.0.1       2008-01-04 12:00:10 0:05:49    Success  local1
2 100.0.0.1       2008-01-04 14:36:52 0:01:03    Success  local1
node1:
--------------------------------------------------------------------------
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           1 100.0.0.1       2008-01-04 14:59:43 1193046:06: Success local1
```

Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users

Syntax

show security firewall-authentication users
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description

Display firewall authentication details about all users.

Options

- none—Display details about all firewall authentication users.
- **node**—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level

view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output

show security firewall-authentication users on page 3665
show security firewall-authentication users node 0 on page 3665
show security firewall-authentication users node all on page 3665

Output Fields

Table 359 lists the output fields for the **show security firewall-authentication users** command. Output fields are listed in the approximate order in which they appear.

Table 359: show security firewall-authentication users Output Fields

Field Name	Field Description
Total users in table	Gives count of how many entries/users the command will display.
Id	Identification number.
Source IP	IP address of the authentication source.
Src zone	User traffic received from the zone.
Dst zone	User traffic destined to the zone.
Profile	Name of profile used for authentication.
Age	Idle timeout for the user.

Copyright © 2016, Juniper Networks, Inc.

Table 359: show security firewall-authentication users Output Fields *(continued)*

Field Name	Field Description
Status	Authentication status success or failure.
User	Name of the user.

Sample Output

show security firewall-authentication users

```
user@host> show security firewall-authentication users
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           1 1111:1212/64        z1       z2       p1       0 Success  local1
```

Sample Output

show security firewall-authentication users node 0

```
user@host> show security firewall-authentication users node 0
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User

           3 100.0.0.1       z1       z2       p1           1 Success  local1
```

Sample Output

show security firewall-authentication users node all

```
user@host> show security firewall-authentication users node all
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 100.0.0.1       z1       z2       p1           1 Success  local1

node1:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           2 100.0.0.1       z1       z2       p1           1 Success  local1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session

Syntax **show security flow session**
[*filter*] **[brief | extensive | summary]**

Release Information Command introduced in Junos OS Release 8.5. Support for filter and view options added in Junos OS Release 10.2. Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2. Policy ID filter added in Junos OS Release 12.3X48-D10. Support for connection tag added in Junos OS Release 15.1X49-D40.

Description Display information about all currently active security sessions on the device.

NOTE: For the normal flow sessions, the **show security flow session** command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. Because of this, the output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific **show** command for examples of the filtered output.

  application—Predefined application name

  application-firewall—Application firewall enabled

  application-firewall-rule-set—Application firewall enabled with the specified rule set

  application-traffic-control—Application traffic control session

  application-traffic-control-rule-set—Application traffic control rule set name and rule name

  conn-tag—Session connection tag

  destination-port—Destination port

  destination-prefix—Destination IP prefix or address

  dynamic-application—Dynamic application

  dynamic-application-group—Dynamic application

  encrypted—Encrypted traffic

  extensive—Display detailed output

Copyright © 2016, Juniper Networks, Inc.

family—Display session by family

idp—IDP enabled sessions

interface—Name of incoming or outgoing interface

logical-system (all | *logical-system-name*)—Name of a specific logical system or **all** to display all logical systems

nat—Display sessions with network address translation

policy-id—Display session information based on policy ID; the range is 1 through 4,294,967,295

protocol—IP protocol number

resource-manager—Resource manager

root-logical-system—Display root logical system as default

security-intelligence—Display security intelligence sessions

services-offload—Display services offload sessions

session-identifier—Display session with specified session identifier

source-port—Source port

source-prefix—Source IP prefix

summary—Display output summary

tunnel—Tunnel sessions

- brief | extensive | summary—Display the specified level of output.
- none—Display information about all active sessions.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- clear security flow session all on page 1902

List of Sample Output
show security flow session on page 3669
show security flow session brief on page 3670
show security flow session extensive on page 3670
show security flow session summary on page 3670

Output Fields Table 26 lists the output fields for the **show security flow session** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 360: show security flow session Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Application traffic control rule-set	AppQoS rule set for this session.
Rule	AppQoS rule for this session.
Forwarding class	The AppQoS forwarding class name for this session that distinguishes the transmission priority
DSCP code point	Differentiated Services (DiffServ) code point (DSCP) value remarked by the matching rule for this session.

Copyright © 2016, Juniper Networks, Inc.

Table 360: show security flow session Output Fields *(continued)*

Field Name	Field Description
Loss priority	One of four priority levels set by the matching rule to control discarding a packet during periods of congestion. A high loss priority means a high probability that the packet could be dropped during a period of congestion.
Rate limiter client to server	The rate-limiter profile assigned to the client-to-server traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Rate limiter server to client	The rate-limiter profile assigned to the server-to-client traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Maximum number of sessions permitted.

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

Copyright © 2016, Juniper Networks, Inc.

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
 Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/10000/2000 --> 203.0.113.11/10000/1000;udp, Conn Tag: 0x0, If:
 reth0.0, Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
   In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
   Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```


Copyright © 2016, Juniper Networks, Inc.

```
Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

Copyright © 2016, Juniper Networks, Inc.

show security idp logical-system policy-association

Syntax show security idp logical-system policy-association

Release Information Command introduced in Junos OS Release 11.3.

Description Display the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.

Required Privilege Level view

Related Documentation
- security-profile on page 3628

List of Sample Output show security idp logical-system policy-association on page 3672

Output Fields Table 361 lists the output fields for the **show security idp logical-system policy-association** command.

Table 361: show security idp logical-system policy-association Output Fields

Field Name	Field Description
Logical system	Name of the logical system to which an IDP policy is assigned.
IDP policy	Name of the IDP policy that is specified in the security profile that is bound to the logical system.

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system          IDP policy
root-logical-system     idp-policy1
lsys1                   idp-policy2
```


Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations

Syntax

show security ike security-associations
peer-address
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number*
sa-type shortcut <detail>

Release Information

Command introduced in Junos OS Release 8.5 . Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for Auto Discovery VPN added in Junos OS Release 12.3X48-D10.

Description

Display information about Internet Key Exchange security associations (IKE SAs).

Options

- none—Display standard information about existing IKE SAs, including index numbers.
- ***peer-address***—(Optional) Display details about a particular SA based on the IPv4 or IPv6 address of the destination peer. This option and **index** provide the same level of output.
- **brief**—(Optional) Display standard information about all existing IKE SAs. (Default)
- **detail**—(Optional) Display detailed information about all existing IKE SAs.
- **family**—(Optional) Display IKE SAs by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc *slot-number***—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index *SA-index-number***—(Optional) Display information for a particular SA based on the index number of the SA. For a particular SA, display the list of existing SAs by using the command with no options. This option and ***peer-address*** provide the same level of output.
- **kmd-instance** —(Optional) Display information about existing IKE SAs in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic *slot-number*** —(Optional) Display information about existing IKE SAs in this PIC slot. This option is used to filter the output.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

List of Sample Output
show security ike security-associations (IPv4) on page 3676
show security ike security-associations (IPv6) on page 3676
show security ike security-associations detail (Branch SRX Series Devices) on page 3677
show security ike security-associations detail (High-End SRX Series Devices) on page 3677
show security ike security-associations family inet6 on page 3677
show security ike security-associations index 8 detail on page 3678
show security ike security-associations 1.1.1.2 on page 3678
show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 3679
show security ike security-associations detail (ADVPN Suggester, Static Tunnel) on page 3679
show security ike security-associations detail (ADVPN Partner, Static Tunnel) on page 3679
show security ike security-associations detail (ADVPN Partner, Shortcut) on page 3679
show security ike security-associations sa-type shortcut (ADVPN) on page 3680
show security ike security-associations sa-type shortcut detail (ADVPN) on page 3680

Output Fields Table 362 lists the output fields for the **show security ike security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 362: show security ike security-associations Output Fields

Field Name	Field Description
IKE Peer or Remote Address	IP address of the destination peer with which the local peer communicates.
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Gateway Name	Name of the IKE gateway.
Location	• **FPC**—Flexible PIC Concentrator (FPC) slot number. • **PIC**—PIC slot number. • **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances are running on each SPU, and any particular IKE negotiation is carried out by a single KMD instance.
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.
State	State of the IKE SAs: • **DOWN**—SA has not been negotiated with the peer. • **UP**—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.

Copyright © 2016, Juniper Networks, Inc.

Table 362: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode or Exchange type	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between one another. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are: • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected. NOTE: IKEv2 protocol does not use the mode configuration for negotiation. Therefore, mode displays the version number of the security association.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Algorithms	IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	
	• **Authentication**—Type of authentication algorithm used:
	  ◦ **sha1**—Secure Hash Algorithm 1 authentication.
	  ◦ **md5**—MD5 authentication.
	• **Encryption**—Type of encryption algorithm used:
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption.
	  ◦ **aes-128-cbc**—AES 128-bit encryption.
	  ◦ **3des-cbc**—3 Data Encryption Standard (DES) encryption.
	  ◦ **des-cbc**—DES encryption.

Diffie-Hellman group	Specifies the IKE Diffie-Hellman group.
Traffic statistics	• **Input bytes**—Number of bytes received. • **Output bytes**—Number of bytes transmitted. • **Input packets**—Number of packets received. • **Output packets**—Number of packets transmitted.

Copyright © 2016, Juniper Networks, Inc.

Table 362: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
Flags	Notification to the key management process of the status of the IKE negotiation: • **caller notification sent**—Caller program notified about the completion of the IKE negotiation. • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.
IPSec security associations	• ***number* created:** The number of SAs created. • ***number* deleted**: The number of SAs deleted.

Phase 2 negotiations in progress	Number of Phase 2 IKE negotiations in progress and status information:
	• **Negotiation type**—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• **Message ID**—Unique identifier for a Phase 2 negotiation.
	• **Local identity**—Identity of the local Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• **Remote identity**—Identity of the remote Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• **Flags**—Notification to the key management process of the status of the IKE negotiation:
	  ◦ **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  ◦ **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  ◦ **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  ◦ **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.

Sample Output

show security ike security-associations (IPv4)

```
user@host> show security ike security-associations
Index  Remote Address    State  Initiator cookie     Responder cookie     Mode
8  1.1.1.2    UP  3a895f8a9f620198     9040753e66d700bb     Main
Index  Remote Address    State  Initiator cookie     Responder cookie     Mode
9  1.2.1.3     UP  5ba96hfa9f65067   1  70890755b65b80b   d  Main
```

show security ike security-associations (IPv6)

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5       UP     e48efd6a444853cf  0d09c59aafb720be  Aggressive     1212::1112
```


Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations detail (Branch SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 25.191.134.245, Index 2577565, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: b869b3424513340a, Responder cookie: 4cb3488cb19397c3
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 25.191.134.241:500, Remote: 25.191.134.245:500
  Lifetime: Expires in 169 seconds
  Peer ike-id: 25.191.134.245
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1012
   Output bytes  :                 1196
   Input  packets:                    4
   Output packets:                    5
  Flags: IKE SA is created
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 25.191.134.241:500, Remote: 25.191.134.245:500
    Local identity: 25.191.134.241
    Remote identity: 25.191.134.245
    Flags: IKE SA is created
```

show security ike security-associations detail (High-End SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 1.1.1.2, Index 914039858, Gateway Name: tropic
  Location: FPC 3, PIC 1, KMD-Instance 3
  Role: Initiator, State: UP
  Initiator cookie: 219a697652bdde37, Responder cookie: b49c30b229d36bcd
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 1.1.1.1:500, Remote: 1.1.1.2:500
  Lifetime: Expires in 26297 seconds
  Peer ike-id: 1.1.1.2
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 1
```

show security ike security-associations family inet6

```
user@host> show security ike security-associations family inet6
```

```
  IKE peer 1212::1112, Index 5, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: e48efd6a444853cf, Responder cookie: 0d09c59aafb720be
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 1212::1111:500, Remote: 1212::1112:500
  Lifetime: Expires in 19518 seconds
  Peer ike-id: not valid
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : sha1
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1568
   Output bytes  :                 2748
   Input  packets:                    6
   Output packets:                   23
  Flags: Caller notification sent
  IPSec security associations: 5 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 2900338624
    Local: 1212::1111:500, Remote: 1212::1112:500
    Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Flags: Caller notification sent, Waiting for done
```

show security ike security-associations index 8 detail

```
user@host> show security ike security-associations index 8 detail
IKE peer 1.1.1.2, Index 8, Gateway Name: tropic
 Role: Responder, State:UP
 Initiator cookie: 3a895f8a9f620198, Responder cookie: 9040753e66d700bb
 Exchange type; main, Authentication method: Pre-shared-keys
 Local: 1.1.1.1:500, Remote: 1.1.1.2:500
 Lifetime: Expired in 381 seconds
 Algorithms:
  Authentication:       md5
  Encryption:      3des-cbc
  Pseudo random function      hmac-md5
     Diffie-Hellman group  : DH-group-5
    Traffic statistics:
  Input bytes:       11268
  Output bytes:        6940
  Input packets:        57
  Output packets:        57
 Flags: Caller notification sent
 IPsec security associations: 0 created, 0 deleted
 Phase 2 negotiations in progress: 1

  Negotiation type: Quick mode, Role: Responder, Message ID: 1765792815
  Local: 1.1.1.1:500, Remote: 1.1.1.2:500
  Local identity: No Id
  Remote identity: No Id
  Flags: Caller notification sent, Waiting for remove
```

show security ike security-associations 1.1.1.2

```
user@host> show security ike security-associations 1.1.1.2
```


Copyright © 2016, Juniper Networks, Inc.

```
Index     State  Initiator cookie  Responder cookie  Mode Remote Address
   8        UP     3a895f8a9f620198  9040753e66d700bb  Main 1.1.1.2
```

show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ike security-associations fpc 6 pic 1 kmd-instance all
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode

1728053250 1.1.1.2      UP     fc959afd1070d10b  bdeb7e8c1ea99483  Main
```

show security ike security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.105, Index 13563297, Gateway Name: zth_hub_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent            :  12
     Suggestion response accepted:  12
     Suggestion response declined:   0
  Role: Responder, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.154:500, Remote: 23.0.0.105:500
  Lifetime: Expires in 26429 seconds
  Peer ike-id: DC=example, CN=test, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.154, Index 4980720, Gateway Name: zth_spoke_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:  12
     Suggestions accepted:  12
     Suggestions declined:   0
  Role: Initiator, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.105:500, Remote: 23.0.0.154:500
  Lifetime: Expires in 26252 seconds
  Peer ike-id: DC=example, CN=vsrxvpn01, OU=SBU, O=Example, L=Sunnyvale, ST=CA,
C=US
```

show security ike security-associations detail (ADVPN Partner, Shortcut)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.106, Index 4980737, Gateway Name:
GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173323
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: e1ed0c655929debc, Responder cookie: 437de6ed784ba63e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.105:500, Remote: 23.0.0.106:500
  Lifetime: Expires in 28796 seconds
  Peer ike-id: DC=example, CN=paulyd, L=Sunnyvale, ST=CA, C=US
```

Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations sa-type shortcut (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut
Initiator cookie  Responder cookie  Mode           Remote Address
4980742Index   State  UP     b56fbe694eaee5b6  064dbccbfa3b2aab  IKEv2
23.0.0.106
```

show security ike security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut detail
IKE peer 23.0.0.106, Index 4980742, Gateway Name:
GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173327
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Role: Partner, Peer Role: Partner
  Role: Responder, State: UP
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec security-associations

Syntax

show security ipsec security-associations
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number>*
sa-type shortcut
vpn-name *vpn-name* <traffic-selector *traffic-selector-name*>

Release Information

Command introduced in Junos OS Release 8.5. Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for the **vpn-name** option added in Junos OS Release 11.4R3. Support for the **traffic-selector** option and traffic selector field added in Junos OS Release 12.1X46-D10. Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10.

Description

Display information about the IPsec security associations (SAs).

Options

- none—Display information about all SAs.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **family**—(Optional) Display SAs by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc** ***slot-number***—(Optional) Display information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index** ***SA-index-number***—(Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- **kmd-instance**—(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number***—(Optional) Display information about existing IPsec SAs in this PIC slot. This option is used to filter the output.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.
- **vpn-name** ***vpn-name***—Name of the VPN. If configured, **traffic-selector** ***traffic-selector-name*** can optionally be specified.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- clear security ipsec security-associations on page 6986
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

List of Sample Output
show security ipsec security-associations (IPv4) on page 3685
show security ipsec security-associations (IPv6) on page 3685
show security ipsec security-associations index 131073 on page 3685
show security ipsec security-associations brief on page 3686
show security ipsec security-associations detail on page 3686
show security ipsec security-associations family inet6 on page 3687
show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 3687
show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel) on page 3687
show security ike sa index 222075191 detail on page 3688
show security ipsec security-associations detail (ADVPN Partner, Static Tunnel) on page 3689
show security ike sa index 788674 detail on page 3690
show security ipsec security-associations sa-type shortcut (ADVPN) on page 3690
show security ipsec security-associations sa-type shortcut detail (ADVPN) on page 3690
show security ipsec security-associations family inet detail on page 3691

Output Fields Table 363 lists the output fields for the **show security ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 363: show security ipsec security-associations

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.
ID	Index number of the SA. You can use this number to get additional information about the SA.
VPN name	IPsec name for VPN.
Gateway	IP address of the remote gateway.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are **hmac-md5-95**, **hmac-sha1-96,** or **ESP**. • An encryption algorithm used to encrypt data traffic. Options are **3des-cbc, aes-128-cbc, aes-192-cbc**, **aes-256-cbc**, or **des-cbc.**

Copyright © 2016, Juniper Networks, Inc.

Table 363: show security ipsec security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
Sta	State has two options, **Installed** and **Not Installed**. • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.
Mon	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays **U** (up) or **D** (down). A hyphen (-) means VPN monitoring is not enabled for this SA.
vsys or Virtual-system	The root system.
Tunnel index	Numeric identifier of the specific IPsec tunnel for the SA.
Local gateway	Gateway address of the local system.
Remote gateway	Gateway address of the remote system.
Traffic selector	Name of the traffic selector.
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote identity	IP address of the destination peer gateway.
DF-bit	State of the don't fragment bit: **set** or **cleared**.
Policy-name	Name of the applicable policy.
Location	**FPC**—Flexible PIC Concentrator (FPC) slot number. **PIC**—PIC slot number. **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance.
Tunnel events	Tunnel event and the number of times the event has occurred. See "Tunnel Events" on page 6797 for descriptions of tunnel events and the action you can take.
Direction	Direction of the SA; it can be inbound or outbound.

Copyright © 2016, Juniper Networks, Inc.

Table 363: show security ipsec security-associations *(continued)*

Field Name	Field Description
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is **AH** or **ESP**, **AUX-SPI** is always 0. • When the value is **AH+ESP**, **AUX-SPI** is always a positive integer.
Mode	Mode of the SA: • **transport**—Protects host-to-host connections. • **tunnel—**Protects connections between security gateways.
Type	Type of the SA: • **manual**—Security parameters require no negotiation. They are static and are configured by the user. • **dynamic**—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.
State	State of the SA: • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.

Protocol	Protocol supported.
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  ◦ **Authentication**—Type of authentication used.
	  ◦ **Encryption**—Type of encryption used.

Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in seconds**—Number of seconds left until the SA expires.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • **Expires in seconds**—Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • **Expires in kilobytes**—Number of kilobytes left until the SA expires.
Anti-replay service	State of the service that prevents packets from being replayed. It can be **Enabled** or **Disabled**.

Copyright © 2016, Juniper Networks, Inc.

Table 363: show security ipsec security-associations *(continued)*

Field Name	Field Description
Replay window size	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Copy-Outer-DSCP	Indicates if copying outer IP header DSCP and ECN to inner IP header is enabled or disabled.

Sample Output

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys
  131075 11.0.28.241    500   ESP:3des/sha1   86758ff0 6918/ unlim   -   0
  131075 11.0.28.241    500   ESP:3des/sha1   3183ff26 6918/ unlim   -   0
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  131074 ESP:3des/sha1 14caf1d9 3597/ unlim   -   root 500   1212::1112
  131074 ESP:3des/sha1 9a4db486 3597/ unlim   -   root 500   1212::1112
```

show security ipsec security-associations index 131073

```
user@host> show security ipsec security-associations index 131073
  ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
```

Copyright © 2016, Juniper Networks, Inc.

```
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations brief

```
user@host> show security ipsec security-associations brief
Total active tunnels: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<16384 1.1.1.1 500 ESP:3des/sha1 af88baa 28795/unlim D 0
>16384 1.1.1.1 500 ESP:3des/sha1 f4e3e5f4 28795/unlim D 0
```

show security ipsec security-associations detail

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Version: IKEv1
DF-bit: clear
, Copy-Outer-DSCP Enabled
Bind-interface: st0.99

Port: 500, Nego#: 8, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
Tunnel events:
Mon Oct 26 2015 22:27:50 -0700: IPSec SA rekey successfully completed (7 times)
Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:41:07 -0700: IKE SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information
 updated (1 times)
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: inbound, SPI: 81ed9998, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled
```


Copyright © 2016, Juniper Networks, Inc.

```
, Replay window size: 64
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: outbound, SPI: 80565248, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```

show security ipsec security-associations family inet6

```
user@host> show security ipsec security-associations family inet6
  Virtual-system: root
  Local Gateway: 1212::1111, Remote Gateway: 1212::1112
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
  Total active tunnels: 1

ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys

<2    1.1.1.2          500   ESP:3des/sha1   67a7d25d 28280/unlim   -   0

>2    1.1.1.2          500   ESP:3des/sha1   a23cbcdc 28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear
  Bind-interface: st0.1

  Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
```

Copyright © 2016, Juniper Networks, Inc.

```
  Tunnel events:
  Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
  Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding
IKE/IPSec SAs are deleted (1 times)
  Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ike sa index 222075191 detail

```
user@host> show security ike sa index 222075191 detail
node0:
--------------------------------------------------------------------------
IKE peer 31.1.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                20474
   Output bytes  :                21091
   Input  packets:                  237
   Output packets:                  237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 31.1.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local
certificate. Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate
 not found. Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

Copyright © 2016, Juniper Networks, Inc.

show security ike sa index 788674 detail

```
user@host> show security ike sa index 788674 detail
IKE peer 11.1.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                22535
   Output bytes  :                21918
   Input  packets:                  256
   Output packets:                  256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 31.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes-cbc-256/sha1 6f164ee0 3580/ unlim - root 500 23.0.0.111
>268173318 ESP:aes-cbc-256/sha1 e6f29cb0 3580/ unlim - root 500 23.0.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
```


Copyright © 2016, Juniper Networks, Inc.

```
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
```

```
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1090 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```


Copyright © 2016, Juniper Networks, Inc.

show security match-policies

Syntax

show security match-policies
destination-ip *ip-address*
destination-port *port-number*
from-zone *zone-name*
protocol *protocol-name* | *protocol-number*
result-count *number*
source-identity *role-name*
source-ip *ip-address*
source-port *port-number*
to-zone *zone-name*

Release Information

Command introduced in Junos OS Release 10.3. Command updated in Junos OS Release 10.4 . Command updated in Junos OS Release 12.1. Command updated to include optional from-zone and to-zone global match options in Junos OS Release 12.1X47-D10.

Description

The **show security match-policies** command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, then the **show security match-policies** command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The **result-count** option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are "shadowed" by the first and are never encountered by matching traffic.

NOTE: **The show security match-policies command is applicable only to security policies; IDP policies are not supported.**

Options

- **destination-ip** ***destination-ip***—Destination IP address of the traffic.
- **destination-port** ***destination-port***—Destination port number of the traffic. Range is 1 through 65,535
- **from-zone** ***from-zone***—Name or ID of the source zone of the traffic.
- **protocol** ***protocol-name*** **|** ***protocol-number***—Protocol name or numeric value of the traffic.
  - **ah** or **51**
  - **egp** or **8**
  - **esp** or **50**
  - **gre** or **47**
  - **icmp** or **1**
  - **igmp** or **2**

Copyright © 2016, Juniper Networks, Inc.

- **igp** or **9**
- **ipip** or **94**
- **ipv6** or **41**
- **ospf** or **89**
- **pgm** or **113**
- **pim** or **103**
- **rdp** or **27**
- **rsvp** or **46**
- **sctp** or **132**
- **tcp** or **6**
- **udp** or **17**
- **vrrp** or **112**

- **result-count** ***number***—(Optional) The number of policy matches to display. Valid range is from 1 through 16. The default value is 1.
- **source-identity** ***role-name***—Source identity of the traffic determined by the user role.
- **source-ip** ***source-ip***—Source IP address of the traffic.
- **source-port** ***source-port***—Source port number of the traffic. Range is 1 through 65,535.
- **to-zone** ***to-zone***—Name or ID of the destination zone of the traffic.

Required Privilege Level view

Related Documentation
- clear security policies statistics on page 1344
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output
Example 1: show security match-policies on page 3696
Example 2: show security match policies ... result-count on page 3696
Example 3: show security match policies ... source-identity on page 3697

Output Fields Table 102 lists the output fields for the **show security match-policies** command. Output fields are listed in the approximate order in which they appear.

Table 364: show security match-policies Output Fields

Field Name	Field Description
Policy	Name of the applicable policy.

Copyright © 2016, Juniper Networks, Inc.

Table 364: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Action or **Action-type**	The action to be taken for traffic that matches the policy's match criteria. Actions include the following: • **permit** • **firewall-authentication** • **tunnel ipsec-vpn** *vpn-name* • **pair-policy** *pair-policy-name* • **source-nat pool** *pool-name* • **pool-set** *pool-set-name* • **interface** • **destination-nat** *name* • **deny** • **reject**
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	An internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
From zone	Name of the source zone.
To zone	Name of the destination zone.
Source addresses	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
Application	Name of a preconfigured or custom application, or **any** if no application is specified.
IP protocol	Numeric value for the IP protocol used by the application, such as 6 for TCP or 1 for ICMP.
ALG	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
Inactivity timeout	Elapsed time without activity after which the application is terminated.
Source-port range	Range of matching source ports defined in the policy.

Copyright © 2016, Juniper Networks, Inc.

Table 364: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Destination-port range	Range of matching destination ports defined in the policy.
Source identities	One or more user roles defined in the matching policy.

Sample Output

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1 to-zone z2 source-ip 10.10.10.1
destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 198.51.100.0/24
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.0/24
    d3: 192.0.2.1/32
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [21-21]
```

Example 2: show security match policies ... result-count

```
user@host> show security match-policies source-ip 10.10.10.1 destination-ip 192.0.2.5 source_port
1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/32
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]
```


Copyright © 2016, Juniper Networks, Inc.

Example 3: show security match policies ... source-identity

```
user@host> show security match-policies from-zone untrust to-zone trust source-ip 10.10.10.1
destination-ip 192.0.2.1 destination_port 21 protocol 6 source-port 1234 source-identity role1
Policy: p1, action-type: permit, State: enabled, Index: 40
      Policy Type: Configured
      Sequence number: 1
      From zone: untrust, To zone: trust
      Source addresses:
        a1: 10.0.0.0/8
      Destination addresses:
        d1: 192.0.2.0/24
      Application: junos-ftp
        IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
          Source port range: [0-0]
          Destination port range: [21-21]
      Source identities: role1
      Per policy TCP Options: SYN check: No, SEQ check: No
```

Copyright © 2016, Juniper Networks, Inc.

show security nat destination rule

Syntax

show security nat destination rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the **application** option field added in Junos OS Release 12.1X47-D10.

Description

Display information about the specified destination Network Address Translation (NAT) rule.

Options

rule-name—Display information about the specified destination NAT rule.

all—Display information about all the destination NAT rules.

logical-system (*logical-system-name* | all)—Display information about the destination NAT rules for the specified logical system or for all logical systems.

root-logical-system—Display information about the destination NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Destination NAT) on page 5092
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat destination rule dst2-rule on page 3699
show security nat destination rule all on page 3700

Output Fields

Table 365 lists the output fields for the **show security nat destination rule** command. Output fields are listed in the approximate order in which they appear.

Table 365: show security nat destination rule Output Fields

Field Name	Field Description
Total destination-nat rules	Number of destination NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly as address names and as address set names in the rule.
Destination NAT rule	Name of the destination NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Table 365: show security nat destination rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the destination NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the destination NAT rule.
From routing instance	Name of the routing instance from which the packets flow.
From interface	Name of the interface from which the packets flow.
From zone	Name of the zone from which the packets flow.
Source addresses	Name of the source addresses that match the rule. The default value is any.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: • **destination NAT pool**—Use user-defined destination NAT pool to perform destination NAT. • **off**—Do not perform destination NAT.
Destination ports	Destination ports number that match the rule. The default value is any.
Application	Indicates whether the application option is configured.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat destination rule dst2-rule

```
user@host>show security nat destination rule dst2-rule

 Destination NAT rule: dst2-rule           Rule-set: dst2
  Description               : The destination rule dst2-rule is for the sales
team
  Rule-Id                   : 1
  Rule position             : 1
  From routing instance     : ri1
                            : ri2
  Match
    Source addresses        : add1
```

Copyright © 2016, Juniper Networks, Inc.

```
                                add2
    Destination addresses    : add9
  Action                     : off

  Destination port           : 0
 Translation hits            : 68
   Successful sessions       : 25
   Failed sessions           : 43
 Number of sessions          : 2
```

Sample Output

show security nat destination rule all

```
user@host> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0

Destination NAT rule: r4                    Rule-set: rs4
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : untrust
  Match
    Source addresses        : 192.0.2.0 - 192.0.2.255
    Destination addresses   : 198.51.100.0 - 198.51.100.255
    Application             : configured
  Action                    : off
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat destination summary

Syntax

show security nat destination summary
<logical-system (*logical-system-name* | all)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

Description

Display a summary of Network Address Translation (NAT) destination pool information.

Options

none—Display summary information about the destination NAT pool.

logical-system (*logical-system-name* | all)—Display summary information about the destination NAT for the specified logical system or for all logical systems.

root-logical-system—Display summary information about the destination NAT for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Destination NAT) on page 5075
- rule (Security Destination NAT) on page 5092
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat destination summary on page 3702

Output Fields

Table 366 lists the output fields for the **show security nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 366: show security nat destination summary Output Fields

Field Name	Field Description
Total destination nat pool number	Number of destination NAT pools.
Pool name	Name of the destination address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.
Total	Number of IP addresses that are in use.
Available	Number of IP addresses that are free for use.
Total destination nat rule number	Number of destination NAT rules.

Copyright © 2016, Juniper Networks, Inc.

Table 366: show security nat destination summary Output Fields *(continued)*

Field Name	Field Description
Total hit times	Number of times a translation in the translation table is used for all the destination NAT rules.
Total fail times	Number of times a translation in the translation table failed to translate for all the destination NAT rules.

Sample Output

show security nat destination summary

```
user@host> show security nat destination summary

Total pools: 2
Pool name           Address                     Routing        Port  Total
                    Range                       Instance             Address
dst-p1              203.0.113.1 -203.0.113.1         default        0     1
dst-p2              2001:db8::1 - 2001:db8::1        default        0     1

Total rules: 171
Rule name           Rule set       From                           Action
dst2-rule           dst2           ri1
                                   ri2
                                   ri3
                                   ri4
                                   ri5
                                   ri6
                                   ri7
dst3-rule           dst3           ri9                            off
                                   ri1
                                   ri2
                                   ri3
                                   ri4
                                   ri5

...
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source rule

Syntax

show security nat source rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 11.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Source port**, **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the **application** output field added in Junos OS Release 12.1X47-D10.

Description

Display information about the specified source Network Address Translation (NAT) rule.

Options

rule-name—Name of the rule.

all—Display information about all the source NAT rules.

logical-system (*logical-system-name* | all)—Display information about the source NAT rules for the specified logical system or for all logical systems source NAT rules.

root-logical-system—Display information about the source NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Source NAT) on page 5093
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source rule r2 on page 3705
show security nat source rule all on page 3705

Output Fields

Table 367 lists the output fields for the **show security nat source rule** command. Output fields are listed in the approximate order in which they appear

Table 367: show security nat source rule Output Fields

Field Name	Field Description
Source NAT rule	Name of the source NAT rule.
Total rules	Number of source NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.

Copyright © 2016, Juniper Networks, Inc.

Table 367: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the source NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the source NAT rule.
From zone	Name of the zone from which the packets flow.
To zone	Name of the zone to which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
To routing instance	Name of the routing instance to which the packets flow.
From interface	Name of the interface from which the packets flow.
To interface	Name of the interface to which the packets flow.
Source addresses	Name of the source addresses that match the rule.
Source port	Source port numbers that match the rule.
Destination address	Name of the destination addresses that match the rule.
Destination ports	Destination port numbers that match the rule.
Application	Indicates whether the application option is configured.
Action	The action taken in regard to a packet that matches the rule's tuples. Actions include the following: • **off**—Do not perform source NAT. • **source NAT pool**—Use user-defined source NAT pool to perform source NAT • **interface**—Use egress interface's IP address to perform source NAT.
Persistent NAT type	Persistent NAT type.
Persistent NAT mapping type	Persistent NAT mapping type.
Inactivity timeout	Inactivity timeout for persistent NAT binding.
Max session number	Maximum number of sessions.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.

Copyright © 2016, Juniper Networks, Inc.

Table 367: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat source rule r2

```
user@host> show security nat source rule r2

 source NAT rule: r2        Rule-set: src-nat
  Description               : The source rule r2 is for the sales team
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : zone1
  To zone                   : zone9
  Match
   Source addresses         : add1
                              add2
   Destination addresses    : add9
                              add10
   Destination port         : 1002            - 1002
  Action                       : off
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits             : 4719
   Successful sessions      : 2000
   Failed sessions          : 2719
 Number of sessions         : 5
```

Sample Output

show security nat source rule all

```
user@host> show security nat source rule all
Logical system: root
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 3/0

source NAT rule: r2                  Rule-set: rs2
  Rule-Id                  : 2
  Rule position            : 1
  From zone                : trust
  To zone                  : untrust
  Match
    Source addresses       : 192.0.2.0 - 192.0.2.255
    Destination addresses  : 203.0.113.0 - 203.0.113.255
                             198.51.100.0 - 198.51.100.2555
    Application            : configured
  Action                      : off
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
  Max session number          : 0
Translation hits              : 0
  Successful sessions         : 0
  Failed sessions             : 0
Number of sessions            : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source summary

Syntax

show security nat source summary
<logical-system (*logical-system-name* | all)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

Description

Display a summary of Network Address Translation (NAT) source information.

Options

none—Display summary source NAT information.

logical-system (*logical-system-name* | all)—Display summary information about the source NAT for the specified logical system or for all logical systems.

root-logical-system—Display summary information about the source NAT for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Source NAT) on page 5076
- rule (Security Source NAT) on page 5093
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source summary on page 3708

Output Fields

Table 368 lists the output fields for the **show security nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 368: show security nat source summary Output Fields

Field Name	Field Description
Total source nat pool number	Number of source NAT pools.
Pool name	Name of the source address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).
Total Address	Number of IP addresses that are in use.
Total source nat rule number	Number of source NAT rules.

Copyright © 2016, Juniper Networks, Inc.

Table 368: show security nat source summary Output Fields *(continued)*

Field Name	Field Description
Total port number usage for port translation pool	Number of ports assigned to the pool.
Maximum port number for port translation pool	Maximum number of NAT or PAT transactions done at any given time.

Sample Output

show security nat source summary

```
user@host> show security nat source summary logical-system all


Logical system: root-logical-system
Total port number usage for port translation pool: 67108864
Maximum port number for port translation pool: 134217728

Logical system: lsys1
Total port number usage for port translation pool: 193536
Maximum port number for port translation pool: 134217728
Total pools: 2

Logical system: root-logical-system
Pool              Address                   Routing   PAT   Total
Name              Range                     Instance        Address
pool1             10.1.1.0-10.1.4.255-        default   yes   2048
                  10.1.5.0-10.1.8.255


Logical system: lsys1
Pool              Address                    Routing   PAT   Total
Name              Range                      Instance        Address
pool2             203.0.113.1-203.0.113.3     default   yes   3


Total rules: 1

Logical system: root-logical-system
Rule name         Rule set      From         To             Action
rule 1            ruleset1      ge-2/2/2.0   ge-2/2/3.0     pool1
rule 1                          ge-2/2/4.0   ge-2/2/5.0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat static rule

Syntax

show security nat static rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.3. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Successful sessions**, **Failed sessions**, **Number of sessions**, **Source addresses,** and **Source ports** output fields added in Junos OS Release 12.1X45-D10. The **Destination NPTv6 addr** and **Destination NPTv6 Netmask** output fields added in Junos OS Release 12.3X48-D25.

Description

Display information about the specified static Network Address Translation (NAT) rule.

Options

rule-name—Name of the rule.

all—Display information about all the static NAT rules.

logical-system (*logical-system-name* | all)—Display information about the static NAT rules for the specified logical system or for all logical systems.

root-logical-system—Display information about the static NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Static NAT) on page 5095
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat static rule on page 3710
show security nat static rule (IPv6) on page 3711
show security nat static rule all on page 3711

Output Fields

Table 369 lists the output fields for the **show security nat static rule** command. Output fields are listed in the approximate order in which they appear.

Table 369: show security nat static rule Output Fields

Field Name	Field Description
Static NAT rule	Name of the static NAT rule.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Rule-set	Name of the rule set. Currently, you can configure 8 rules within the same rule set.

Copyright © 2016, Juniper Networks, Inc.

Table 369: show security nat static rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the static NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the rule that indicates the order in which it applies to traffic.
From interface	Name of the interface from which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
From zone	Name of the zone from which the packets flow.
Destination addresses	Name of the destination addresses that match the rule.
Destination NPTv6 addr	Destination address that matches the rule.
Source addresses	Name of the source addresses that match the rule.
Host addresses	Name of the host addresses that match the rule.
Netmask	Subnet IP address.
Destination NPTv6 Netmask	Subnet IPv6 address.
Host routing-instance	Name of the host routing instance.
Destination port	Destination port numbers that match the rule. The default value is any.
Source port	Source port numbers that match the rule.
Total static-nat rules	Number of static NAT rules.
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat static rule

```
user@host> show security nat static rule sta-r2

 Static NAT rule: sta-r2                  Rule-set: sta-nat
```


Copyright © 2016, Juniper Networks, Inc.

```
  Description               : The static rule sta-r2 is for the sales team
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : zone9
  Destination addresses     : add3
  Host addresses            : add4
  Netmask                   : 24
  Host routing-instance     : N/A
  Translation hits          : 2
   Successful sessions     : 2
   Failed sessions         : 0
 Number of sessions        : 2
```

Sample Output

show security nat static rule (IPv6)

```
user@host> show security nat static rule r1

Static NAT rule: r1                  Rule-set: rs1
   Rule-Id                   : 1
   Rule position             : 1
   From zone                 : trust
   Destination NPTv6 addr    : 2001:db8::
   Destination NPTv6 Netmask : 48
   Host addresses            : 2001:db8::3000:
   Netmask                   : 48
   Host routing-instance     : N/A
   Translation hits          : 0
      Successful sessions     : 0
      Failed sessions         : 0
   Number of sessions        : 0
```

Sample Output

show security nat static rule all

```
user@host> show security nat static rule all

Static NAT rule: r1                  Rule-set: rs1
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : trust
  Source addresses          : 192.0.2.0 -192.0.2.3
: addr1
  Source ports              : 200 - 300
  Destination addresses     : 198.51.100.0
  Host addresses            : 203.0.113.0
  Netmask                   : 24
  Host routing-instance     : N/A
  Translation hits          : 4
   Successful sessions     : 4
   Failed sessions         : 0
 Number of sessions        : 4
Static NAT rule: r2                  Rule-set: rs1
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : trust
  Source addresses          : 192.0.2.0 -192.0.2.255  Destination addresses
  : 203.0.113.1
  Destination ports         : 100 - 200
```

Copyright © 2016, Juniper Networks, Inc.

```
  Host addresses             : 192.0.2.1
  Host ports                 : 300 - 400
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 4
   Successful sessions      : 4
   Failed sessions          : 0
Number of sessions          : 4
```


Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 3716
show security policies policy-name p1 detail on page 3717
show security policies (services-offload) on page 3718
show security policies detail on page 3718
show security policies detail (TCP Options) on page 3719
show security policies policy-name p1 (Negated Address) on page 3719
show security policies policy-name p1 detail (Negated Address) on page 3720
show security policies global on page 3720

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 370: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 370: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 370: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```


Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                 272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
      Initial direction:                9072                  272 bps
      Reply direction  :                9072                  272 bps
    Input  packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Output packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Session rate       :                 108                    3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

show security screen statistics

Syntax

show security screen statistics (zone *zone-name* | interface *interface-name*)
<logical-system (*logical-system-name* | all)>
<node (*node-id* | all | local | primary)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 8.5. **node** options added in Junos OS Release 9.0. **logical-system all** option added in Junos OS Release 11.2R6. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

Description

Display intrusion detection service (IDS) security screen statistics.

Options

- **zone** ***zone-name*** —Display screen statistics for this security zone.
- **interface** *interface-name*—Display screen statistics for this interface.
- **logical-system**—(Optional) Display screen statistics for configured logical systems.
  - ***logical-system-name***—Display screen statistics for the named logical system.
  - **all**—Display screen statistics for all logical systems, including the master (root) logical system.
- **node**—(Optional) For chassis cluster configurations, display screen statistics on a specific node.
  - ***node-id***—Identification number of a node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.
- **root-logical-system**—(Optional) Display screen statistics for the master logical system only.

Required Privilege Level

view

Related Documentation

- clear security screen statistics on page 1022
- clear security screen statistics interface on page 1023
- clear security screen statistics zone on page 1025
- Example: Configuring Multiple Screening Options on page 849

List of Sample Output

show security screen statistics zone scrzone on page 3724
show security screen statistics zone untrust (IPv6) on page 3724
show security screen statistics interface ge-0/0/3 on page 3725
show security screen statistics interface ge-0/0/1 (IPv6) on page 3725
show security screen statistics interface ge-0/0/1 node primary on page 3726
show security screen statistics zone trust logical-system all on page 3726

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 69 lists the output fields for the **show security screen statistics** command. Output fields are listed in the approximate order in which they appear.

Table 371: show security screen statistics Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.

Copyright © 2016, Juniper Networks, Inc.

Table 371: show security screen statistics Output Fields *(continued)*

Field Name	Field Description
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.
UDP address sweep	Number of UDP address sweeps.
IPv6 extension header	Number of packets filtered for the defined IPv6 extension headers.
IPv6 extension hop by hop option	Number of packets filtered for the defined IPv6 hop-by-hop option types.
IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).

Copyright © 2016, Juniper Networks, Inc.

Table 371: show security screen statistics Output Fields *(continued)*

ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                            Statistics
  ICMP flood                                    0
  UDP  flood                                    0
  TCP  winnuke                                  0
  ……
  IPv6   extension header                        0
  IPv6   extension hop by hop option             0
  IPv6   extension destination option            0
  IPv6   extension header limit                  0
  IPv6   malformed header                        0
```


Copyright © 2016, Juniper Networks, Inc.

```
  ICMPv6 malformed packet                    0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              91
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
  IP source route option                     0
  TCP land attack                            0
  TCP SYN fragment                           0
  TCP no flag                                0
  IP unknown protocol                        0
  IP bad options                             0
  IP record route option                     0
  IP timestamp option                        0
  IP security option                         0
  IP loose source route option               0
  IP strict source route option              0
  IP stream option                           0
  ICMP fragment                              0
  ICMP large packet                          0
  TCP SYN FIN                                0
  TCP FIN no ACK                             0
  Source session limit                       0
  TCP SYN-ACK-ACK proxy                      0
  IP block fragment                          0
  Destination session limit                  0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                      0
  UDP flood                                       0
  ……
  IPv6 extension header                           0
  IPv6 extension hop by hop option                0
  IPv6 extension destination option               0
  IPv6 extension header limit                     0
  IPv6 malformed header                           0
  ICMPv6 malformed packet                         0
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                              Statistics
  ICMP flood                                 1
  UDP flood                                  1
  TCP winnuke                                1
  TCP port scan                              1
  ICMP address sweep                         1
  TCP sweep                                  1
  UDP sweep                                  1
  IP tear drop                               1
  TCP SYN flood                              1
  IP spoofing                                1
  ICMP ping of death                         1
  IP source route option                     1
  TCP land attack                            1
  TCP SYN fragment                           1
  TCP no flag                                1
  IP unknown protocol                        1
  IP bad options                             1
  IP record route option                     1
  IP timestamp option                        1
  IP security option                         1
  IP loose source route option               1
  IP strict source route option              1
  IP stream option                           1
  ICMP fragment                              1
  ICMP large packet                          1
  TCP SYN FIN                                1
  TCP FIN no ACK                             1
  Source session limit                       1
  TCP SYN-ACK-ACK proxy                      1
  IP block fragment                          1
  Destination session limit                  1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                              Statistics
  ICMP flood                                 0
  UDP flood                                  0
  TCP winnuke                                0
  TCP port scan                              0
  ICMP address sweep                         0
  TCP sweep                                  0
  UDP sweep                                  0
  IP tear drop                               0
  TCP SYN flood                              0
  IP spoofing                                0
  ICMP ping of death                         0
```


Copyright © 2016, Juniper Networks, Inc.

```
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls1
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
  IP source route option                    0
  TCP land attack                           0
  TCP SYN fragment                          0
  TCP no flag                               0
  IP unknown protocol                       0
  IP bad options                            0
  IP record route option                    0
  IP timestamp option                       0
  IP security option                        0
  IP loose source route option              0
  IP strict source route option             0
  IP stream option                          0
  ICMP fragment                             0
  ICMP large packet                         0
  TCP SYN FIN                               0
  TCP FIN no ACK                            0
  Source session limit                      0
  TCP SYN-ACK-ACK proxy                     0
  IP block fragment                         0
  Destination session limit                 0

Logical system: ls2
Screen statistics:

IDS attack type                             Statistics
```

```
ICMP flood                              0
UDP flood                               0
TCP winnuke                             0
TCP port scan                           0
ICMP address sweep                      0
TCP sweep                               0
UDP sweep                               0
IP tear drop                            0
TCP SYN flood                           0
IP spoofing                             0
ICMP ping of death                      0
IP source route option                  0
TCP land attack                         0
TCP SYN fragment                        0
TCP no flag                             0
IP unknown protocol                     0
IP bad options                          0
IP record route option                  0
IP timestamp option                     0
IP security option                      0
IP loose source route option            0
IP strict source route option           0
IP stream option                        0
ICMP fragment                           0
ICMP large packet                       0
TCP SYN FIN                             0
TCP FIN no ACK                          0
Source session limit                    0
TCP SYN-ACK-ACK proxy                   0
IP block fragment                       0
Destination session limit               0
```


Copyright © 2016, Juniper Networks, Inc.

show system security-profile

Syntax show system security-profile (all-resource | *resource*) <detail | terse> <logical-system (all | *logical-system-name*)> <root-logical-system> <summary>

Release Information Command introduced in Junos OS Release 11.2. Support for application firewall added in Junos OS Release 11.3. Option to display all resources for a logical system added in Junos OS Release 11.. Resource information for ports in source NAT pools with port translation added in Release Junos OS 11.4.

Description Display information about a resource allocated to the logical system in a security profile. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

This command can be used by the master administrator to display resource information for the master logical system or user logical system. This command can also be used by the user logical system administrator to display resource information for a user logical system.

Options Either specify **all-resource** to display information about all resources allocated for the logical system, or specify one of the following resources:

- address-book—Address books.
- appfw-rule-set—Application firewall rule set entries.
- appfw-rule—Application firewall rule entries.
- auth-entry—Firewall authentication entries.
- cpu—CPU utilization.
- flow-gate—Flow gates, also known as pinholes.
- flow-session—Flow sessions.
- nat-cone-binding—Network Address Translation (NAT) cone bindings.
- nat-destination-pool—NAT destination pools.
- nat-destination-rule—NAT destination rules.
- nat-nopat-address—NAT without port address translations.
- nat-pat-address—NAT with port address translations.
- nat-pat-portnum—NAT source port numbers for port translation
- nat-port-ol-ipnumber—NAT port overloading IP numbers.
- nat-rule-referenced-prefix—NAT rule referenced IP-prefixes.
- nat-source-pool—NAT source pools.
- nat-source-rule—NAT source rules.
- nat-static-rule—NAT static rules.

Copyright © 2016, Juniper Networks, Inc.

- policy—Security policies.
- policy-with-count—Security policies with a count.
- scheduler—Schedulers.
- zone—Security zones.

detail | terse—(Optional) Display the specified level of output.

The following options are available only to the master administrator:

- logical-system—Display resource information for a specified user logical system. Specify **all** to display resource information for all logical systems, including the master logical system.
- root-logical-system—Display resource information for the master (root) logical system.
- summary—Display summary information about the resource for all logical systems.

Required Privilege Level view

Related Documentation
- security-profile-resources on page 3631

List of Sample Output
show system security-profile all-resource on page 3731
show system security-profile policy on page 3731
show system security-profile cpu on page 3731
show system security-profile cpu logical-system all on page 3732
show system security-profile cpu summary on page 3732
show system security-profile nat-pat-portnum on page 3732
show system security-profile nat-pat-portnum summary on page 3733

Output Fields Table 372 lists the output fields for the **show system security-profile** command. Output fields are listed in the approximate order in which they appear.

Table 372: show system security-profile Output Fields

Field Name	Field Description
logical system name	Name of the logical system.
security profile name	Name of the security profile bound to the logical system.
usage	Number of resources that are currently being used by the logical system.
reserved	Number of resources that are guaranteed to be available to the logical system.
maximum	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
CPU control	**TRUE** if CPU control is enabled or **FALSE** if CPU control is not enabled.

Copyright © 2016, Juniper Networks, Inc.

Table 372: show system security-profile Output Fields *(continued)*

Field Name	Field Description
CPU control target	Upper limit for CPU utilization on the device. The default value is 80 percent.
CPU name	Central point (CP) or services processing unit (SPU). CP utilization and average utilization of all SPUs is shown. The **detail** option shows CPU utilization on each SPU.
drop rate	Packets dropped for CPU control.

Sample Output

show system security-profile all-resource

```
user@host> show system security-profile all-resource

resource                                  usage    reserved     maximum

[logical system name:   root-logical-system]
[security profile name: Default-Profile]
address-book                                  0           0         512
auth-entry                                    0           0  2147483647
cpu on CP                                 0.00%       1.00%      80.00%
cpu on SPU                                0.00%       1.00%      80.00%
flow-gate                                     0           0      524288
flow-session                                  2           0     6291456
nat-cone-binding                              0           0       65536
nat-destination-pool                          0           0        4096
nat-destination-rule                          0           0        8192
nat-nopat-address                             0           0     1048576
nat-pat-address                               0           0        2048
nat-port-ol-ipnumber                          0           0           4
nat-rule-referenced-prefix                    0           0     1048576
nat-source-pool                               0           0        2048
nat-source-rule                               0           0        8192
nat-static-rule                               0           0       20480
policy                                        0           0       40000
policy-with-count                             0           0        1024
scheduler                                     0           0          64
zone                                          0           0         512
```

show system security-profile policy

```
user@host> show system security-profile policy
logical system name   security profile name       usage    reserved     maximum

ls-product-design     ls-design-profile               0          40          50
```

show system security-profile cpu

```
user@host> show system security-profile cpu
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%
     0.00%
```

Copyright © 2016, Juniper Networks, Inc.

```
root-logical-system    Default-Profile SPU            0.00%          1.00%
     0.00%
```

show system security-profile cpu logical-system all

```
user@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    Default-Profile CP             0.00%          1.00%
     0.00%
root-logical-system    Default-Profile SPU            0.00%          1.00%
     0.00%
ls-product-design      ls-design-profile CP           0.00%          0.00%
     0.00%
ls-product-design      ls-design-profile SPU          0.00%          0.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP        0.00%          0.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU       0.00%          0.00%
     0.00%
```

Should the above output actually look as follows?

```
logical system name    security profile name    usage    reserved    maximum
root-logical-system    Default-Profile         67108864     0      134217728
lsys1                  profile1                193536     6000     134217728
```

show system security-profile cpu summary

```
user@host> show system security-profile cpu summary
CPU control: TRUE
CPU control target: 80.00%

CPU type                :     CP
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system

CPU type                :    SPU
global used amount      :  0.00%
global maximum quota    : 80.00%
global available amount : 80.00%
total logical systems   :      3
total security profiles :      3
heaviest usage / user   :  0.00%      / root-logical-system
lightest usage / user   :  0.00%      / root-logical-system
```

show system security-profile nat-pat-portnum

```
user@host> show system security-profile cpu nat-pat-portnum
CPU control: TRUE
CPU control target: 80.00%
logical system name    security profile name        usage(%)    reserved(%)
maximum
root-logical-system    Default-Profile CP             67108864          0
    134217728
```


Copyright © 2016, Juniper Networks, Inc.

show system security-profile nat-pat-portnum summary

```
user@host> show system security-profile nat-pat-portnum summary
global used amount    :67302400
global maximum quota    :134217728
global available amount   :66915328
total logical systems    :2
total security profiles   :1
heaviest usage / user    :193536 / lsys1
```

Copyright © 2016, Juniper Networks, Inc.

show security softwires

Syntax show security softwires <softwire-name *softwire-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced in Junos OS Release 10.4. The **logical-system** option introduced in Junos OS Release 12.1.

Description Display a summary of information of all the softwire concentrators and details on concentrators with specified name.

Options **softwire-name *softwire-name***—Display the details of the specified softwire concentrator.

logical-system (all | *logical-system-name*)—Display softwire information for all logical systems or for a specified logical system. This option is only available to the master administrator.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

Sample Output

```
user@host> show security softwires
Softwire Name             SC Address         Status   Number of SI connected
  SC-CSSI-1                  3001::1          Connected      2
  SC-CSSI-str00              3100::1          Active         0
  SC-CSSI-str01              3101::1          Inactive       0
  SC-CSSI-str02              3001::1          Connected      2520
```

```
user@host> show security softwires softwire-name SC-CSSI-1
Name of softwire: SC-CSSI-1
   SC status: Connected
   SC address: 3001::1
   Zone: trust
   VR ID: 0
 SI Address               SI Status                  SPU
 3001::2                  Active                     spu-1
 3001::2                  Active                     spu-21
 SI number: 2
```

```
user@host> show security softwires logical-system ls-product-design
Softwire Name             SC Address         Status   Number of SI connected
  sc_1                       3000::1          Connected      1
```


Copyright © 2016, Juniper Networks, Inc.

show security zones

Syntax
show security zones
<detail | terse>
< *zone-name* >

Release Information
Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.

Description
Display information about security zones.

Options
- none—Display information about all zones.
- **detail | terse**—(Optional) Display the specified level of output.
- ***zone-name*** —(Optional) Display information about the specified zone.

Required Privilege Level
view

Related Documentation
- Security Zones and Interfaces Overview on page 1049
- Supported System Services for Host Inbound Traffic on page 1061
- security-zone on page 385

List of Sample Output
show security zones on page 3736
show security zones abc on page 3736
show security zones abc detail on page 3736
show security zones terse on page 3737

Output Fields
Table 32 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 373: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                     Type
my-internal              Security
my-external              Security
dmz                      Security
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

MPLS Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 53

Overview

- Introduction to MPLS on page 3743

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 157

Introduction to MPLS

- MPLS Overview on page 3743
- MPLS Configuration Overview on page 3748
- Example: Deleting Security Services on page 3748
- Example: Enabling MPLS on page 3750

MPLS Overview

Multiprotocol Label Switching (MPLS) is a method for engineering traffic patterns by assigning short labels to network packets that describe how to forward them through the network. MPLS is independent of routing tables or any routing protocol and can be used for unicast packets.

The MPLS framework supports traffic engineering and the creation of virtual private networks (VPNs). Traffic is engineered (controlled) primarily by the use of signaling protocols to establish label-switched paths (LSPs). VPN support includes Layer 2 and Layer 3 VPNs and Layer 2 circuits.

When you enable your device to allow MPLS traffic, the device performs packet-based processing and functions as a standard Junos router.

CAUTION: When packet forwarding mode is changed to MPLS, all flow-based security features are deactivated, and the device performs packet-based processing only. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device. However, MPLS can be enabled in flow-based packet forwarding mode for selected traffic using firewall filters.

This overview contains the following topics:

- Label Switching on page 3744
- Label-Switched Paths on page 3744
- Label-Switching Routers on page 3745
- Labels on page 3746
- Label Operations on page 3746

Copyright © 2016, Juniper Networks, Inc.

- Penultimate Hop Popping on page 3747
- LSP Establishment on page 3747

Label Switching

In a traditional IP network, packets are transmitted with an IP header that includes a source and destination address. When a router receives such a packet, it examines its forwarding tables for the next-hop address associated with the packet's destination address and forwards the packet to the next-hop location.

In an MPLS network, each packet is encapsulated with an MPLS header. When a router receives the packet, it copies the header as an index into a separate MPLS forwarding table. The MPLS forwarding table consists of pairs of inbound interfaces and path information. Each pair includes forwarding information that the router uses to forward the traffic and modify, when necessary, the MPLS header.

Because the MPLS forwarding table has far fewer entries than the more general forwarding table, the lookup consumes less processing time and processing power. The resultant savings in time and processing are a significant benefit for traffic that uses the network to transit between outside destinations only.

Label-Switched Paths

Label-switched paths (LSPs) are unidirectional routes through a network or autonomous system (AS). In normal IP routing, the packet has no predetermined path. Instead, each router forwards a packet to the next-hop address stored in its forwarding table, based only on the packet's destination address. Each subsequent router then forwards the packet using its own forwarding table.

In contrast, MPLS routers within an AS determine paths through a network through the exchange of MPLS traffic engineering information. Using these paths, the routers direct traffic through the network along an established route. Rather than selecting the next hop along the path as in IP routing, each router is responsible for forwarding the packet to a predetermined next-hop address.

Figure 138 shows a typical LSP topology.

Copyright © 2016, Juniper Networks, Inc.

Figure 138: Typical LSP Topology

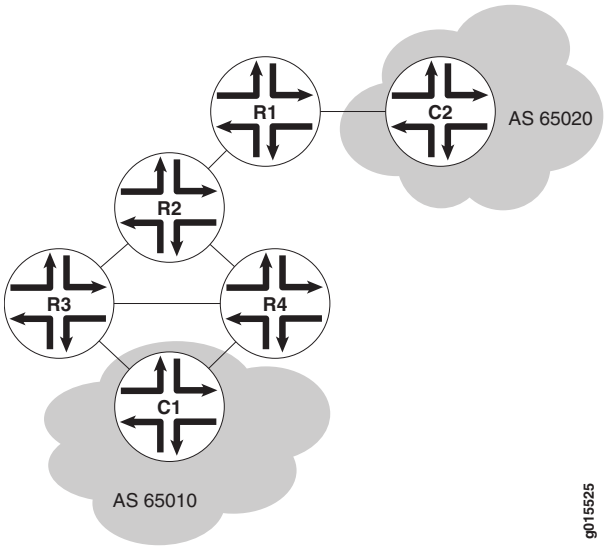

In the topology shown in Figure 138, traffic is forwarded from Host C1 to the transit network with standard IP forwarding. When the traffic enters the transit network, it is switched across a preestablished LSP through the network. In this example, an LSP might switch the traffic from Router R4 to Router R2 to Router R1. When the traffic exits the network, it is forwarded to its destination by IP routing protocols.

Label-Switching Routers

Routers that are part of the LSP are label-switching routers (LSRs). Each LSR must be configured with MPLS so that it can interpret MPLS headers and perform the MPLS operations required to pass traffic through the network. An LSP can include four types of LSRs:

- Inbound router—The only entry point for traffic into MPLS. Native IPv4 packets are encapsulated into the MPLS protocol by the inbound router. Each LSP can have only one inbound router. Inbound routers are also known as ingress routers.
- Transit router—Any router in the middle of an LSP. An individual LSP can contain between 0 and 253 transit routers. Transit routers forward MPLS traffic along the LSP, using only the MPLS header to determine how the packet is routed.
- Penultimate router—The second-to-last router in the LSP. The penultimate router in an LSP is responsible for stripping the MPLS header from the packet before forwarding it to the outbound router.
- Outbound router—The endpoint for the LSP. The outbound router receives MPLS packets from the penultimate router and performs an IP route lookup. The router then forwards the packet to the next hop of the route. Each LSP can have only one outbound router. Outbound routers are also known as egress routers.

Copyright © 2016, Juniper Networks, Inc.

Labels

To forward traffic through an MPLS network, MPLS routers encapsulate packets and assign and manage headers known as *labels*. A label is a 20–bit unsigned integer in the range 0 through 1,048,575. The routers use the labels to index the MPLS forwarding tables that determine how packets are routed through the network.

When a network's inbound router receives traffic, it inserts an MPLS label between the IP packet and the appropriate Layer 2 header for the physical link. The label contains an index value that identifies a next-hop address for the particular LSP. When the next-hop transit router receives the packet, it uses the index in the MPLS label to determine the next-hop address for the packet and forwards the packet to the next router in the LSP.

As each packet travels through the transit network, every router along the way performs a lookup on the MPLS label and forwards the packet accordingly. When the outbound router receives a packet, it examines the header to determine that it is the final router in the LSP. The outbound router then removes the MPLS header, performs a regular IP route lookup, and forwards the packet with its IP header to the next-hop address.

Label Operations

Each LSR along an LSP is responsible for examining the MPLS label, determining the LSP next hop, and performing the required label operations. LSRs can perform five label operations:

- Push—Adds a new label to the top of the packet. For IPv4 packets arriving at the inbound router, the new label is the first label in the label stack. For MPLS packets with an existing label, this operation adds a label to the stack and sets the stacking bit to 0, indicating that more MPLS labels follow the first.

  When it receives the packet, the inbound router performs an IP route lookup on the packet. Because the route lookup yields an LSP next hop, the inbound router performs a label push on the packet, and then forwards the packet to the LSP next hop.

- Swap—Replaces the label at the top of the label stack with a new label.

  When a transit router receives the packet, it performs an MPLS forwarding table lookup. The lookup yields the LSP next hop and the path index of the link between the transit router and the next router in the LSP.

- Pop—Removes the label from the top of the label stack. For IPv4 packets arriving at the penultimate router, the entire MPLS label is removed from the label stack. For MPLS packets with an existing label, this operation removes the top label from the label stack and modifies the stacking bit as necessary—sets it to 1, for example, if only a single label remains in the stack.

  If multiple LSPs terminate at the same outbound router, the router performs MPLS label operations for all outbound traffic on the LSPs. To share the operations among multiple routers, most LSPs use penultimate hop popping (PHP).

- Multiple push—Adds multiple labels to the top of the label stack. This action is equivalent to performing multiple push operations.

Copyright © 2016, Juniper Networks, Inc.

The multiple push operation is used with label stacking, which is beyond the scope of this topic.

- Swap and push—Replaces the top label with a new label and then pushes a new label to the top of the stack.

  The swap and push operation is used with label stacking, which is beyond the scope of this topic.

Penultimate Hop Popping

Multiple LSPs terminating at a single outbound router load the router with MPLS label operations for all their outbound traffic. Penultimate hop popping (PHP) transfers the operation from the outbound router to penultimate routers.

With PHP, the penultimate router is responsible for popping the MPLS label and forwarding the traffic to the outbound router. The outbound router then performs an IP route lookup and forwards the traffic. For example, if four LSPs terminate at the same outbound router and each has a different penultimate router, label operations are shared across four routers.

LSP Establishment

An MPLS LSP is established by one of two methods: static LSPs and dynamic LSPs.

Static LSPs

Like a static route, a static LSP requires each router along the path to be configured explicitly. You must manually configure the path and its associated label values. Static LSPs require less processing by the LSRs because no signaling protocol is used. However, because paths are statically configured, they cannot adapt to network conditions. Topology changes and network outages can create black holes in the LSP that exist until you manually reconfigure the LSP.

Dynamic LSPs

Dynamic LSPs use signaling protocols to establish themselves and propagate LSP information to other LSRs in the network. You configure the inbound router with LSP information that is transmitted throughout the network when you enable the signaling protocols across the LSRs. Because the LSRs must exchange and process signaling packets and instructions, dynamic LSPs consume more resources than static LSPs. However, dynamic LSPs can avoid the network black holes of static LSPs by detecting topology changes and outages and propagating them throughout the network.

Related Documentation

- MPLS Configuration Overview on page 3748
- Example: Deleting Security Services on page 3748
- Example: Enabling MPLS on page 3750
- MPLS Traffic Engineering and Signaling Protocols Overview on page 3755
- MPLS VPN Overview on page 3791

Copyright © 2016, Juniper Networks, Inc.

MPLS Configuration Overview

When you first install Junos OS on your device, MPLS is disabled by default. You must explicitly configure your device to allow MPLS traffic to pass through. Complete the following steps for all devices in your MPLS network that are running Junos OS.

To enable MPLS:

1. Delete all configured security services from the device. If you do not complete this step, you will get a commit failure. See "Example: Deleting Security Services" on page 3748.
2. Enable MPLS on the device. See "Example: Enabling MPLS" on page 3750.
3. Commit the configuration.
4. Reboot the device.
5. Configure MPLS features such as traffic engineering, VPNs, and VPLS. See:
   - MPLS Traffic Engineering and Signaling Protocols Overview on page 3755
   - MPLS VPN Overview on page 3791
   - CLNS Overview on page 3815
   - VPLS Overview on page 3841

CAUTION: When packet forwarding mode is changed to MPLS, all flow-based security features are deactivated, and the device performs packet-based processing only. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device. However, MPLS can be enabled in flow-based packet forwarding mode for selected traffic using firewall filters.

Related Documentation

- MPLS Overview on page 3743
- Example: Deleting Security Services on page 3748
- Example: Enabling MPLS on page 3750

Example: Deleting Security Services

This example shows how to delete configured services in the security level of the configuration hierarchy.

- Requirements on page 3749
- Overview on page 3749
- Configuration on page 3749
- Verification on page 3749

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, save your current configuration to a temporary file. Do this prior to removing all configurations from the security level of the configuration hierarchy and deleting the inherited configurations.

Overview

In this example, you save your current configuration in the var/tmp/ directory with an appropriate filename and .cfg extension—for example, *curfeb08.cfg*. Then you remove all configurations from the **security** level of the configuration hierarchy, and delete all global groups and inherited configurations.

Configuration

Step-by-Step Procedure

To delete the configured services in the security level of the configuration hierarchy:

1. Save your current configuration.

   ```
   [edit]
   user@host# save /var/tmp/curfeb08.cfg
   ```

2. Remove all configurations in the **security** level of the configuration hierarchy.

   ```
   [edit]
   user@host# delete security
   ```

3. Remove all inherited configurations in the security level of the configuration hierarchy.

   ```
   [edit]
   user@host# delete groups global security
   ```

CAUTION: Do not commit after deleting the security configurations. A commit without any security configurations leaves the router unreachable through the management port.

Verification

To verify the configuration is working properly, enter the **show groups global security** command.

Related Documentation

- MPLS Overview on page 3743
- MPLS Configuration Overview on page 3748
- Example: Enabling MPLS on page 3750

Copyright © 2016, Juniper Networks, Inc.

Example: Enabling MPLS

This example shows how to enable MPLS for packet-based processing. It also shows how to enable the MPLS family and MPLS process on all of the transit interfaces in the network.

NOTE: When MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, IP packets, and IPsec VPNs are unavailable on the device.

Before changing from flow mode to packet mode, you must remove all security policies remaining under flow mode. To prevent management connection loss, you must bind the management interface to zones and enable host-inbound traffic to prevent the device from losing connectivity.

For information about configuring zones, see *Building Blocks Feature Guide for Security Devices*.

Requirements

Before you begin, delete all configured security services. See "Example: Deleting Security Services" on page 3748.

Overview

The instructions in this topic describe how to enable MPLS on the device. You must enable MPLS on the device before including a device running Junos OS in an MPLS network.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security forwarding-options family mpls mode packet-based
set interfaces ge-1/0/0 unit 0 family mpls
set protocols mpls ge-1/0/0 unit 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To enable MPLS:

1. Enable MPLS for packet-based processing.

```
[edit security forwarding-options]
user@host# set family mpls mode packet-based
```


Copyright © 2016, Juniper Networks, Inc.

2. Enable the MPLS family on each transit interface that you want to include in the MPLS network.

```
[edit interfaces]
user@host# set interfaces ge-1/0/0 unit 0 family mpls
```

3. Enable the MPLS process on all of the transit interfaces in the MPLS network.

```
[edit protocols mpls]
user@host# set interface ge-1/0/0 unit 0
```

Results From configuration mode, confirm your configuration by entering the **show security forwarding-options** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

NOTE: If you enable MPLS for packet-based processing by using the command **set security forward-option family mpls mode packet**, the mode will not change immediately and the system will display the following messages:

warning: Reboot may required when try reset flow inet mode

warning: Reboot may required when try reset mpls flow mode please check security flow status for detail.

You need to reboot your device for the configuration to take effect.

CAUTION: If you disable MPLS and switch back to using the security services (flow-based processing), the mode will not change immediately and the system will display warning messages instructing you to restart your device. You must reboot your device for the configuration to take effect. This will also result in management sessions being reset and transit traffic getting interrupted.

```
[edit]
user@host# show security forwarding-options
family {
   mpls {
      mode packet-based;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying MPLS Is Enabled at the Protocols Level on page 3752
- Verifying MPLS Is Enabled at the Interfaces Level on page 3752

Verifying MPLS Is Enabled at the Protocols Level

Purpose Verify that MPLS is enabled at the protocols level.

Action From operational mode, enter the **show protocols** command.

Verifying MPLS Is Enabled at the Interfaces Level

Purpose Verify that MPLS is enabled at the interfaces level.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- MPLS Overview on page 3743
- MPLS Configuration Overview on page 3748
- Example: Deleting Security Services on page 3748

Copyright © 2016, Juniper Networks, Inc.

PART 54

Configuring Traffic Engineering

- Configuring MPLS Traffic Engineering and Signaling Protocols on page 3755

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 158

Configuring MPLS Traffic Engineering and Signaling Protocols

- MPLS Traffic Engineering and Signaling Protocols Overview on page 3755
- Understanding the LDP Signaling Protocol on page 3756
- Example: Configuring LDP-Signaled LSPs on page 3757
- Understanding the RSVP Signaling Protocol on page 3761
- Example: Configuring RSVP-Signaled LSPs on page 3765
- Understanding Point-to-Multipoint LSPs on page 3768
- Point-to-Multipoint LSP Configuration Overview on page 3770
- Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on page 3770

MPLS Traffic Engineering and Signaling Protocols Overview

Traffic engineering facilitates efficient and reliable network operations while simultaneously optimizing network resources and traffic performance. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) to a potentially less congested physical path across a network. To support traffic engineering, besides source routing, the network must do the following:

- Compute a path at the source by taking into account all the constraints, such as bandwidth and administrative requirements.
- Distribute the information about network topology and link attributes throughout the network once the path is computed.
- Reserve network resources and modify link attributes.

When transit traffic is routed through an IP network, MPLS is often used to engineer its passage. Although the exact path through the transit network is of little importance to either the sender or the receiver of the traffic, network administrators often want to route traffic more efficiently between certain source and destination address pairs. By adding a short label with specific routing instructions to each packet, MPLS switches packets from router to router through the network rather than forwarding packets based on next-hop lookups. The resulting routes are called label-switched paths (LSPs). LSPs control the passage of traffic through the network and speed traffic forwarding.

Copyright © 2016, Juniper Networks, Inc.

You can create LSPs manually, or through the use of signaling protocols. Signaling protocols are used within an MPLS environment to establish LSPs for traffic across a transit network. Junos OS supports two signaling protocols—LDP and the Resource Reservation Protocol (RSVP).

MPLS traffic engineering uses the following components:

- MPLS LSPs for packet forwarding
- IGP extensions for distributing information about the network topology and link attributes
- Constrained Shortest Path First (CSPF) for path computation and path selection
- RSVP extensions to establish the forwarding state along the path and to reserve resources along the path

Junos OS also supports traffic engineering across different OSPF regions.

Related Documentation

- Understanding the LDP Signaling Protocol on page 3756
- Understanding the RSVP Signaling Protocol on page 3761
- Understanding Point-to-Multipoint LSPs on page 3768

Understanding the LDP Signaling Protocol

LDP is a signaling protocol that runs on a device configured for MPLS support. The successful configuration of both MPLS and LDP initiates the exchange of TCP packets across the LDP interfaces. The packets establish TCP-based LDP sessions for the exchange of MPLS information within the network. Enabling both MPLS and LDP on the appropriate interfaces is sufficient to establish LSPs.

LDP is a simple, fast-acting signaling protocol that automatically establishes LSP adjacencies within an MPLS network. Routers then share LSP updates such as hello packets and LSP advertisements across the adjacencies. Because LDP runs on top of an IGP such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces. After both are configured, LDP begins transmitting and receiving LDP messages through all LDP-enabled interfaces. Because of LDP's simplicity, it cannot perform the true traffic engineering which RSVP can perform. LDP does not support bandwidth reservation or traffic constraints.

When you configure LDP on a label-switching router (LSR), the router begins sending LDP discovery messages out all LDP-enabled interfaces. When an adjacent LSR receives LDP discovery messages, it establishes an underlying TCP session. An LDP session is then created on top of the TCP session. The TCP three-way handshake ensures that the LDP session has bidirectional connectivity. After they establish the LDP session, the LDP neighbors maintain, and terminate, the session by exchanging messages. LDP advertisement messages allow LSRs to exchange label information to determine the next hops within a particular LSP. Any topology changes, such as a router failure, generate LDP notifications that can terminate the LDP session or generate additional LDP advertisements to propagate an LSP change.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- MPLS Traffic Engineering and Signaling Protocols Overview on page 3755
- Example: Configuring LDP-Signaled LSPs on page 3757

Example: Configuring LDP-Signaled LSPs

This example shows how to create and configure LDP instances within an MPLS network.

- Requirements on page 3757
- Overview on page 3757
- Configuration on page 3758
- Verification on page 3759

Requirements

Before you begin:

- Configure network interfaces. See *Interfaces Feature Guide for Security Devices*.
- Configure an IGP across your network. (The LDP configuration is added to the existing IGP configuration and included in the MPLS configuration.)
- Configure a network to use LDP for LSP establishment by enabling MPLS on all transit interfaces in the MPLS network.

NOTE: **Because LDP runs on top of an IGP such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces.**

Overview

To configure LDP-signaled LSPs, you must enable the MPLS family on all transit interfaces in the MPLS network and include all the transit interfaces under the [**protocols mpls**] and [**protocols ldp**] hierarchy levels.

In this example, you enable the MPLS family and create an LDP instance on all the transit interfaces. Additionally, you enable the MPLS process on all the transit interfaces in the MPLS network. In this example, you configure a sample network as shown in Figure 139.

Copyright © 2016, Juniper Networks, Inc.

Figure 139: Typical LDP-Signaled LSP

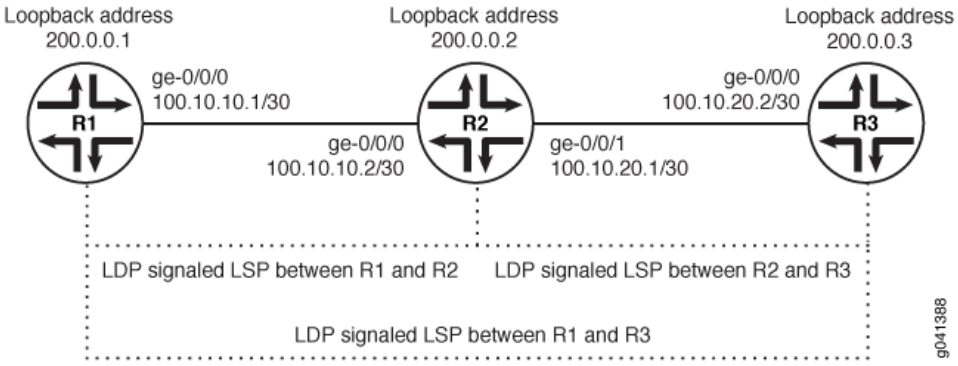

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level and then enter **commit** from configuration mode.

For Router R1, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
```

For Router R2, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
set interfaces ge-0/0/1 unit 0 family mpls
set protocols mpls ge-0/0/1 unit 0
set protocols ldp interface ge-0/0/0.1 unit 0
```

For Router R3, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
```

Step-by-Step Procedure

To enable LDP instances within an MPLS network:

1. Enable the MPLS family on the transit interface on Router R1.

   ```
   [edit]
   user@R1# set interfaces ge-0/0/0 unit 0 family mpls
   ```

2. Enable the MPLS process on the transit interface.

   ```
   [edit]
   user@R1# set protocols mpls interface ge-0/0/0 unit 0
   ```

3. Create the LDP instance on the transit interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@R1# set protocols ldp interface ge-0/0/0 unit 0
```

Results Confirm your configuration by entering the **show** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@R1# show
...
  interfaces {
    ge-0/0/0 {
      unit 0 {
        family inet {
          address 10.100.37.20/24;
        }
        family mpls;
      }
    }
    }
  ...
    protocols {
      mpls {
        interface all;
      }
      ldp {
        interface ge-0/0/0.0;
      }
    }
```

If you are done configuring the device, enter the **commit** command from the configuration mode to activate the configuration.

Verification

Confirm that the configuration is working properly.

- Verifying LDP Neighbors on page 3759
- Verifying LDP Sessions on page 3760
- Verifying the Presence of LDP-Signaled LSPs on page 3761
- Verifying Traffic Forwarding over the LDP-Signaled LSP on page 3761

Verifying LDP Neighbors

Purpose Verify that each device shows the appropriate LDP neighbors.

Action From the CLI, enter the **show ldp neighbor** command.

```
user@r2> show ldp neighbor
```

```
Address     Interface       Label space ID      Hold time
100.10.10.1    ge-0/0/0.0      200.0.0.1:0           11
100.10.20.2   ge-0/0/1.0      200.0.0.2:0           14
```

The output shows the IP addresses of the neighboring interfaces along with the interface through which the neighbor adjacency is established. Verify the following information:

- Each interface on which LDP is enabled is listed.
- Each neighboring LDP interface address is listed with the appropriate corresponding LDP interface.
- Under **Label space ID**, the appropriate loopback address for each neighbor appears.

Verifying LDP Sessions

Purpose Verify that a TCP-based LDP session has been established between all LDP neighbors. Also, verify that the modified keepalive value is active.

Action From the CLI, enter the **show ldp session detail** command.

```
user@r1> show ldp session detail
200.0.0.2, State: Operational, Connection: Open, Hold time: 22
  Session ID: 200.0.0.1:0--200.0.0.2:0
  Next keepalive in 9 seconds
  Active, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 200.0.0.1, Remote address: 200.0.0.2
  Up for 01:58:49
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: disabled, Helper mode: enabled
  Remote - Restart: disabled, Helper mode: enabled
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Nonstop routing state: Not in sync
  Next-hop addresses received:
 100.10.10.2
 100.10.20.1
```

The output shows the detailed information, including session IDs, keepalive interval, and next-hop addresses, for each established LDP session. Verify the following information:

- Each LDP neighbor address has an entry, listed by loopback address.
- The state for each session is **Operational**, and the connection for each session is **Open**. A state of **Nonexistent** or a connection of **Closed** indicates a problem with one of the following:
  - LDP configuration
  - Passage of traffic between the two devices
  - Physical link between the two routers
- For **Keepalive interval**, the appropriate value, **10**, appears.

Copyright © 2016, Juniper Networks, Inc.

Verifying the Presence of LDP-Signaled LSPs

Purpose Verify that each Juniper Networks device's **inet.3** routing table has an LSP for the loopback address on each of the other routers.

Action From the CLI, enter the **show route table inet.3** command.

```
user@r1> run show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
200.0.0.2/32       *[LDP/9] 05:20:20, metric 1
                    > via ge-0/0/0.0, Push 300640
200.0.0.3/32       *[LDP/9] 05:20:20, metric 1
                    > via ge-0/0/0.0, Push 300704
```

The output shows the LDP routes that exist in the **inet.3** routing table. Verify that an LDP-signaled LSP is associated with the loopback addresses of the other routers in the MPLS network.

Verifying Traffic Forwarding over the LDP-Signaled LSP

Purpose Verify that the LDP path between R1 and R3 is complete over the LDP-signaled LSP.

Action From the CLI on R1, enter the **traceroute mpls ldp 200.0.0.3** command.

```
user@r1> traceroute mpls ldp 200.0.0.3
Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl Label Protocol Address Previous Hop Probe Status
1 300704 LDP 200.0.0.1 (null) Unhelpful
2 200.0.0.2 200.0.0.1 Unhelpful
3 200.0.0.3 200.0.0.2 Egress
Path 1 via ge-0/0/0.0 destination 127.0.0.64
```

The output shows the path between R1 and R3 using an LDP signalled LSP.

Related Documentation

- Understanding the LDP Signaling Protocol on page 3756
- *Routing Protocols Overview for Security Devices*

Understanding the RSVP Signaling Protocol

RSVP is a signaling protocol that handles bandwidth allocation and true traffic engineering across an MPLS network. Like LDP, RSVP uses discovery messages and advertisements to exchange LSP path information between all hosts. However, RSVP also includes a set of features that control the flow of traffic through an MPLS network. Whereas LDP is restricted to using the configured IGP's shortest path as the transit path through the network, RSVP uses a combination of the Constrained Shortest Path First (CSPF) algorithm and Explicit Route Objects (EROs) to determine how traffic is routed through the network.

Basic RSVP sessions are established in exactly the same way that LDP sessions are established. By configuring both MPLS and RSVP on the appropriate transit interfaces,

Copyright © 2016, Juniper Networks, Inc.

you enable the exchange of RSVP packets and the establishment of LSPs. However, RSVP also lets you configure link authentication, explicit LSP paths, and link coloring.

This topic contains the following sections:

- RSVP Fundamentals on page 3762
- Bandwidth Reservation Requirement on page 3762
- Explicit Route Objects on page 3762
- Constrained Shortest Path First on page 3763
- Link Coloring on page 3764

RSVP Fundamentals

RSVP uses unidirectional and simplex flows through the network to perform its function. The inbound router initiates an RSVP path message and sends it downstream to the outbound router. The path message contains information about the resources needed for the connection. Each router along the path begins to maintain information about a reservation.

When the path message reaches the outbound router, resource reservation begins. The outbound router sends a reservation message upstream to the inbound router. Each router along the path receives the reservation message and sends it upstream, following the path of the original path message. When the inbound router receives the reservation message, the unidirectional network path is established.

The established path remains open as long as the RSVP session is active. The session is maintained by the transmission of additional path and reservation messages that report the session state every 30 seconds. If a router does not receive the maintenance messages for 3 minutes, it terminates the RSVP session and reroutes the LSP through another active router.

Bandwidth Reservation Requirement

When a bandwidth reservation is configured, reservation messages propagate the bandwidth value throughout the LSP. Routers must reserve the bandwidth specified across the link for the particular LSP. If the total bandwidth reservation exceeds the available bandwidth for a particular LSP segment, the LSP is rerouted through another LSR. If no segments can support the bandwidth reservation, LSP setup fails and the RSVP session is not established.

Explicit Route Objects

Explicit Route Objects (EROs) limit LSP routing to a specified list of LSRs. By default, RSVP messages follow a path that is determined by the network IGP's shortest path. However, in the presence of a configured ERO, the RSVP messages follow the path specified.

EROs consist of two types of instructions: loose hops and strict hops. When a loose hop is configured, it identifies one or more transit LSRs through which the LSP must be routed. The network IGP determines the exact route from the inbound router to the first loose

Copyright © 2016, Juniper Networks, Inc.

hop, or from one loose hop to the next. The loose hop specifies only that a particular LSR be included in the LSP.

When a strict hop is configured, it identifies an exact path through which the LSP must be routed. Strict-hop EROs specify the exact order of the routers through which the RSVP messages are sent.

You can configure loose-hop and strict-hop EROs simultaneously. In this case, the IGP determines the route between loose hops, and the strict-hop configuration specifies the exact path for particular LSP path segments.

Figure 140 shows a typical RSVP-signaled LSP that uses EROs.

Figure 140: Typical RSVP-Signaled LSP with EROs

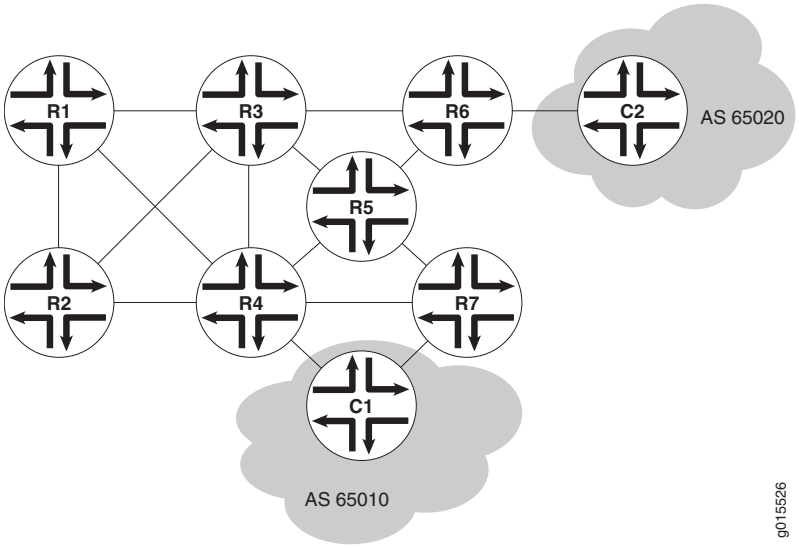

In the topology shown in Figure 140, traffic is routed from Host C1 to Host C2. The LSP can pass through Routers R4 or Router R7. To force the LSP to use R4, you can set up either a loose-hop or strict-hop ERO that specifies R4 as a hop in the LSP. To force a specific path through Router R4, R3, and R6, configure a strict-hop ERO through the three LSRs.

Constrained Shortest Path First

Whereas IGPs use the Shortest Path First (SPF) algorithm to determine how traffic is routed within a network, RSVP uses the Constrained Shortest Path First (CSPF) algorithm to calculate traffic paths that are subject to the following constraints:

- LSP attributes—Administrative groups such as link coloring, bandwidth requirements, and EROs
- Link attributes—Colors on a particular link and available bandwidth

These constraints are maintained in the traffic engineering database (TED). The database provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors.

Copyright © 2016, Juniper Networks, Inc.

In determining which path to select, CSPF follows these rules:

- Computes LSPs one at a time, beginning with the highest priority LSP—the one with the lowest setup priority value. Among LSPs of equal priority, CSPF starts with those that have the highest bandwidth requirement.
- Prunes the traffic engineering database of links that are not full duplex and do not have sufficient reservable bandwidth.
- If the LSP configuration includes the **include** statement, prunes all links that do not share any included colors.
- If the LSP configuration includes the **exclude** statement, prunes all links that contain excluded colors. If a link does not have a color, it is accepted.
- Finds the shortest path toward the LSP's outbound router, taking into account any EROs. For example, if the path must pass through Router A, two separate SPF algorithms are computed: one from the inbound router to Router A and one from Router A to the outbound router.
- If several paths have equal cost, chooses the one with a last-hop address the same as the LSP's destination.
- If several equal-cost paths remain, selects the path with the least number of hops.
- If several equal-cost paths remain, applies CSPF load-balancing rules configured on the LSP.

Link Coloring

RSVP allows you to configure administrative groups for CSPF path selection. An administrative group is typically named with a color, assigned a numeric value, and applied to the RSVP interface for the appropriate link. Lower numbers indicate higher priority.

After configuring the administrative group, you can either exclude, include, or ignore links of that color in the TED:

- If you exclude a particular color, all segments with an administrative group of that color are excluded from CSPF path selection.
- If you include a particular color, only those segments with the appropriate color are selected.
- If you neither exclude nor include the color, the metrics associated with the administrative groups and applied on the particular segments determine the path cost for that segment.

The LSP with the lowest total path cost is selected into the TED.

Related Documentation

- MPLS Traffic Engineering and Signaling Protocols Overview on page 3755
- Example: Configuring RSVP-Signaled LSPs on page 3765

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring RSVP-Signaled LSPs

This example shows how to create an LSP between routers in an IP network using RSVP as the signaling protocol.

- Requirements on page 3765
- Overview and Topology on page 3765
- Configuration on page 3766
- Verification on page 3767

Requirements

Before you begin, delete security services from the device. See "Example: Deleting Security Services" on page 3748.

Overview and Topology

Using RSVP as a signaling protocol, you can create LSPs between routers in an IP network. In this example, you configure a sample network as shown in Figure 141.

Figure 141: Typical RSVP-Signaled LSP

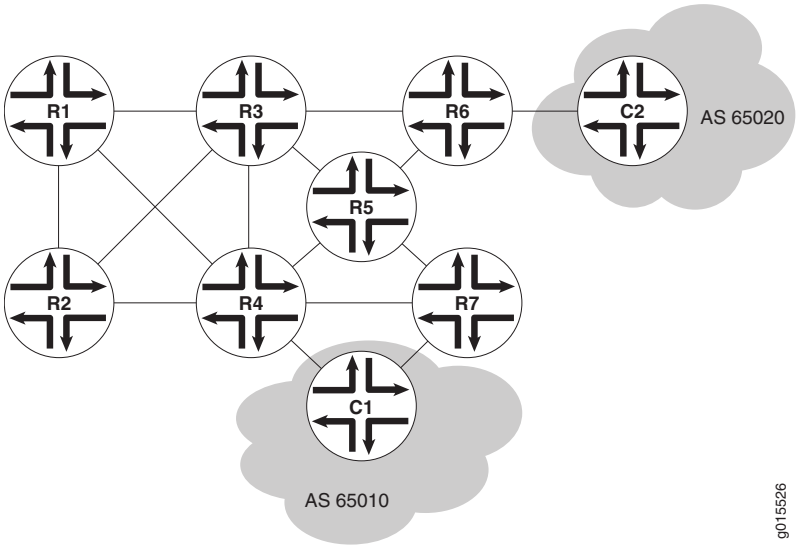

To establish an LSP between routers, you must individually enable the MPLS family and configure RSVP on each of the transit interfaces in the MPLS network. This example shows how to enable MPLS and configure RSVP on the ge-0/0/0 transit interface. Additionally, you must enable the MPLS process on all of the MPLS interfaces in the network.

This example shows how to define an LSP from R1 to R7 on the ingress router (R1) using R7's loopback address (10.0.9.7). The configuration reserves 10 Mbps of bandwidth. Additionally, the configuration disables the CSPF algorithm, ensuring that Hosts C1 and C2 use the RSVP-signaled LSP that correspond to the network IGP's shortest path.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols rsvp interface ge-0/0/0.0
set protocols mpls label-switched-path r1-r7 to 10.0.9.7
set protocols mpls label-switched-path r1-r7 bandwidth 10m
set protocols mpls interface all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure RSVP:

1. Enable the MPLS family on all transit interfaces in the MPLS network.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family mpls
   ```

2. Enable RSVP on each transit interface in the MPLS network.

   ```
   [edit]
   user @host# set protocols rsvp interface ge-0/0/0
   ```

3. Enable the MPLS process on the transit interface in the MPLS network.

   ```
   [edit]
   user@host# set protocols mpls interface ge-0/0/0
   ```

4. Define the LSP on the ingress router.

   ```
   [edit protocols mpls]
   user@host# set label-switched-path r1-r7 to 10.0.9.7
   ```

5. Reserve 10 Mbps of bandwidth on the LSP.

   ```
   [edit protocols mpls]
   user @host# set label-switched-path r1-r7 bandwidth 10m
   ```

Results

Confirm your configuration by entering the **show** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@host# show
...
  interfaces {
  ge-0/0/0 {
     family mpls;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  }
  ...
  protocols {
    rsvp {
      interface ge-0/0/0.0;
    }
    mpls {
      label-switched-path r1-r7 {
        to 10.0.9.7;
        bandwidth 10m;
      }
      interface all;
    }
  }
...
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying RSVP Neighbors on page 3767
- Verifying RSVP Sessions on page 3767
- Verifying the Presence of RSVP-Signaled LSPs on page 3768

Verifying RSVP Neighbors

Purpose Verify that each device shows the appropriate RSVP neighbors—for example, that Router R1 in Figure 141 lists both Router R3 and Router R2 as RSVP neighbors.

Action From the CLI, enter the **show rsvp neighbor** command.

```
user@r1> show rsvp neighbor
RSVP neighbor: 2 learned
Address            Idle Up/Dn LastChange HelloInt HelloTx/Rx
10.0.6.2              0  3/2       13:01        3   366/349
10.0.3.3              0  1/0       22:49        3   448/448
```

The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.

Verifying RSVP Sessions

Purpose Verify that an RSVP session has been established between all RSVP neighbors. Also, verify that the bandwidth reservation value is active.

Action From the CLI, enter the **show rsvp session detail** command.

```
user@r1> show rsvp session detail
Ingress RSVP: 1 sessions
```

Copyright © 2016, Juniper Networks, Inc.

```
10.0.9.7
  From: 10.0.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-r7, LSPpath: Primary
  Bidirectional, Upstream label in: -, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100000
  Resv style: 1 FF, Label in: -, Label out: 100000
  Time left:    -,  Since: Thu Jan 26 17:57:45 2002
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 17 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 10.0.4.13 (ge-0/0/1.0) 1467 pkts
  RESV rcvfrom: 10.0.4.13  (ge-0/0/1.0) 1467 pkts
  Record route: <self> 10.0.4.13  10.0.2.1  10.0.8.10
```

The output shows the detailed information, including session IDs, bandwidth reservation, and next-hop addresses, for each established RSVP session. Verify the following information:

- Each RSVP neighbor address has an entry for each neighbor, listed by loopback address.
- The state for each LSP session is **Up**.
- For **Tspec**, the appropriate bandwidth value, **10Mbps**, appears.

Verifying the Presence of RSVP-Signaled LSPs

Purpose Verify that the routing table of the entry (ingress) router has a configured LSP to the loopback address of the other router. For example, verify that the **inet.3** routing table of the R1 entry router in Figure 141 has a configured LSP to the loopback address of Router R7.

Action From the CLI, enter the **show route table inet.3** command.

```
user@r1> show route table inet.3
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.7/32         *[RSVP/7] 00:05:29, metric 10
                    > to 10.0.4.17 via ge-0/0/0.0, label-switched-path r1-r7
```

The output shows the RSVP routes that exist in the **inet.3** routing table. Verify that an RSVP-signaled LSP is associated with the loopback address of the exit (egress) router, R7, in the MPLS network.

Related Documentation

- Understanding the RSVP Signaling Protocol on page 3761

Understanding Point-to-Multipoint LSPs

A point-to-multipoint MPLS label-switched path (LSP) is an LDP-signaled or RSVP-signaled LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the inbound (ingress) router. Packet replication takes

Copyright © 2016, Juniper Networks, Inc.

place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in Figure 142. Device PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Device PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Device P1 replicates the packet and forwards it to Routers PE2 and PE3. Device P2 sends the packet to Device PE4.

Figure 142: Point-to-Multipoint LSPs

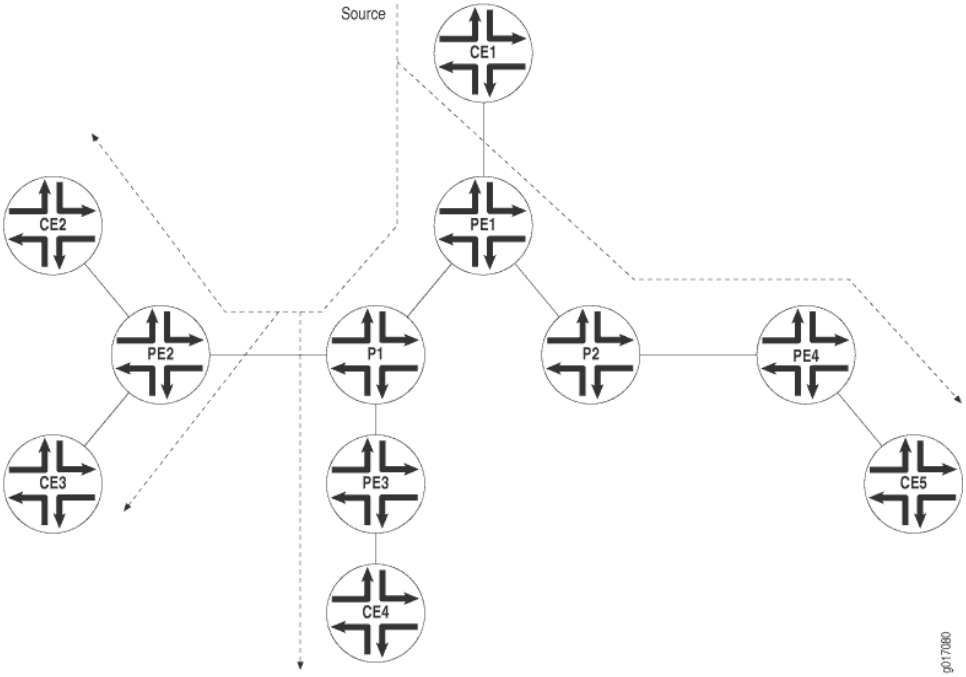

Following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP allows you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an outbound (egress) router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any primary paths fail, traffic can be quickly switched to the bypass.
- You can configure subpaths either statically or dynamically.
- You can enable graceful restart on point-to-multipoint LSPs.

Related Documentation

- MPLS Traffic Engineering and Signaling Protocols Overview on page 3755

Copyright © 2016, Juniper Networks, Inc.

- Point-to-Multipoint LSP Configuration Overview on page 3770

Point-to-Multipoint LSP Configuration Overview

To set up a point-to-multipoint LSP:

1. Configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers.
2. Specify a pathname on the primary LSP and this same path name on each branch LSP.

NOTE: By default, the branch LSPs are dynamically signaled by means of Constrained Shortest Path First (CSPF) and require no configuration. You can alternatively configure the branch LSPs as static paths.

Related Documentation

- Understanding Point-to-Multipoint LSPs on page 3768

Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP

This example shows how to configure a collection of paths to create an RSVP-signaled point-to-multipoint label-switched path (LSP).

- Requirements on page 3770
- Overview on page 3770
- Configuration on page 3771
- Verification on page 3786

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, multiple routing devices serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. On the provider edge (PE), Device PE1 is the ingress node. The branches go from PE1 to PE2, PE1 to PE3, and PE1 to PE4. Static unicast routes on the ingress node (PE1) point to the egress nodes.

This example also demonstrates static routes with a next hop that is a point-to-multipoint LSP, using the **p2mp-lsp-next-hop** statement. This is useful when implementing filter-based forwarding.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Another option is to use the **lsp-next-hop** statement to configure a regular point-to-point LSP to be the next hop. Though not shown in this example, you can optionally assign an independent preference and metric to the next hop.

Topology Diagram

Figure 143 shows the topology used in this example.

Figure 143: RSVP-Signaled Point-to-Multipoint LSP

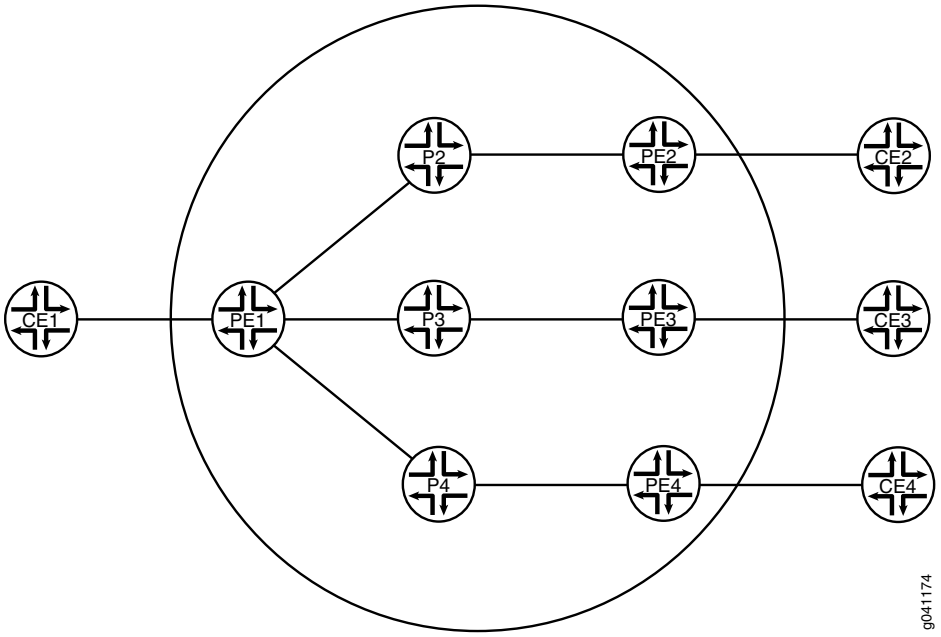

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device PE1

```
set interfaces ge-2/0/2 unit 0 description PE1-to-CE1
set interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set interfaces fe-2/0/10 unit 1 description PE1-to-P2
set interfaces fe-2/0/10 unit 1 family inet address 2.2.2.1/24
set interfaces fe-2/0/10 unit 1 family mpls
set interfaces fe-2/0/9 unit 8 description PE1-to-P3
set interfaces fe-2/0/9 unit 8 family inet address 6.6.6.1/24
set interfaces fe-2/0/9 unit 8 family mpls
set interfaces fe-2/0/8 unit 9 description PE1-to-P4
set interfaces fe-2/0/8 unit 9 family inet address 3.3.3.1/24
set interfaces fe-2/0/8 unit 9 family mpls
set interfaces lo0 unit 1 family inet address 100.10.10.10/32
set protocols rsvp interface fe-2/0/10.1
```

Copyright © 2016, Juniper Networks, Inc.

```
set protocols rsvp interface fe-2/0/9.8
set protocols rsvp interface fe-2/0/8.9
set protocols rsvp interface lo0.1
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path PE1-PE2 to 100.50.50.50
set protocols mpls label-switched-path PE1-PE2 link-protection
set protocols mpls label-switched-path PE1-PE2 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE3 to 100.70.70.70
set protocols mpls label-switched-path PE1-PE3 link-protection
set protocols mpls label-switched-path PE1-PE3 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE4 to 100.40.40.40
set protocols mpls label-switched-path PE1-PE4 link-protection
set protocols mpls label-switched-path PE1-PE4 p2mp p2mp1
set protocols mpls interface fe-2/0/10.1
set protocols mpls interface fe-2/0/9.8
set protocols mpls interface fe-2/0/8.9
set protocols mpls interface lo0.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface fe-2/0/10.1
set protocols ospf area 0.0.0.0 interface fe-2/0/9.8
set protocols ospf area 0.0.0.0 interface fe-2/0/8.9
set protocols ospf area 0.0.0.0 interface lo0.1
set routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set routing-options router-id 100.10.10.10
```

Device CE1

```
set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30
set interfaces ge-1/3/2 unit 0 description CE1-to-PE1
set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10
```

Device CE2

```
set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30
set interfaces ge-1/3/3 unit 0 description CE2-to-PE2
set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10
```

Device CE3

```
set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30
set interfaces ge-2/0/1 unit 0 description CE3-to-PE3
set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10
```

Device CE4

```
set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30
set interfaces ge-3/1/3 unit 0 description CE4-to-PE4
set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9
```

Configuring the Ingress Label-Switched Router (LSR) (Device PE1)

Step-by-Step Procedure

To configure Device PE1:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit interfaces]
user@PE1# set ge-2/0/2 unit 0 description PE1-to-CE1
user@PE1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
```


Copyright © 2016, Juniper Networks, Inc.

```
user@PE1# set fe-2/0/10 unit 1 description PE1-to-P2
user@PE1# set fe-2/0/10 unit 1 family inet address 2.2.2.1/24
user@PE1# set fe-2/0/10 unit 1 family mpls
user@PE1# set fe-2/0/9 unit 8 description PE1-to-P3
user@PE1# set fe-2/0/9 unit 8 family inet address 6.6.6.1/24
user@PE1# set fe-2/0/9 unit 8 family mpls
user@PE1# set fe-2/0/8 unit 9 description PE1-to-P4
user@PE1# set fe-2/0/8 unit 9 family inet address 3.3.3.1/24
user@PE1# set fe-2/0/8 unit 9 family mpls
user@PE1# set lo0 unit 1 family inet address 100.10.10.10/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit protocols]
user@PE1# set rsvp interface fe-2/0/10.1
user@PE1# set rsvp interface fe-2/0/9.8
user@PE1# set rsvp interface fe-2/0/8.9
user@PE1# set rsvp interface lo0.1
user@PE1# set mpls interface fe-2/0/10.1
user@PE1# set mpls interface fe-2/0/9.8
user@PE1# set mpls interface fe-2/0/8.9
user@PE1# set mpls interface lo0.1
user@PE1# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/10.1
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/9.8
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/8.9
user@PE1# set ospf area 0.0.0.0 interface lo0.1
```

3. Configure the MPLS point-to-multipoint LSPs.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 to 100.50.50.50
user@PE1# set mpls label-switched-path PE1-PE2 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE3 to 100.70.70.70
user@PE1# set mpls label-switched-path PE1-PE3 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE4 to 100.40.40.40
user@PE1# set mpls label-switched-path PE1-PE4 p2mp p2mp1
```

4. (Optional) Enable link protection on the LSPs.

   Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 link-protection
user@PE1# set mpls label-switched-path PE1-PE3 link-protection
user@PE1# set mpls label-switched-path PE1-PE4 link-protection
```

5. Enable MPLS to perform traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set mpls traffic-engineering bgp-igp
```

   This causes the ingress routes to be installed in the **inet.0** routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

Copyright © 2016, Juniper Networks, Inc.

6. Enable traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

7. Configure the router ID.

```
[edit routing-options]
user@PE1# set router-id 100.10.10.10
```

8. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```
[edit routing-options]
user@PE1# set static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
```

9. If you are done configuring the device, commit the configuration.

```
[edit]
user@PE1# commit
```

Configuring the Transit and Egress LSRs (Devices P2, P3, P4, PE2, PE3, and PE4)

Step-by-Step Procedure

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@P2# set interfaces fe-2/0/10 unit 2 description P2-to-PE1
user@P2# set interfaces fe-2/0/10 unit 2 family inet address 2.2.2.2/24
user@P2# set interfaces fe-2/0/10 unit 2 family mpls
user@P2# set interfaces fe-2/0/9 unit 10 description P2-to-PE2
user@P2# set interfaces fe-2/0/9 unit 10 family inet address 5.5.5.1/24
user@P2# set interfaces fe-2/0/9 unit 10 family mpls
user@P2# set interfaces lo0 unit 2 family inet address 100.20.20.20/32

user@PE2# set interfaces ge-2/0/3 unit 0 description PE2-to-CE2
user@PE2# set interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@PE2# set interfaces fe-2/0/10 unit 5 description PE2-to-P2
user@PE2# set interfaces fe-2/0/10 unit 5 family inet address 5.5.5.2/24
user@PE2# set interfaces fe-2/0/10 unit 5 family mpls
user@PE2# set interfaces lo0 unit 5 family inet address 100.50.50.50/32

user@P3# set interfaces fe-2/0/10 unit 6 description P3-to-PE1
user@P3# set interfaces fe-2/0/10 unit 6 family inet address 6.6.6.2/24
user@P3# set interfaces fe-2/0/10 unit 6 family mpls
user@P3# set interfaces fe-2/0/9 unit 11 description P3-to-PE3
user@P3# set interfaces fe-2/0/9 unit 11 family inet address 7.7.7.1/24
user@P3# set interfaces fe-2/0/9 unit 11 family mpls
user@P3# set interfaces lo0 unit 6 family inet address 100.60.60.60/32
```


Copyright © 2016, Juniper Networks, Inc.

```
user@PE3# set interfaces ge-2/0/1 unit 0 description PE3-to-CE3
user@PE3# set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@PE3# set interfaces fe-2/0/10 unit 7 description PE3-to-P3
user@PE3# set interfaces fe-2/0/10 unit 7 family inet address 7.7.7.2/24
user@PE3# set interfaces fe-2/0/10 unit 7 family mpls
user@PE3# set interfaces lo0 unit 7 family inet address 100.70.70.70/32

user@P4# set interfaces fe-2/0/10 unit 3 description P4-to-PE1
user@P4# set interfaces fe-2/0/10 unit 3 family inet address 3.3.3.2/24
user@P4# set interfaces fe-2/0/10 unit 3 family mpls
user@P4# set interfaces fe-2/0/9 unit 12 description P4-to-PE4
user@P4# set interfaces fe-2/0/9 unit 12 family inet address 4.4.4.1/24
user@P4# set interfaces fe-2/0/9 unit 12 family mpls
user@P4# set interfaces lo0 unit 3 family inet address 100.30.30.30/32

user@PE4# set interfaces ge-2/0/0 unit 0 description PE4-to-CE4
user@PE4# set interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@PE4# set interfaces fe-2/0/10 unit 4 description PE4-to-P4
user@PE4# set interfaces fe-2/0/10 unit 4 family inet address 4.4.4.2/24
user@PE4# set interfaces fe-2/0/10 unit 4 family mpls
user@PE4# set interfaces lo0 unit 4 family inet address 100.40.40.40/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@P2# set protocols rsvp interface fe-2/0/10.2
user@P2# set protocols rsvp interface fe-2/0/9.10
user@P2# set protocols rsvp interface lo0.2
user@P2# set protocols mpls interface fe-2/0/10.2
user@P2# set protocols mpls interface fe-2/0/9.10
user@P2# set protocols mpls interface lo0.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/9.10
user@P2# set protocols ospf area 0.0.0.0 interface lo0.2

user@PE2# set protocols rsvp interface fe-2/0/10.5
user@PE2# set protocols rsvp interface lo0.5
user@PE2# set protocols mpls interface fe-2/0/10.5
user@PE2# set protocols mpls interface lo0.5
user@PE2# set protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@PE2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.5
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.5

user@P3# set protocols rsvp interface fe-2/0/10.6
user@P3# set protocols rsvp interface fe-2/0/9.11
user@P3# set protocols rsvp interface lo0.6
user@P3# set protocols mpls interface fe-2/0/10.6
user@P3# set protocols mpls interface fe-2/0/9.11
user@P3# set protocols mpls interface lo0.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/9.11
user@P3# set protocols ospf area 0.0.0.0 interface lo0.6

user@PE3# set protocols rsvp interface fe-2/0/10.7
```

Copyright © 2016, Juniper Networks, Inc.

```
user@PE3# set protocols rsvp interface lo0.7
user@PE3# set protocols mpls interface fe-2/0/10.7
user@PE3# set protocols mpls interface lo0.7
user@PE3# set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@PE3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.7
user@PE3# set protocols ospf area 0.0.0.0 interface lo0.7

user@P4# set protocols rsvp interface fe-2/0/10.3
user@P4# set protocols rsvp interface fe-2/0/9.12
user@P4# set protocols rsvp interface lo0.3
user@P4# set protocols mpls interface fe-2/0/10.3
user@P4# set protocols mpls interface fe-2/0/9.12
user@P4# set protocols mpls interface lo0.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/9.12
user@P4# set protocols ospf area 0.0.0.0 interface lo0.3

user@PE4# set protocols rsvp interface fe-2/0/10.4
user@PE4# set protocols rsvp interface lo0.4
user@PE4# set protocols mpls interface fe-2/0/10.4
user@PE4# set protocols mpls interface lo0.4
user@PE4# set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@PE4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.4
user@PE4# set protocols ospf area 0.0.0.0 interface lo0.4
```

3. Enable traffic engineering for OSPF.

```
[edit]
user@P2# set protocols ospf traffic-engineering
user@P2# set protocols ospf traffic-engineering
user@P3# set protocols ospf traffic-engineering
user@PE2# set protocols ospf traffic-engineering
user@PE3# set protocols ospf traffic-engineering
user@PE4# set protocols ospf traffic-engineering
```

   This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```
[edit]
user@P2# set routing-options router-id 100.20.20.20
user@P3# set routing-options router-id 100.60.60.60
user@P4# set routing-options router-id 100.30.30.30
user@PE2# set routing-options router-id 100.50.50.50
user@PE3# set routing-options router-id 100.70.70.70
user@PE4# set routing-options router-id 100.40.40.40
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Copyright © 2016, Juniper Networks, Inc.

Device PE1

```
user@PE1# show interfaces
ge-2/0/2 {
  unit 0 {
    description R1-to-CE1;
    family inet {
      address 10.0.244.10/30;
    }
  }
}
fe-2/0/10 {
  unit 1 {
    description PE1-to-P2;
    family inet {
      address 2.2.2.1/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 8 {
    description PE1-to-P2;
    family inet {
      address 6.6.6.1/24;
    }
    family mpls;
  }
}
fe-2/0/8 {
  unit 9 {
    description PE1-to-P3;
    family inet {
      address 3.3.3.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 1 {
    family inet {
      address 100.10.10.10/32;
    }
  }
}

user@PE1# show protocols
rsvp {
  interface fe-2/0/10.1;
  interface fe-2/0/9.8;
  interface fe-2/0/8.9;
  interface lo0.1;
}
mpls {
  traffic-engineering bgp-igp;
  label-switched-path PE1-to-PE2 {
    to 100.50.50.50;
```

Copyright © 2016, Juniper Networks, Inc.

```
        link-protection;
        p2mp p2mp1;
      }
      label-switched-path PE1-to-PE3 {
        to 100.70.70.70;
        link-protection;
        p2mp p2mp1;
      }
      label-switched-path PE1-to-PE4 {
        to 100.40.40.40;
        link-protection;
        p2mp p2mp1;
      }
      interface fe-2/0/10.1;
      interface fe-2/0/9.8;
      interface fe-2/0/8.9;
      interface lo0.1;
    }
    ospf {
      traffic-engineering;
      area 0.0.0.0 {
        interface ge-2/0/2.0;
        interface fe-2/0/10.1;
        interface fe-2/0/9.8;
        interface fe-2/0/8.9;
        interface lo0.1;
      }
    }
```

```
user@PE1# show routing-options
static {
  route 5.5.5.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
  route 7.7.7.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
  route 4.4.4.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
}
router-id 100.10.10.10;
```

Device P2

```
user@P2# show interfaces
fe-2/0/10 {
  unit 2 {
    description P2-to-PE1;
    family inet {
      address 2.2.2.2/24;
    }
    family mpls;
  }
fe-2/0/9 {
  unit 10 {
    description P2-to-PE2;
    family inet {
      address 5.5.5.1/24;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    family mpls;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 100.20.20.20/32;
    }
  }
}
```

```
user@P2# show protocols
rsvp {
  interface fe-2/0/10.2;
  interface fe-2/0/9.10;
  interface lo0.2;
}
mpls {
  interface fe-2/0/10.2;
  interface fe-2/0/9.10;
  interface lo0.2;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.2;
    interface fe-2/0/9.10;
    interface lo0.2;
  }
}
```

```
user@P2# show routing-options
router-id 100.20.20.20;
```

Device P3

```
user@P3# show interfaces
fe-2/0/10 {
  unit 6 {
    description P3-to-PE1;
    family inet {
      address 6.6.6.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 11 {
    description P3-to-PE3;
    family inet {
      address 7.7.7.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 6 {
```

Copyright © 2016, Juniper Networks, Inc.

```
      family inet {
        address 100.60.60.60/32;
      }
    }
  }
```

```
user@P3# show protocols
rsvp {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
mpls {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.6;
    interface fe-2/0/9.11;
    interface lo0.6;
  }
}
```

```
user@P2# show routing-options
router-id 100.60.60.60;
```

Device P4

```
user@P4# show interfaces
fe-2/0/10 {
  unit 3 {
    description P4-to-PE1;
    family inet {
      address 3.3.3.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 12 {
    description P4-to-PE4;
    family inet {
      address 4.4.4.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 3 {
    family inet {
      address 100.30.30.30/32;
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

```
user@P4# show protocols
rsvp {
    interface fe-2/0/10.3;
    interface fe-2/0/9.12;
    interface lo0.3;
}
mpls {
    interface fe-2/0/10.3;
    interface fe-2/0/9.12;
    interface lo0.3;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-2/0/10.3;
        interface fe-2/0/9.12;
        interface lo0.3;
    }
}
```

```
user@P3# show routing-options
router-id 100.30.30.30;
```

Device PE2

```
user@PE2# show interfaces
    ge-2/0/3 {
        unit 0 {
            description PE2-to-CE2;
            family inet {
                address 10.0.224.10/30;
            }
        }
    }
    fe-2/0/10 {
        unit 5 {
            description PE2-to-P2;
            family inet {
                address 5.5.5.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 5 {
            family inet {
                address 100.50.50.50/32;
            }
        }
    }
}
```

```
user@PE2# show protocols
rsvp {
    interface fe-2/0/10.5;
    interface lo0.5;
}
mpls {
```

```
    interface fe-2/0/10.5;
    interface lo0.5;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/3.0;
        interface fe-2/0/10.5;
        interface lo0.5;
    }
}
```

```
user@PE2# show routing-options
router-id 100.50.50.50;
```

Device PE3

```
user@PE3# show interfaces
    ge-2/0/1 {
        unit 0 {
            description PE3-to-CE3;
            family inet {
                address 10.0.134.10/30;
            }
        }
    }
    fe-2/0/10 {
        unit 7 {
            description PE3-to-P3;
            family inet {
                address 7.7.7.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 7 {
            family inet {
                address 100.70.70.70/32;
            }
        }
    }
}
```

```
user@PE3# show protocols
rsvp {
    interface fe-2/0/10.7;
    interface lo0.7;
}
mpls {
    interface fe-2/0/10.7;
    interface lo0.7;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/1.0;
        interface fe-2/0/10.7;
```


Copyright © 2016, Juniper Networks, Inc.

```
    interface lo0.7;
  }
}

user@PE3# show routing-options
router-id 100.70.70.70;
```

Device PE4

```
user@PE4# show interfaces
  ge-2/0/0 {
    unit 0 {
      description PE4-to-CE4;
      family inet {
        address 10.0.104.9/30;
      }
    }
  }
  fe-2/0/10 {
    unit 4 {
      description PE4-to-P4;
      family inet {
        address 4.4.4.2/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 4 {
      family inet {
        address 100.40.40.40/32;
      }
    }
  }
}

user@PE4# show protocols
rsvp {
  interface fe-2/0/10.4;
  interface lo0.4;
}
mpls {
  interface fe-2/0/10.4;
  interface lo0.4;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/0.0;
    interface fe-2/0/10.4;
    interface lo0.4;
  }
}

user@PE4# show routing-options
router-id 100.40.40.40;
```

Copyright © 2016, Juniper Networks, Inc.

Configuring Device CE1

Step-by-Step Procedure

To configure Device CE1:

1. Configure an interface to Device PE1.

```
[edit interfaces]
user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
user@CE1# set ge-1/3/2 unit 0 description CE1-to-PE1
```

2. Configure static routes from Device CE1 to the three other customer networks, with Device PE1 as the next hop.

```
[edit routing-options]
user@CE1# set static route 10.0.104.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.134.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.224.8/30 next-hop 10.0.244.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE1# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE1# show interfaces
ge-1/3/2 {
   unit 0 {
      family inet {
         address 10.0.244.9/30;
         description CE1-to-PE1;
      }
   }
}

user@CE1# show routing-options
static {
   route 10.0.104.8/30 next-hop 10.0.244.10;
   route 10.0.134.8/30 next-hop 10.0.244.10;
   route 10.0.224.8/30 next-hop 10.0.244.10;
}
```

Configuring Device CE2

Step-by-Step Procedure

To configure Device CE2:

1. Configure an interface to Device PE2.

```
[edit interfaces]
user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
user@CE2# set ge-1/3/3 unit 0 description CE2-to-PE2
```

2. Configure a static route from Device CE2 to CE1, with Device PE2 as the next hop.

```
[edit routing-options]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE2# show interfaces
ge-1/3/3 {
  unit 0 {
    family inet {
      address 10.0.224.9/30;
      description CE2-to-PE2;
    }
  }
}
```

```
user@CE2# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.224.10;
}
```

Configuring Device CE3

Step-by-Step Procedure To configure Device CE3:

1. Configure an interface to Device PE3.

```
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-PE3
```

2. Configure a static route from Device CE3 to CE1, with Device PE3 as the next hop.

```
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE3# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE3# show interfaces
ge-2/0/1 {
  unit 0 {
    family inet {
      address 10.0.134.9/30;
      description CE3-to-PE3;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
user@CE3# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.134.10;
}
```

Configuring Device CE4

Step-by-Step Procedure

To configure Device CE4:

1. Configure an interface to Device PE4.

```
[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-PE4
```

2. Configure a static route from Device CE4 to CE1, with Device PE4 as the next hop.

```
[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE4# commit
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE4# show interfaces
ge-3/1/3 {
  unit 0 {
    family inet {
      address 10.0.104.10/30;
      description CE4-to-PE4;
    }
  }
}
user@CE4# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.104.9;
}
```

Verification

Confirm that the configuration is working properly.

- Verifying Connectivity on page 3787
- Verifying the State of the Point-to-Multipoint LSP on page 3787
- Checking the Forwarding Table on page 3788

Copyright © 2016, Juniper Networks, Inc.

Verifying Connectivity

Purpose Make sure that the devices can ping each other.

Action Run the **ping** command from CE1 to the interface on CE2 connecting to PE2.

```
user@CE1> ping 10.0.224.9
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to PE3.

```
user@CE1> ping 10.0.134.9
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to PE4.

```
user@CE1> ping 10.0.104.10
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

Verifying the State of the Point-to-Multipoint LSP

Purpose Make sure that the ingress, transit, and egress LSRs are in the Up state.

Copyright © 2016, Juniper Networks, Inc.

Action Run the **show mpls lsp p2mp** command on all of the LSRs. Only the ingress LSR is shown here.

```
user@PE1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To              From            State Rt P     ActivePath       LSPname
100.40.40.40    100.10.10.10    Up     0 *                      PE1-PE4
100.70.70.70    100.10.10.10    Up     0 *                      PE1-PE3
100.50.50.50    100.10.10.10    Up     0 *                      PE1-PE2
Total 3 displayed, Up 3, Down 0
...
```

Checking the Forwarding Table

Purpose Make sure that the routes are set up as expected by running the **show route forwarding-table** command. Only the routes to the remote customer networks are shown here.

Action

```
user@PE1> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
...
10.0.104.8/30      user     0 3.3.3.2            ucst  1006     6 fe-2/0/8.9
10.0.134.8/30      user     0 6.6.6.2            ucst  1010     6 fe-2/0/9.8
10.0.224.8/30      user     0 2.2.2.2            ucst  1008     6 fe-2/0/10.1
...
```

Related Documentation

- Example: Configuring RSVP-Signaled LSPs on page 3765

Copyright © 2016, Juniper Networks, Inc.

PART 55

Configuring MPLS VPNs

- Introduction to MPLS VPNs on page 3791
- Configuring MPLS Layer 2 VPNs on page 3799
- Configuring MPLS Layer 2 Circuit VPNs on page 3805
- Configuring MPLS Layer 3 VPNs on page 3809

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 159

Introduction to MPLS VPNs

- MPLS VPN Overview on page 3791
- Configuring a BGP Session for MPLS VPNs (CLI Procedure) on page 3794
- Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure) on page 3795
- Configuring Routing Options for MPLS VPNs (CLI Procedure) on page 3795
- Configuring a Routing Instance for MPLS VPNs (CLI Procedure) on page 3796

MPLS VPN Overview

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. Instead of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. VPNs are a cost-effective alternative to expensive dedicated lines. The type of VPN is determined by the connections it uses and whether the customer network or the provider network performs the virtual tunneling.

You can configure a router running Junos OS to participate in several types of VPNs. This topic discusses MPLS VPNs.

This topic contains the following sections:

- MPLS VPN Topology on page 3791
- MPLS VPN Routing on page 3793
- VRF Instances on page 3793
- Route Distinguishers on page 3793

MPLS VPN Topology

There are many ways to set up an MPLS VPN and direct traffic through it. Figure 144 shows a typical MPLS VPN topology.

Copyright © 2016, Juniper Networks, Inc.

Figure 144: Typical VPN Topology

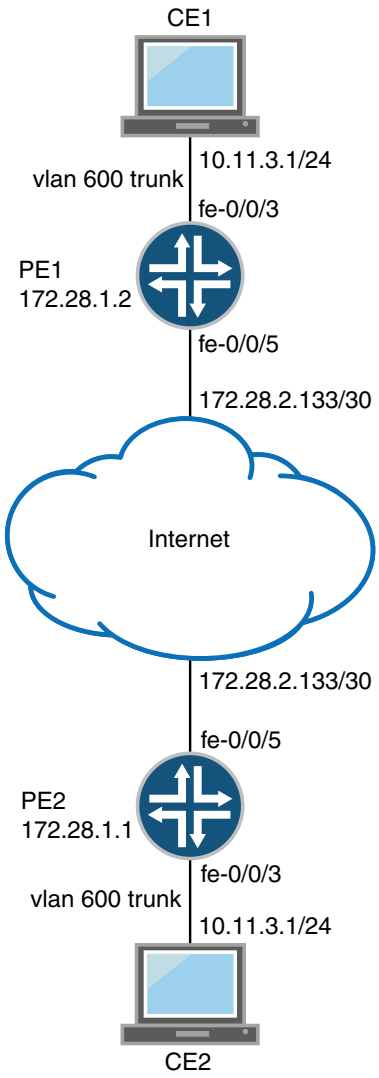

There are three primary types of MPLS VPNs: Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. All types of MPLS VPNs share certain components:

- The provider edge (PE) routers in the provider's network connect to the customer edge (CE) routers located at customer sites. PE routers support VPN and MPLS label functionality. Within a single VPN, pairs of PE routers are connected through a virtual tunnel, typically a label-switched path (LSP).
- Provider routers within the core of the provider's network are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support LSP functionality as part of the tunnel support, but do not support VPN functionality.
- CE routers are the routers or switches located at the customer site that connect to the provider's network. CE routers are typically IP routers, but they can also be Asynchronous Transfer Mode (ATM), Frame Relay, or Ethernet switches.

Copyright © 2016, Juniper Networks, Inc.

All VPN functions are performed by the PE routers. Neither CE routers nor provider routers are required to perform any VPN functions.

MPLS VPN Routing

VPNs tunnel traffic as follows from one customer site to another customer site, using a public network as a transit network, when certain requirements are met:

1. Traffic is forwarded by standard IP forwarding from the CE routers to the PE routers.
2. The PE routers establish an LSP through the provider network.
3. The inbound PE router receives traffic, and it performs a route lookup. The lookup yields an LSP next hop, and the traffic is forwarded along the LSP.
4. The traffic reaches the outbound PE router, and the PE router pops the MPLS label and forwards the traffic with standard IP routing.

VRF Instances

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of MPLS VPNs, each VPN has a VPN routing and forwarding (VRF) instance.

A VRF instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation.

A separate VRF table is created for each VPN that has a connection to a CE router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF instance, and with routes received from other PE routers in the same VPN.

Route Distinguishers

Because a typical transit network is configured to handle more than one VPN, the provider routers are likely to have multiple VRF instances configured. As a result, depending on the origin of the traffic and any filtering rules applied to the traffic, the BGP routing tables can contain multiple routes for a particular destination address. Because BGP requires that exactly one BGP route per destination be imported into the forwarding table, BGP must have a way to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs.

A route distinguisher is a locally unique number that identifies all route information for a particular VPN. Unique numeric identifiers allow BGP to distinguish between routes that are otherwise identical.

Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

Copyright © 2016, Juniper Networks, Inc.

- ***as-number:number***, where ***as-number*** is an autonomous system (AS) number (a 2–byte value) in the range 1 through 65,535, and ***number*** is any 4–byte value. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.
- ***ip-address:number***, where ***ip-address*** is an IP address (a 4–byte value) and ***number*** is any 2–byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a public IP address in your assigned prefix range.

The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VRF table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes.

The PE router then exports the route in IBGP sessions to the other provider routers. Route export is governed by any routing policy that has been applied to the particular VRF table. To propagate the routes through the provider network, the PE router must also convert the route to VPN format, which includes the route distinguisher.

When the outbound PE router receives the route, it strips off the route distinguisher and advertises the route to the connected CE router, typically through standard BGP IPv4 route advertisements.

Related Documentation

- Understanding MPLS Layer 2 VPNs on page 3799
- Understanding MPLS Layer 3 VPNs on page 3809
- Understanding MPLS Layer 2 Circuits on page 3805

Configuring a BGP Session for MPLS VPNs (CLI Procedure)

NOTE: This section is valid for Layer 2 VPNs and Layer 3 VPNs, but not Layer 2 circuits.

To configure an IBGP session, perform the following steps on each PE router:

1. Configure BGP.

   [edit]
   user@host# **edit protocols bgp group *group-name***

2. Set the BGP type to internal.

   [edit protocols bgp group *group-name*]
   user@host# **set type internal**

3. Specify the loopback interface.

   [edit protocols bgp group *group-name*]
   user@host# **set local-address *loopback-interface-ip-address***

Copyright © 2016, Juniper Networks, Inc.

4. Set the Layer 2 or Layer 3 VPN family type to unicast.

   ```
   [edit protocols bgp group group-name]
   user@host# set family family-type unicast
   ```

   Replace ***family-type*** with **l2vpn** for a Layer 2 VPN or **inet–vpn** for a Layer 3 VPN.

5. Enter the loopback address of the neighboring PE router.

   ```
   [edit protocols bgp]
   user@host# set neighbor ip-address
   ```

6. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810

Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure)

To configure RSVP and OSPF:

1. Configure OSPF with traffic engineering support on the PE routers.

   ```
   [edit]
   user@host# edit protocols ospf traffic-engineering shortcuts
   ```

   NOTE: You must configure the IGP at the **[edit protocols]** level, not within the routing instance at the **[edit routing-instances]** level.

2. Enable RSVP on interfaces that participate in the LSP. For PE routers, enable interfaces on the source and destination points. For provider routers, enable interfaces that connect the LSP between the PE routers.

   ```
   [edit]
   user@host# edit protocols rsvp interface interface-name
   ```

3. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810
- MPLS Layer 2 Circuit Configuration Overview on page 3806

Configuring Routing Options for MPLS VPNs (CLI Procedure)

To configure routing options for a VPN:

Copyright © 2016, Juniper Networks, Inc.

1. Configure the AS number.

   ```
   [edit]
   user@host# set routing-options autonomous-system as-number
   ```

2. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810
- MPLS Layer 2 Circuit Configuration Overview on page 3806

Configuring a Routing Instance for MPLS VPNs (CLI Procedure)

To configure a VPN routing instance on each PE router:

1. Create the routing instance.

   ```
   [edit]
   user@host# edit routing-instances routing-instance-name
   ```

2. Create a routing instance description. (This text appears in the output of the **show route instance detail** command.)

   ```
   [edit routing-instances routing-instance-name]
   user@host# set description "text"
   ```

3. Specify the instance type, either **l2vpn** for Layer 2 VPNs or **vrf** for Layer 3 VPNs.

   ```
   [edit routing-instances routing-instance-name]
   user@host# set instance-type instance-type
   ```

4. Specify the interface of the remote PE router.

   ```
   [edit routing-instances routing-instance-name]
   user@host# set interface interface-name
   ```

5. Specify the route distinguisher using one of the following commands:

   ```
   [edit routing-instances routing-instance-name]
   user@host# set route-distinguisheras-number:number
   user@host# set route-distinguisher ip-address:number
   ```

6. Specify the policy for the Layer 2 VRF table.

   ```
   [edit routing-instances routing-instance-name]
   user@host# set vrf-import import-policy-name vrf-export export-policy-name
   ```

7. Specify the policy for the Layer 3 VRF table.

   ```
   [edit routing-instances routing-instance-name]
   user@host# set vrf-target target:community-id
   ```

   Where *community-id* is either *as-number:number* or *ip-address:number*.

8. Commit the configuration if you are finished configuring the device.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# commit
```

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 160

Configuring MPLS Layer 2 VPNs

- Understanding MPLS Layer 2 VPNs on page 3799
- MPLS Layer 2 VPN Configuration Overview on page 3799
- Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure) on page 3801
- Configuring Interfaces for Layer 2 VPNs (CLI Procedure) on page 3802
- Verifying an MPLS Layer 2 VPN Configuration on page 3803

Understanding MPLS Layer 2 VPNs

In an MPLS Layer 2 VPN, traffic is forwarded to the provider edge (PE) router in Layer 2 format, carried by MPLS through an label-switched path (LSP) over the service provider network, and then converted back to Layer 2 format at the receiving customer edge (CE) router.

Routing occurs on the customer routers, typically on the CE router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE router receiving the traffic sends it across the network to the PE router on the outbound side. The PE routers need no information about the customer's routes or routing topology, and need only to determine the virtual tunnel through which to send the traffic.

Implementing a Layer 2 VPN on the router is similar to implementing a VPN using a Layer 2 technology such as Asynchronous Transfer Mode (ATM) or Frame Relay.

Related Documentation

- MPLS VPN Overview on page 3791
- MPLS Layer 2 VPN Configuration Overview on page 3799

MPLS Layer 2 VPN Configuration Overview

To configure MPLS Layer 2 VPN functionality on a router running Junos OS, you must enable support on the provider edge (PE) router and configure the PE router to distribute routing information to other routers in the VPN, as explained in the following steps. However, because the tunnel information is maintained at both PE routers, neither the provider core routers nor the customer edge (CE) routers need to maintain any VPN information in their configuration databases.

Copyright © 2016, Juniper Networks, Inc.

To configure an MPLS Layer 2 VPN:

1. Determine all of the routers that you want to participate in the VPN, and then complete the initial configuration of their interfaces. See *Interfaces Feature Guide for Security Devices*.
2. For all of the routers in the VPN configuration, update the interface configurations to enable participation in the Layer 2 VPN. As part of the interface configuration, you must configure the MPLS address family for each interface that uses LDP or RSVP. See "Configuring Interfaces for Layer 2 VPNs (CLI Procedure)" on page 3802.
3. For all of the routers in the VPN configuration, configure the appropriate protocols.
   a. MPLS—For PE routers and provider routers, use MPLS to advertise the Layer 2 VPN interfaces that communicate with other PE routers and provider routers.
   b. BGP and internal BGP (IBGP)—For PE routers, configure a BGP session to enable the routers to exchange information about routes originating and terminating in the VPN. (The PE routers use this information to determine which labels to use for traffic destined to the remote sites. The IBGP session for the VPN runs through the loopback address.) In addition, CE routers require a BGP connection to the PE routers. See "Configuring a BGP Session for MPLS VPNs (CLI Procedure)" on page 3794.
   c. IGP and a signaling protocol—For PE routers, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but you cannot use them for interfaces between PE routers and CE routers.

      In addition, configure an IGP such as OSPF or static routes for PE routers to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

      See "Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)" on page 3808 and "Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure)" on page 3795.
4. For all of the routers in the VPN configuration, configure routing options. The only required routing option for VPNs is the AS number. You must specify it on each router involved in the VPN. See "Configuring Routing Options for MPLS VPNs (CLI Procedure)" on page 3795.
5. For each PE router in the VPN configuration, configure a routing instance for each VPN. The routing instance should have the same name on each PE router. Each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP distinguish between potentially identical network layer reachable information [NLRI] messages received from different

 Copyright © 2016, Juniper Networks, Inc.

VPNs.) See "Configuring a Routing Instance for MPLS VPNs (CLI Procedure)" on page 3796.

6. For each PE router in the VPN configuration, configure a VPN routing policy if you are not using a route target. Within the policy, describe which packets are sent and received across the VPN and specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. If the routing instance uses a policy for accepting and rejecting packets instead of a route target, you must specify the import and export routing policies and the community on each PE router. See "Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure)" on page 3801.

Related Documentation

- Verifying an MPLS Layer 2 VPN Configuration on page 3803

Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure)

These instructions show how to configure a Layer 2 VPN routing policy on the PE routers in the VPN.

After configuring an import routing policy for a Layer 2 VPN, configure an export routing policy for the Layer 2 VPN. Configure this export policy on the PE routers in the VPN. The export routing policy defines how routes are exported from the PE router routing table. An export policy is applied to routes sent to other PE routers in the VPN. The export policy must also evaluate all routes received over the routing protocol session with the CE router. The export policy must also contain a second term for rejecting all other routes.

To configure a Layer 2 VPN routing policy on a PE router:

1. Configure the import routing policy.

   ```
   [edit]
   user@host# edit policy-options policy-statement import-policy-name
   ```

2. Define the import policy's term for accepting packets.

   ```
   [edit edit policy-options policy-statement import-policy-name]
   user@host# set term term-name-accept from protocol bgp community
     community-name
   user@host# set term term-name-accept then accept
   ```

3. Define the import policy's term for rejecting packets.

   ```
   [edit edit policy-options policy-statement import-policy-name]
   user@host# set term term-name-reject then reject
   ```

4. Configure the export routing policy.

   ```
   [edit]
   user@host# edit policy-options policy-statement export-policy-name
   ```

5. Define the export policy's term for accepting packets.

   ```
   [edit policy-options policy-statement export-policy-name]
   user@host# set term term-name-accept from community add community-name
   user@host# set termterm-name-accept then accept
   ```

Copyright © 2016, Juniper Networks, Inc.

6. Define the export policy's term for rejecting packets.

   [edit policy-options policy-statement *export-policy-name*]
   user@host# set term *term-name-reject* from community add *community-name*
   user@host# set term *term-name-reject* then reject

7. Define the export policy's community using one of the following commands.

   [edit policy-options policy-statement *export-policy-name*]
   user@host# community *community-name* target: *as-number*
   user@host# community *community-name* target: *ip-address*:*number*

8. Commit the configuration if you are finished configuring the device.

   [edit]
   user@host# commit

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810

Configuring Interfaces for Layer 2 VPNs (CLI Procedure)

Configuring the router interfaces that participate in the VPN is similar to configuring them for other uses, with a few requirements for the VPN. Perform the following tasks for each interface involved in the VPN, except Layer 3 loopback interfaces, which do not require other configuration.

To configure an interface for an MPLS VPN:

1. Configure IPv4 on all of the routers' interfaces.

   - For all interfaces except loopback interfaces and Layer 2 VPN interfaces facing a CE router:

     [edit]
     user@host# edit interfaces *interface-name* unit *logical_interface* family inet address *ipv4_address*

   - For a loopback address on a Layer 2 configuration:

     [edit]
     user@host# edit interfaces lo0 unit *logical_interface* family inet address *ipv4_address* primary

   - For a Layer 2 VPN interface facing a CE router:

     [edit]
     user@host# set interfaces *interface-name* vlan-tagging encapsulation vlan-ccc unit *logical_interface* encapsulation vlan-ccc vlan-id *id-number*

2. Configure the MPLS address family on the PE router or provider router interfaces that communicate with other PE routers or provider routers (and not loopback addresses).

   [edit interfaces *interface*]
   user@host# set unit *logical_interface* family mpls

Copyright © 2016, Juniper Networks, Inc.

3. Configure encapsulation for the interfaces on the PE routers that communicate with the CE routers in Layer 2 VPNs and Layer 2 circuits. If multiple logical units are configured, the encapsulation type is needed at the interface level only. It is always required at the unit level.

   [edit interfaces *interface*]
   user@host# **set encapsulation *encapsulation_type***
   user@host# **set unit *logical_interface* encapsulation *encapsulation_type***

4. Enable protocol mpls on CE facing interface.

   [edit interfaces *interface*]
   user@host# **set protocols mpls interface *interface-name***

5. Commit the configuration if you are finished configuring the device.

   [edit]
   user@host# **commit**

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810
- MPLS Layer 2 Circuit Configuration Overview on page 3806

Verifying an MPLS Layer 2 VPN Configuration

Purpose Verify the connectivity of MPLS Layer 2 VPNs using the **ping mpls** command. This command helps to verify that a VPN has been enabled by testing the integrity of the VPN connection between the PE routers. It does not test the connection between a PE router and CE router.

Action

- To ping an interface configured for the Layer 2 VPN on the PE router, use the following command:

  ping mpls l2vpn interface *interface-name*

- To ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier to test the integrity of the Layer 2 VPN connection (specified by identifiers) between the two PE routers, use the following command:

  ping mpls l2vpn instance *l2vpn-instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number*

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 161

Configuring MPLS Layer 2 Circuit VPNs

- Understanding MPLS Layer 2 Circuits on page 3805
- MPLS Layer 2 Circuit Configuration Overview on page 3806
- Configuring an MPLS Layer 2 Circuit (CLI Procedure) on page 3807
- Verifying an MPLS Layer 2 Circuit Configuration on page 3807
- Configuring an IGP and the LDP Signaling Protocol (CLI Procedure) on page 3808

Understanding MPLS Layer 2 Circuits

An MPLS Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by means of MPLS or another tunneling technology on the service provider network. The Layer 2 circuit creates a virtual connection to direct traffic between two customer edge (CE) routers across a service provider network. The main difference between a Layer 2 VPN and a Layer 2 circuit is the method of setting up the virtual connection. As with a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router to the local CE router. All Layer 2 circuits using a particular remote PE router neighbor is identified by its IP address and is usually the endpoint destination for the label-switched path (LSP) tunnel transporting the Layer 2 circuit.

Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. Based on the virtual circuit ID and the neighbor relationship, an LDP label is bound to an LDP circuit. LDP uses the binding for sending traffic on that Layer 2 circuit to the remote CE router.

Related Documentation

- MPLS VPN Overview on page 3791
- MPLS Layer 2 Circuit Configuration Overview on page 3806

Copyright © 2016, Juniper Networks, Inc.

MPLS Layer 2 Circuit Configuration Overview

To configure an MPLS Layer 2 circuit:

1. Determine all of the routers that you want to participate in the circuit, and then complete the initial configuration of their interfaces. See the *Interfaces Feature Guide for Security Devices*.
2. For all of the routers in the circuit configuration, update the interface configurations to enable participation in the Layer 2 circuit.
   a. On the interface communicating with the other provider edge (PE) router, specify MPLS and IPv4, and include the IP address. For the loopback interface, specify **inet**, and include the IP address. For IPv4, designate the loopback interface as primary so it can receive control packets. (Because it is always operational, the loopback interface is best able to perform the control function.)
   b. On the PE router interface facing the customer edge (CE) router, specify a circuit cross-connect (CCC) encapsulation type. The type of encapsulation depends on the interface type. For example, an Ethernet interface uses **ethernet-ccc**. (The encapsulation type determines how the packet is constructed for that interface.)
   c. On the CE router interface that faces the PE router, specify **inet** (for IPv4), and include the IP address. In addition, specify a routing protocol such as Open Shortest Path First (OSPF), which specifies the area and IP address of the router interface.

   See "Configuring Interfaces for Layer 2 VPNs (CLI Procedure)" on page 3802.
3. For all of the routers in the circuit configuration, configure the appropriate protocols.
   a. MPLS—For PE routers and provider routers, use MPLS to advertise the Layer 2 circuit interfaces that communicate with other PE routers and provider routers.
   b. BGP—For PE routers, configure a BGP session.
   c. IGP and a signaling protocol—For PE routers, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but cannot use them for interfaces between PE routers and CE routers.

      In addition, configure an IGP such as OSPF or static routes on the PE routers to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

Copyright © 2016, Juniper Networks, Inc.

See "Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)" on page 3808 and "Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure)" on page 3795.

4. For all of the routers in the circuit configuration, configure routing options. The only required routing option for circuits is the autonomous system (AS) number. You must specify it on each router involved in the circuit. See "Configuring Routing Options for MPLS VPNs (CLI Procedure)" on page 3795.
5. For PE routers, configure Layer 2 circuits on the appropriate interfaces. See "Configuring an MPLS Layer 2 Circuit (CLI Procedure)" on page 3807.

Related Documentation

- Verifying an MPLS Layer 2 Circuit Configuration on page 3807

Configuring an MPLS Layer 2 Circuit (CLI Procedure)

To configure a Layer 2 circuit on a PE router:

1. Enable a Layer 2 circuit on the appropriate interface.

   ```
   [edit]
   user@host# edit protocols l2circuit neighbor interface-name interface interface-name
   ```

2. Enter the circuit ID number.

   ```
   [edit protocols l2circuit neighbor interface-name interface interface-name]
   user@host# set virtual-circuit-id id-number
   ```

   For **neighbor**, specify the local loopback address, and for **interface**, specify the interface name of the remote PE router.

3. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- MPLS Layer 2 Circuit Configuration Overview on page 3806

Verifying an MPLS Layer 2 Circuit Configuration

Purpose To verify the connectivity of MPLS Layer 2 circuits, use the **ping mpls** command. This command helps to verify that the circuit has been enabled by testing the integrity of the Layer 2 circuit between the source and destination routers.

Action

- To ping an interface configured for the Layer 2 circuit on the PE router, enter the following command:

  ping mpls l2circuit interface *interface-name*

- To ping a combination of the IPv4 prefix and the virtual circuit ID on the destination PE router, enter the following command:

  ping mpls l2circuit virtual-circuit*prefix virtual-circuit-id*

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- MPLS Layer 2 Circuit Configuration Overview on page 3806

Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)

The following instructions show how to configure LDP and OSPF on PE routers and provider routers. Within the task, you specify which interfaces to enable for LDP. Perform this step on each PE router interface and provider router interface that communicates with other PE routers and provider routers. For OSPF, you configure at least one area on at least one of the router's interfaces. (An AS can be divided into multiple areas.) These instructions use the backbone area **0.0.0.0** and show how to enable traffic engineering for Layer 2 VPN circuits.

To configure LDP and OSPF:

1. Enable the ldp protocol.

   ```
   [edit]
   user@host# edit protocols ldp
   ```

   NOTE: You must configure the IGP at the **[protocols]** level of the configuration hierarchy, not within the routing instance at the **[routing-instances]** level of the configuration hierarchy.

2. Specify which interfaces to enable for LDP.

   ```
   [edit protocols ldp]
   user@host# edit interface interface-name
   ```

3. Configure OSPF for each interface that uses LDP.

   ```
   [edit]
   user@host# edit protocols ospf area 0.0.0.0 interface interface-name
   ```

4. (Layer 2 VPN circuits only) Enable traffic engineering.

   ```
   [edit protocols ospf]
   user@host# set traffic engineering
   ```

5. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810
- MPLS Layer 2 Circuit Configuration Overview on page 3806

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 162

Configuring MPLS Layer 3 VPNs

- Understanding MPLS Layer 3 VPNs on page 3809
- MPLS Layer 3 VPN Configuration Overview on page 3810
- Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure) on page 3811
- Verifying an MPLS Layer 3 VPN Configuration on page 3812

Understanding MPLS Layer 3 VPNs

An MPLS Layer 3 VPN operates at the Layer 3 level of the OSI model, the Network layer. The VPN is composed of a set of sites that are connected over a service provider's existing public Internet backbone. The sites share common routing information and the connectivity of the sites is controlled by a collection of policies.

In an MPLS Layer 3 VPN, routing occurs on the service provider's routers. The provider routers route and forward VPN traffic at the entry and exit points of the transit network. The service provider network must learn the IP addresses of devices sending traffic across the VPN and the routes must be advertised and filtered throughout the provider network. As a result, Layer 3 VPNs require information about customer routes and a more extensive VPN routing and forwarding (VRF) policy configuration than a Layer 2 VPN. This information is used to share and filter routes that originate or terminate in the VPN.

The MPLS Layer 3 VPN requires more processing power on the provider edge (PE) routers than a Layer 2 VPN, because the Layer 3 VPN has larger routing tables for managing network traffic on the customer sites. Route advertisements originate at the customer edge (PE) routers and are shared with the inbound PE routers through standard IP routing protocols, typically BGP. Based on the source address, the PE router filters route advertisements and imports them into the appropriate VRF table.

The provider router uses OSPF and LDP to communicate with the PE routers. For OSPF, the provider router interfaces that communicate with the PE routers are specified, as well as the loopback interface. For the PE routers, the loopback interface is in passive mode, meaning it does not send OSPF packets to perform the control function.

Related Documentation

- MPLS VPN Overview on page 3791

Copyright © 2016, Juniper Networks, Inc.

MPLS Layer 3 VPN Configuration Overview

To configure MPLS Layer 3 VPN functionality on a router running Junos OS, you must enable support on the provider edge (PE) router and configure the PE router to distribute routing information to other routers in the VPN, as explained in the following steps. However, because the tunnel information is maintained at both PE routers, neither the provider core routers nor the customer edge (CE) routers need to maintain any VPN information in their configuration databases.

To configure an MPLS Layer 3 VPN:

1. Determine all of the routers that you want to participate in the VPN, and then complete the initial configuration of their interfaces. See the *Interfaces Feature Guide for Security Devices*.
2. For all of the routers in the VPN configuration, update the interface configurations to enable participation in the Layer 3 VPN. As part of the interface configuration, you must configure the MPLS address family for each interface that uses LDP or RSVP. See "Configuring Interfaces for Layer 2 VPNs (CLI Procedure)" on page 3802.
3. For all of the routers in the VPN configuration, configure the appropriate protocols.
   a. MPLS—If you are using RSVP, use MPLS to advertise the Layer 3 VPN interfaces on the PE routers and provider routers that communicate with other PE routers and provider routers.
   b. BGP, EBGP, and internal BGP (IBGP)—For PE routers, configure a BGP session to enable the routers to exchange information about routes originating and terminating in the VPN. (The PE routers use this information to determine which labels to use for traffic destined to the remote sites. The IBGP session for the VPN runs through the loopback address.) In addition, CE routers require a BGP connection to the PE routers. See "Configuring a BGP Session for MPLS VPNs (CLI Procedure)" on page 3794.
   c. IGP and a signaling protocol—For PE routers and provider, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but cannot use them for interfaces between PE routers and CE routers.

      In addition, configure an IGP such as OSPF or static routes on the PE routers in order to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

Copyright © 2016, Juniper Networks, Inc.

See "Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)" on page 3808 and "Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure)" on page 3795.

4. For all of the routers in the VPN configuration, configure routing options. The only required routing option for VPNs is the autonomous system (AS) number. You must specify it on each router involved in the VPN. See "Configuring Routing Options for MPLS VPNs (CLI Procedure)" on page 3795.
5. For each PE router in the VPN configuration, configure a routing instance for each VPN. The routing instance should have the same name on each PE router. Each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP distinguish between potentially identical network layer reachable information [NLRI] messages received from different VPNs.) See "Configuring a Routing Instance for MPLS VPNs (CLI Procedure)" on page 3796.
6. For CE routers, configure a routing policy. In addition, if you are not using a route target, configure a VPN routing policy for each PE router in the VPN configuration. Within the policy, describe which packets are sent and received across the VPN and specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. See "Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure)" on page 3811.

Related Documentation

- Verifying an MPLS Layer 3 VPN Configuration on page 3812

Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure)

To configure a Layer 3 VPN routing policy on a CE router:

1. Configure the routing policy for the loopback interface.

   ```
   [edit]
   user@host# edit policy-options policy-statement policy-name
   ```

2. Define the term for accepting packets.

   ```
   [edit policy-options policy-statement policy-name]
   user@host# set term term-name-accept from protocol direct route-filter
     local-loopback-address/netmask exact
   user@host# set termterm-name-accept then accept
   ```

3. Define the term for rejecting packets.

   ```
   [edit policy-options policy-statement policy-name]
   user@host# set termterm-name-reject then reject
   ```

4. Commit the configuration if you are finished configuring the device.

   ```
   [edit]
   user@host# commit
   ```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- MPLS Layer 3 VPN Configuration Overview on page 3810

Verifying an MPLS Layer 3 VPN Configuration

Purpose Verify the connectivity of MPLS Layer 3 VPNs using the **ping mpls** command. This command helps to verify that a VPN has been enabled by testing the integrity of the VPN connection between the source and destination routers. The destination prefix corresponds to a prefix in the Layer 3 VPN. However, ping tests only whether the prefix is present in a PE VRF table.

Action To a combination of a IPv4 destination prefix and a Layer 3 VPN name on the destination PE router, use the following command:

ping mpls l3vpn *l3vpn-name* prefix *prefix* count *count*

Related Documentation

- MPLS Layer 3 VPN Configuration Overview on page 3810

Copyright © 2016, Juniper Networks, Inc.

PART 56

Configuring CLNS VPNs

- Introduction to CLNS on page 3815
- Configuring ES-IS for CLNS on page 3819
- Configuring IS-IS for CLNS on page 3823
- Configuring Static Routes for CLNS on page 3827
- Configuring BGP for CLNS on page 3831

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 163

Introduction to CLNS

- CLNS Overview on page 3815
- CLNS Configuration Overview on page 3815

CLNS Overview

Connectionless Network Service (CLNS) is a Layer 3 protocol similar to IP version 4 (IPv4) for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network. CLNS and its related OSI protocols, Intermediate System-to-Intermediate System (IS-IS) and End System-to-Intermediate System (ES-IS), are International Organization for Standardization (ISO) standards.

You can configure devices running Junos OS as provider edge (PE) routers within a CLNS network. CLNS networks can be connected over an IP MPLS network core using Border Gateway Protocol (BGP) and MPLS Layer 3 virtual private networks (VPNs). See RFC 2547, *BGP/MPLS VPNs*.

CLNS uses network service access points (NSAPs), similar to IP addresses found in IPv4, to identify end systems (hosts) and intermediate systems (routers). ES-IS enables the hosts and routers to discover each other. IS-IS is the interior gateway protocol (IGP) that carries ISO CLNS routes through a network.

For more information about CLNS, see the ISO 8473 standards.

Related Documentation

- CLNS Configuration Overview on page 3815
- Understanding ES-IS for CLNS on page 3819
- Understanding IS-IS for CLNS on page 3823
- Understanding Static Routes for CLNS on page 3827
- Understanding BGP for CLNS VPNs on page 3831

CLNS Configuration Overview

To configure CLNS:

1. Configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.
2. If applicable, configure BGP and VPNs. See:

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring BGP for CLNS VPNs on page 3832
- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 3 VPN Configuration Overview on page 3810

3. Configure a VPN routing instance. You typically configure ES-IS, IS-IS, and CLNS static routes using a VPN routing instance. See "Example: Configuring a VPN Routing Instance for CLNS" on page 3834.
4. Configure one or more of the following protocols for CLNS (depending on your network).

   - ES-IS—If a device is a PE router within a CLNS island that contains any end systems, you must configure ES-IS on the device. If a CLNS island does not contain any end systems, you do not need to configure ES-IS on a device. See "Example: Configuring ES-IS for CLNS" on page 3820.

     NOTE: ES-IS is enabled only if either ES-IS or IS-IS is configured on the router. ES-IS must not be disabled. If ES-IS is not explicitly configured, the interface sends and receives only intermediate system hello (ISH) messages. If ES-IS is explicitly configured and disabled, the interface does not send or receive ES-IS packets. If ES-IS is explicitly configured and not disabled, the interface sends and receives ISH messages as well as ES-IS packets.

     One of the interfaces that is configured for ES-IS must be configured with an ISO address for hello messages. The ISO address family must be configured on an interface to support ES-IS on that interface.

   - IS-IS—You can configure IS-IS to exchange CLNS routes within a CLNS island. See "Example: Configuring IS-IS for CLNS" on page 3823.

     NOTE: If you have a pure CLNS island—an island that does not contain any IP devices—you must disable IPv4 and IPv6 routing. Also, to export BGP routes into IS-IS, you must configure and apply an export policy.

   - Static routes—If some devices in your network do not support IS-IS, you must configure CLNS static routes. You can use static routing with or without IS-IS. You might also consider using static routes if your network is simple. See "Example: Configuring Static Routes for CLNS" on page 3827.
   - BGP—See "Example: Configuring BGP for CLNS VPNs" on page 3832.

NOTE: Many of the configuration statements used to configure CLNS and routing protocols can be included at different hierarchy levels in the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *Junos OS Routing Protocols Library for Security Devices*
- CLNS Overview on page 3815
- Verifying a CLNS VPN Configuration on page 3835

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 164

Configuring ES-IS for CLNS

- Understanding ES-IS for CLNS on page 3819
- Example: Configuring ES-IS for CLNS on page 3820

Understanding ES-IS for CLNS

End System-to-Intermediate System (ES-IS) is a protocol that resolves Layer 3 ISO network service access points (NSAP) to Layer 2 addresses. ES-IS has an equivalent role as Address Resolution Protocol (ARP) in IP version 4 (IPv4).

ES-IS provides the basic interaction between Connectionless Network Service (CLNS) hosts (end systems) and routers (intermediate systems). ES-IS allows hosts to advertise NSAP addresses to other routers and hosts attached to the network. Those routers can then advertise the address to the rest of the network by using Intermediate System-to-Intermediate System (IS-IS). Routers use ES-IS to advertise their network entity title (NET) to hosts and routers that are attached to that network.

ES-IS routes are exported to Layer 1 IS-IS by default. You can also export ES-IS routes into Layer 2 IS-IS by configuring a routing policy. ES-IS generates and receives end system hello (ESH) hello messages when the protocol is configured on an interface. ES-IS is a resolution protocol that allows a network to be fully ISO integrated at both the network layer and the data layer.

The resolution of Layer 3 ISO NSAPs to Layer 2 subnetwork point of attachments (SNPAs) by ES-IS is equivalent to ARP within an IPv4 network. If a device is a provider edge (PE) router within a CLNS island that contains any end systems, you must configure ES-IS on the device.

For more information about ES-IS, see the ISO 9542 standard.

Related Documentation

- CLNS Overview on page 3815
- Example: Configuring ES-IS for CLNS on page 3820

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring ES-IS for CLNS

This example shows how to create a routing instance and enable ES-IS for CLNS on all interfaces.

- Requirements on page 3820
- Overview on page 3820
- Configuration on page 3820
- Verification on page 3821

Requirements

Before you begin, configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing-instance called aaaa, set the end system configuration timer for the interfaces to 180, and set a preference value to 30 for ES-IS.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa protocols esis interface all end-system-configuration-timer
  180
set routing-instances aaaa protocols esis preference 30
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure ES-IS for CLNS:

1. Configure the routing instance.

   ```
   [edit]
   user@host# edit routing-instances aaaa
   ```

2. Enable ES-IS on all interfaces.

   ```
   [edit routing-instances aaaa]
   user@host# set protocols esis interface all
   ```

3. Configure the end system configuration timer.

   ```
   [edit routing-instances aaaa]
   user@host# set protocols esis interface all end-system-configuration-timer 180
   ```

4. Configure the preference value.

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-instances aaaa]
user@host# set protocols esis preference 30
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
aaaa {
  protocols {
    esis {
      preference 30;
      interface all {
        end-system-configuration-timer 180;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Routing-Instance for CLNS on page 3821
- Verifying ES-IS for CLNS on page 3821

Verifying Routing-Instance for CLNS

Purpose Verify that the policy options are enabled for the routing instance.

Action From operational mode, enter the **show routing-instances** command.

Verifying ES-IS for CLNS

Purpose Verify that ES-IS is enabled.

Action From operational mode, enter the **show protocols** command.

Related Documentation

- CLNS Configuration Overview on page 3815
- Understanding ES-IS for CLNS on page 3819
- Verifying a CLNS VPN Configuration on page 3835

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 165

Configuring IS-IS for CLNS

- Understanding IS-IS for CLNS on page 3823
- Example: Configuring IS-IS for CLNS on page 3823

Understanding IS-IS for CLNS

Intermediate System-to-Intermediate System (IS-IS) extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

Related Documentation

- CLNS Overview on page 3815
- Example: Configuring IS-IS for CLNS on page 3823

Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable IS-IS protocol on all interfaces.

- Requirements on page 3823
- Overview on page 3823
- Configuration on page 3824
- Verification on page 3825

Requirements

Before you begin, configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing-instance called aaaa, enable IS-IS on all interfaces, and define BGP export policy name (dist-bgp), family (ISO), and protocol (BP), and apply the export policy to IS-IS.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure IS-IS for CLNS:

1. Configure the routing instance.

   ```
   [edit]
   user@host# edit routing-instances aaaa
   ```

2. Enable CLNS routing.

   ```
   [edit routing-instances aaaa]
   user@host# set protocols isis clns-routing
   ```

3. Enable IS-IS on all interfaces.

   ```
   [edit routing-instances aaaa]
   user@host# set protocols isis interface all
   ```

4. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network .

   ```
   [edit routing-instances aaaa]
   user@host# set protocols isis no-ipv4-routing no-ipv6-routing
   ```

5. Define the BGP export policy name, family, and protocol.

   ```
   [edit policy-options]
   user@host# set policy-statement dist-bgp from family iso protocol bgp
   ```

6. Define the action for the export policy.

   ```
   [edit policy-options]
   user@host# set policy-statement dist-bgp then accept
   ```

7. Apply the export policy to IS-IS.

   ```
   [edit routing-instances aaaa]
   user@host# set protocols isis export dist-bgp
   ```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show routing-instances
aaaa {
  protocols {
    isis {
      export dist-bgp;
      no-ipv4-routing;
      no-ipv6-routing;
      clns-routing;
      interface all;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Routing-Instance for CLNS on page 3825
- Verifying IS-IS for CLNS on page 3825

Verifying Routing-Instance for CLNS

Purpose Verify that the policy options are enabled for the routing instance.

Action From operational mode, enter the **show routing-instances** command.

Verifying IS-IS for CLNS

Purpose Verify that IS-IS is enabled.

Action From operational mode, enter the **show protocols** command.

Related Documentation

- CLNS Configuration Overview on page 3815
- Understanding IS-IS for CLNS on page 3823
- Verifying a CLNS VPN Configuration on page 3835

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 166

Configuring Static Routes for CLNS

- Understanding Static Routes for CLNS on page 3827
- Example: Configuring Static Routes for CLNS on page 3827

Understanding Static Routes for CLNS

The Connectionless Network Service (CLNS) is an ISO Layer 3 protocol that uses network service access point (NSAP) reachability information instead of IPv4 or IPv6 prefixes.

You can configure static routes to exchange CLNS routes within a CLNS island. A CLNS island is typically an IS-IS level 1 area that is part of a single IGP routing domain. An island can contain more than one area. CLNS islands can be connected by VPNs.

Related Documentation

- Example: Configuring Static Routes for CLNS on page 3827

Example: Configuring Static Routes for CLNS

This example shows how to configure static routes for CLNS.

- Requirements on page 3827
- Overview on page 3827
- Configuration on page 3828
- Verification on page 3829

Requirements

Before you begin, configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

In this example, you configure static routes for CLNS. In the absence of an interior gateway protocol (IGP) on a certain link, a routing device might need to be configured with static routes for CLNS prefixes to be reachable by way of that link. This might be useful, for example, at an autonomous system (AS) boundary.

Copyright © 2016, Juniper Networks, Inc.

When you configure static routes for CLNS, consider the following tasks:

- Specify the **iso.0** routing table option to configure a primary instance CLNS static route.
- Specify the **instance-name.iso.0** routing table option to configure a CLNS static route for a particular routing instance.
- Specify the **route** ***nsap-prefix*** statement to configure the destination for the CLNS static route.
- Specify the **next-hop (*interface-name* | *iso-net*)** statement to configure the next hop, specified as an ISO network entity title (NET) or interface name.
- Include the **qualified-next-hop (*interface-name* | *iso-net*)** statement to configure a secondary backup next hop, specified as an ISO network entity title or interface name.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
set routing-options rib iso.0 static iso-route
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.eee0/152
  qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference
  20
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.eee0/152
  qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

Step-by-Step Procedure

To configure static routes for CLNS:

1. Configure the routes.

```
[edit routing-options rib iso.0 static]
user@host# set iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
user@host# set iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152
  next-hop t1-0/2/2.0
user@host# set iso-route 47.0005.80ff.f800.0000.eee0/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference 20
user@host# set iso-route 47.0005.80ff.f800.0000.eee0/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by issuing the **show routing-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show routing-options
rib iso.0 {
  static {
    iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
      47.0005.80ff.f800.0000.0108.0001.1921.6800.4212;
    iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0;
    iso-route 47.0005.80ff.f800.0000.eee0/152 {
      qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 {
        preference 20;
        metric 10;
      }
    }
  }
}
```

Verification

Checking the Routing Table

Purpose Make sure that the expected routes appear in the routing table.

Action

```
user@host> show route table iso.0

iso.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152
                   *[Static/5] 00:00:25
                    > via t1-0/2/2.0
47.0005.80ff.f800.0000.eee0/84
                   *[Static/20] 00:04:01, metric 10, metric2 10
                    > to #75 0.12.0.34.0.56 via fe-0/0/1.0
47.0005.80ff.f800.0000.ffff.ffff/104
                   *[Static/5] 00:04:01, metric2 0
                    > via t1-0/2/2.0
```

Meaning The static routes appear in the routing table.

Related Documentation

- CLNS Configuration Overview on page 3815
- Understanding Static Routes for CLNS on page 3827

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 167

Configuring BGP for CLNS

- Understanding BGP for CLNS VPNs on page 3831
- Example: Configuring BGP for CLNS VPNs on page 3832
- Example: Configuring a VPN Routing Instance for CLNS on page 3834
- Verifying a CLNS VPN Configuration on page 3835

Understanding BGP for CLNS VPNs

BGP extensions allow BGP to carry Connectionless Network Service (CLNS) virtual private network (VPN) network layer reachability information (NLRI) between provider edge (PE) routers. Each CLNS route is encapsulated into a CLNS VPN NLRI and propagated between remote sites in a VPN.

CLNS is a Layer 3 protocol similar to IP version 4 (IPv4). CLNS uses network service access points (NSAPs) to address end systems. This allows for a seamless autonomous system (AS) based on International Organization for Standardization (ISO) NSAPs.

A single routing domain consisting of ISO NSAP devices are considered to be CLNS islands. CLNS islands are connected together by VPNs.

You can configure BGP to exchange ISO CLNS routes between PE routers connecting various CLNS islands in a VPN using multiprotocol BGP extensions. These extensions are the ISO VPN NLRIs.

Each CLNS network island is treated as a separate VPN routing and forwarding instance (VRF) instance on the PE router.

You can configure CLNS on the global level, group level, and neighbor level.

Related Documentation

- CLNS Overview on page 3815
- Example: Configuring BGP for CLNS VPNs on page 3832

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring BGP for CLNS VPNs

This example shows how to create a BGP group for CLNS VPNs, define the BGP peer neighbor address for the group, and define the family.

- Requirements on page 3832
- Overview on page 3832
- Configuration on page 3832
- Verification on page 3833

Requirements

Before you begin, configure the network interfaces. See the *Interfaces Feature Guide for Security Devices*.

Overview

In this example, you create the BGP group called pedge-pedge, define the BGP peer neighbor address for the group as 10.255.245.215, and define the BGP family.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols bgp group pedge-pedge neighbor 10.255.245.213
set protocols bgp family iso-vpn unicast
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure BGP for CLNS VPNs:

1. Configure the BGP group and define the BGP peer neighbor address.

   ```
   [edit protocols bgp]
   user@host# set group pedge-pedge neighbor 10.255.245.213
   ```

2. Define the family.

   ```
   [edit protocols bgp]
   user@host# set family iso-vpn unicast
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```


Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying the Neighbor Status

Purpose Display information about the BGP peer.

Action From operational mode, run the **show bgp neighbor 10.255.245.213** command. Look for **iso-vpn-unicast** in the output.

```
user@host> show bgp neighbor 10.255.245.213
Peer: 10.255.245.213+179 AS 200 Local: 10.255.245.214+3770 AS 100
Type: External State: Established Flags: <ImportEval Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
Address families configured: iso-vpn-unicast
Local Address: 10.255.245.214 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.245.213 Local ID: 10.255.245.214 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes: 3
Received prefixes: 3
Suppressed due to damping: 0
Advertised prefixes: 3
Table aaaa.iso.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes: 3
Received prefixes: 3
Suppressed due to damping: 0
Last traffic (seconds): Received 6 Sent 5 Checked 5
Input messages: Total 1736 Updates 4 Refreshes 0 Octets 33385
Output messages: Total 1738 Updates 3 Refreshes 0 Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0
```

Related Documentation

- CLNS Configuration Overview on page 3815
- Understanding BGP for CLNS VPNs on page 3831
- Verifying a CLNS VPN Configuration on page 3835

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring a VPN Routing Instance for CLNS

This example shows how to create a CLNS routing instance and set the instance type for Layer 3 VPNs.

- Requirements on page 3834
- Overview on page 3834
- Configuration on page 3834
- Verification on page 3835

Requirements

Before you begin, configure the network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

The following example shows how to create a CLNS routing instance called aaaa and set the instance type to VRF for Layer 3 VPNs. Within the example, you specify that the lo0.1 interface, e1–2/0/0.0 interface, and t1–3/0/0.0 interface all belong to the routing instance. The route distinguisher is set as 10.255.245.1:1 and the policy for the Layer 3 VRF table is set as target:11111:1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa instance-type vrf
set routing-instances aaaa interface lo0.1
set routing-instances aaaa interface ge-0/0/3
set routing-instances aaaa interface ge-0/0/2
set routing-instances aaaa route-distinguisher 10.255.245.1:1
set routing-instances aaaa vrf-target target:11111:1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a VPN routing instance:

1. Create the routing instance.

   ```
   [edit]
   user@host# edit routing-instances aaaa
   ```

2. Specify the routing instance type.

   ```
   [edit routing-instances aaaa]
   user@host# set instance-type vrf
   ```


Copyright © 2016, Juniper Networks, Inc.

3. Specify the interfaces that belong to the routing instance.

```
[edit routing-instances aaaa]
user@host# set interface lo0.1
user@host# set interface ge-0/0/3
user@host# set interface ge-0/0/2
```

4. Specify the route distinguisher.

```
[edit routing-instances aaaa]
user@host# set route-distinguisher 10.255.245.1:1
```

5. Specify the policy for the Layer 3 VRF table.

```
[edit routing-instances aaaa]
user@host# set vrf-target target:11111:1
```

6. Enable family ISO on the interfaces edit interfaces interface-name unit-id.

```
[edit routing-instances aaaa]
user@host# set family ISO
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit ]
user@host# show routing-instances
instance-type vrf;
interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface lo0.1;
route-distinguisher 10.255.245.1:1;
vrf-target target:11111:1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configured CLNS Routing Instance

Purpose Confirm that the configuration is working properly.

Verify that the CLNS routing instance is configured.

Action From operational mode, enter the **show routing-instances** command.

Related Documentation

- CLNS Configuration Overview on page 3815
- Verifying a CLNS VPN Configuration on page 3835

Verifying a CLNS VPN Configuration

Purpose Verify that the device is configured correctly for CLNS VPNs.

Copyright © 2016, Juniper Networks, Inc.

Action From configuration mode in the CLI, enter the **show** command.

```
[edit]
user@host# show
interfaces {
   e1–2/0/0.0 {
      unit 0 {
         family inet {
            address 192.168.37.51/31;
         }
         family iso;
         family mpls;
      }
   }
   t1–3/0/0.0 {
      unit 0 {
         family inet {
            address 192.168.37.24/32;
         }
         family iso;
         family mpls;
      }
   }
   lo0 {
      unit 0 {
         family inet {
            address 127.0.0.1/32;
            address 10.255.245.215/32;
         }
         family iso {
            address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4215.00;
         }
      }
      unit 1 {
         family iso {
            address 47.0005.80ff.f800.0000.0108.aaa2.1921.6800.4215.00;
         }
      }
   }
}
routing-options {
   autonomous-system 230;
}
protocols {
   bgp {
      group pedge-pedge {
         type internal;
         local-address 10.255.245.215;
         neighbor 10.255.245.212 {
            family iso-vpn {
               unicast;
            }
         }
      }
   }
}
```


Copyright © 2016, Juniper Networks, Inc.

```
policy-options {
   policy-statement dist-bgp {
      from {
         protocol bgp;
         family iso;
      }
      then accept;
   }
}
routing-instances {
   aaaa {
      instance-type vrf;
      interface lo0.1;
      interface e1–2/0/0.0;
      interface t1–3/0/0.0;
      route-distinguisher 10.255.245.1:1;
      vrf-target target:11111:1;
      routing-options {
         rib aaaa.iso.0 {
            static {
               iso-route 47.0005.80ff.f800.0000.bbbb.1022/104
                  next-hop 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00;
            }
         }
      }
      protocols {
         esis {
            interface all;
         }
         isis {
            export dist-bgp;
            no-ipv4–routing;
            no-ip64–routing;
            clns–routing;
            interface all;
         }
      }
   }
}
```

Related Documentation

- CLNS Configuration Overview on page 3815

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 57

Configuring VPLS

- Introduction to VPLS on page 3841
- Configuring Interfaces on page 3849
- Configuring Routing Instances on page 3859
- Configuring Routing and Signaling Protocols on page 3867
- Configuring Encapsulation on page 3905

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 168

Introduction to VPLS

- VPLS Overview on page 3841
- VPLS Configuration Overview on page 3846

VPLS Overview

Virtual private LAN service (VPLS) is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet LAN sites to each other across an MPLS backbone. For customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider's network.

VPLS, in its implementation and configuration, has much in common with an MPLS Layer 2 VPN. In a VPLS topology, a packet originating within a customer's network is sent first to a customer edge (CE) device (for example, a router or Ethernet switch). It is then sent to a provider edge (PE) router within the service provider's network. The packet traverses the service provider's network over an MPLS label-switched path (LSP). It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.

The difference is that for VPLS, packets can traverse the service provider's network in point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to all the PE routers participating in a VPLS routing instance. In contrast, a Layer 2 VPN forwards packets in point-to-point fashion only. The paths carrying VPLS traffic between each PE router participating in a routing instance are signaled using BGP.

NOTE: **The RSVP automatic mesh feature with multiple RSVP neighbors on a single LAN is not supported on branch SRX Series devices because RSVP runs on WAN links in a service provider network. Most of these WAN interfaces are point-to-point and are rarely seen in LAN networks.**

This topic contains the following sections:

- Sample VPLS Topology on page 3842
- VPLS on PE Routers on page 3843
- Using an Ethernet Switch as the VPLS CE Device on page 3845
- VPLS Exceptions on SRX Series Devices on page 3845

Copyright © 2016, Juniper Networks, Inc.

Sample VPLS Topology

Figure 145 shows a basic VPLS topology.

Figure 145: Basic VPLS Topology

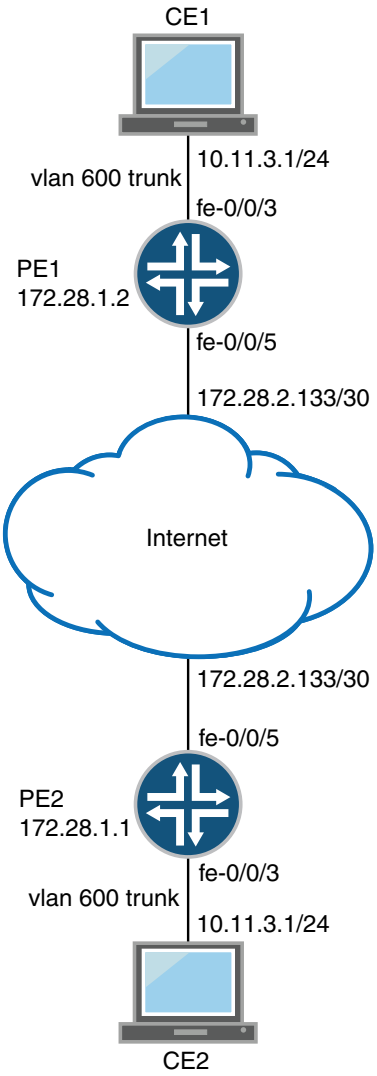

In this sample, the PE routers use the same autonomous system (AS). Within the AS, routing information is communicated through an interior gateway protocol (IGP). Outside the AS, routing information is shared with other ASs through BGP. The PE routers must use the same signaling protocols to communicate.

Copyright © 2016, Juniper Networks, Inc.

VPLS on PE Routers

Within a VPLS configuration, a device running Junos OS can act as a PE router. Junos OS passes the VPLS traffic through the following ports and PIMs on the Juniper Networks device to CE routers in the VPLS network:

- Built-in Ethernet ports on front panel
- Gigabit Ethernet uPIMs
- Gigabit Ethernet ePIMs
- Fast Ethernet PIMs
- Fast Ethernet ePIMs

NOTE: **Ports on uPIMs and ePIMs must be in routing mode before you can configure the corresponding interfaces for VPLS.**

Because a VPLS carries Ethernet traffic across a service provider network, it must mimic an Ethernet network in some ways. When a PE router configured with a VPLS routing instance receives a packet from a CE device, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, it forwards the packet to the appropriate PE router or CE device. If it does not, it broadcasts the packet to all other PE routers and CE devices that are members of that VPLS routing instance. In both cases, the CE device receiving the packet must be different from the one sending the packet.

When a PE router receives a packet from another PE router, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, the PE router either forwards the packet or drops it depending on whether the destination is a local or remote CE device:

- If the destination is a local CE device, the PE router forwards the packet to it.
- If the destination is a remote CE device (connected to another PE router), the PE router discards the packet.

If the PE router cannot determine the destination of the VPLS packet, it floods the packet to all attached CE devices. Figure 146 illustrates this process.

Copyright © 2016, Juniper Networks, Inc.

Figure 146: Flooding a Packet with an Unknown Destination

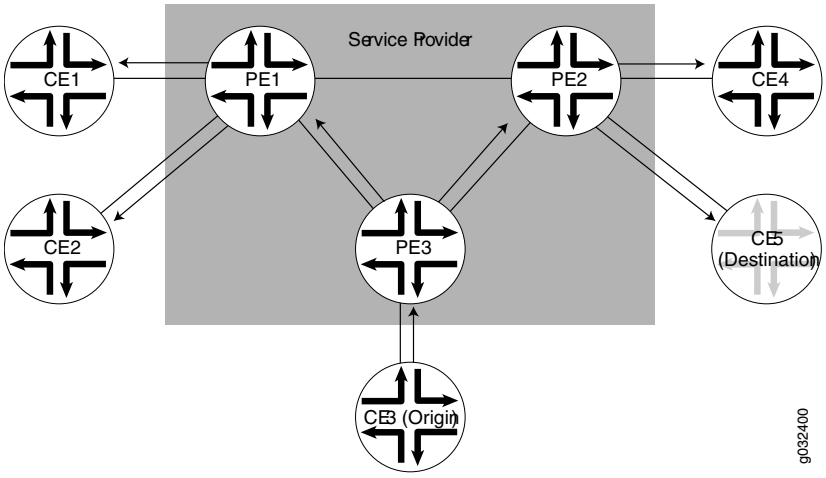

A VPLS interface can be directly connected to an Ethernet switch. Layer 2 information gathered by an Ethernet switch, for example, MAC addresses and interface ports, is included in the VPLS routing instance table.

An MPLS label-switched interface (LSI) label is used as the inner label for VPLS. This label maps to a VPLS routing instance on the ingress PE router. On the egress PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

One restriction on flooding behavior in VPLS is that traffic received from remote PE routers is never forwarded to other PE routers. This restriction helps prevent loops in the core network. However, if a CE Ethernet switch has two or more connections to the same PE router, you must enable the Spanning Tree Protocol (STP) on the CE switch to prevent loops.

NOTE: Under certain circumstances, VPLS PE routers might duplicate an Internet Control Message Protocol (ICMP) reply from a CE device when a PE router has to flood an ICMP request because the destination MAC address has not yet been learned. The duplicate ICMP reply can be triggered when a CE device with promiscuous mode enabled is connected to a PE router. The PE router automatically floods the promiscuous mode enabled CE device, which then returns the ICMP request to the VPLS PE routers. The VPLS PE routers consider the ICMP request to be new and flood the request again, creating a duplicate ping reply.

Copyright © 2016, Juniper Networks, Inc.

Using an Ethernet Switch as the VPLS CE Device

For VPLS configurations, the CE device does not necessarily need to be a router. You can link the PE routers directly to Ethernet switches. However, be aware of the following configuration issues:

- When you configure VPLS routing instances and establish two or more connections between a CE Ethernet switch and a PE router, you must enable the Spanning Tree Protocol (STP) on the switch to prevent loops.
- Junos OS allows standard bridge protocol data unit (BPDU) frames to pass through emulated Layer 2 connections, such as those configured with Layer 2 VPNs, Layer 2 circuits, and VPLS instances. However, CE Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routing platforms configured for these emulated Layer 2 connections.

VPLS Exceptions on SRX Series Devices

The VPLS implementation on SRX Series device is similar to VPLS implementations on M Series, T Series, and MX Series routers, with the following exceptions:

- SRX Series devices do not support aggregated Ethernet interfaces. Therefore, aggregated Ethernet interfaces between CE devices and PE routers are not supported for VPLS routing instances on SRX Series devices.
- SRX Series devices do not support aggregated Ethernet interfaces. Therefore, aggregated Ethernet interfaces between PE devices and PE routers are not supported for VPLS routing instances on SRX Series devices.
- VPLS multihoming, which allows connecting a CE device to multiple PE routers to provide redundant connectivity, is not supported on SRX Series devices.
- SRX Series devices do not support BGP mesh groups.
- SRX Series devices support only the following encapsulation types on VPLS interfaces that face CE devices: extended VLAN VPLS, Ethernet VPLS, and VLAN VPLS. Ethernet VPLS over ATM LLC encapsulation is not supported.
- Virtual ports are generated dynamically on a Tunnel Services PIC on some Juniper Networks routing platforms. SRX Series devices do not support Tunnel Services modules or virtual ports.
- The VPLS implementation on SRX Series devices does not support dual-tagged frames. Therefore, VLAN rewrite operations are not supported on dual-tagged frames. VLAN rewrite operations such as pop-pop, pop-swap, push-push, swap-push, and swap-swap, which are supported on M Series and T Series routing platforms, are not supported on SRX Series devices.

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- Understanding VPLS Interfaces on page 3849
- Understanding VPLS Routing Instances on page 3859

Copyright © 2016, Juniper Networks, Inc.

- Understanding VPLS VLAN Encapsulation on page 3905
- Understanding VPLS VLAN Encapsulation on a Logical Interface on page 3906
- VPLS Configuration Overview on page 3846

VPLS Configuration Overview

To configure VPLS functionality, you must enable VPLS support on the provider edge (PE) routers. You must also configure PE routers to distribute routing information to the other PE routers in the VPLS and configure the circuits between the PE routers and the customer edge (CE) devices, as explained in the steps that follow.

NOTE: **Many configuration procedures for VPLS are identical to the procedures for Layer 2 and Layer 3 VPNs.**

To configure VPLS:

1. Determine which uPIM and ePIM ports correspond to the interfaces that will carry the VPLS traffic and enable routing mode on those ports.
2. Configure the interfaces that will carry the VPLS traffic between the PE router and CE devices. On the PE router interfaces that are facing the CE devices, specify a VPLS encapsulation type. The type of encapsulation depends on the interface type. See"Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
3. Create a VPLS routing instance on each PE router that is participating in the VPLS. For each VPLS routing instance, specify which interfaces will carry the VPLS traffic between the PE and CE devices. On the CE device interface that faces the PE router, you must specify inet (for IPv4), and include the IP address. Additionally, each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP identify overlapping network layer reachability information (NLRI) messages from different VPNs.) See "Example: Configuring the VPLS Routing Instance" on page 3862.
4. Configure routing options on the PE router. See "Example: Configuring Routing Options on the VPLS PE Router" on page 3903.
5. Configure MPLS LSPs between the PE routers. See "Example: Configuring MPLS on the VPLS PE Router" on page 3869.
6. Configure RSVP on the PE routers. Enable RSVP for all connections that participate in the MPLS LSP. See "Example: Configuring RSVP on the VPLS PE Router" on page 3868.
7. Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See "Example: Configuring BGP on the VPLS PE Router" on page 3902.

Copyright © 2016, Juniper Networks, Inc.

8. Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
9. Configure VLAN encapsulation. See "Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces" on page 3909,"Example: Configuring VPLS VLAN Encapsulation" on page 3906, and "Example: Configuring Extended VLAN VPLS Encapsulation" on page 3910.

Related Documentation

- MPLS Layer 2 VPN Configuration Overview on page 3799
- MPLS Layer 2 VPN Configuration Overview on page 3799
- Example: Configuring MPLS on the VPLS PE Router on page 3869
- Example: Configuring RSVP on the VPLS PE Router on page 3868
- Example: Configuring BGP on the VPLS PE Router on page 3902
- Example: Configuring OSPF on the VPLS PE Router on page 3867

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 169

Configuring Interfaces

- Understanding VPLS Interfaces on page 3849
- Example: Configuring Routing Interfaces on the VPLS PE Router on page 3851
- Example: Configuring the Interface to the VPLS CE Device on page 3852
- VPLS Filters and Policers Overview on page 3853
- Example: Configuring VPLS Filters on page 3853
- Example: Configuring VPLS Policers on page 3855

Understanding VPLS Interfaces

For each VPLS routing instance on a PE router, you specify which interfaces are to be used to carry VPLS traffic between the PE and CE devices.

This topic contains the following sections:

- Interface Name on page 3849
- Encapsulation Type on page 3849
- Flexible VLAN Tagging on page 3850
- VLAN Rewrite on page 3850

Interface Name

Specify both the physical and logical portions of the interface name, in the following format:

physical.logical

For example, in ge-1/2/1.2, ge-1/0/1 is the physical portion of the interface name and 2 is the logical portion. If you do not specify the logical portion of the interface name, 0 is set by default. A logical interface can be associated with only one routing instance.

Encapsulation Type

The physical link-layer encapsulation type for a VPLS interface can be one of the following:

Copyright © 2016, Juniper Networks, Inc.

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol Identifier (TPID) values.
- **extended-vlan-vpls**—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. All VLAN IDs from 1 through 1023 are valid for VPLS VLANs on Fast Ethernet interfaces, and all VLAN IDs from 1 through 4094 are valid for VPLS VLANs on Gigabit Ethernet interfaces.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. You must configure this encapsulation type on both the physical interface and the logical interface. VLAN IDs 1 through 511 are reserved for normal Ethernet VLANs, IDs 512 through 1023 are reserved for VPLS VLANs on Fast Ethernet interfaces, and IDs 512 through 4094 are reserved for VPLS VLANs on Gigabit Ethernet interfaces.
- **flexible-ethernet-services**—Use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, and VPLS encapsulations on a single physical port. Aggregated Ethernet bundles cannot use this encapsulation type.

  For flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

Flexible VLAN Tagging

For untagged packets to be accepted on an 802.1Q VLAN-tagged port, specify the native VLAN ID with the flexible VLAN tagging option. (No other flexible VLAN tagging features are supported.)

VLAN Rewrite

You can rewrite VLAN tags on VPLS interfaces. Rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between CE devices that share a VLAN ID.

You can configure rewrite operations to stack (push), remove (pop), or rewrite (swap) tags on single-tagged frames. If a port is not configured for VLAN tagging, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- pop—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- push—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- swap—Replace the VLAN tag at the top of the VLAN tag stack with a user-specified VLAN tag value.

Copyright © 2016, Juniper Networks, Inc.

You perform VLAN rewrite operations by applying input and output VLAN maps at the ingress and egress, respectively, of the interface. For incoming frames, use the input-vlan-map; for outgoing frames, use the output-vlan-map.

The VPLS implementation on SRX Series devices does not support dual-tagged frames. Therefore, VLAN rewrite operations are not supported on dual-tagged frames. VLAN rewrite operations such as pop-pop, pop-swap, push-push, swap-push, and swap-swap, which are supported on M Series and T Series routing platforms, are not supported on SRX Series devices.

Related Documentation

- Example: Configuring Routing Interfaces on the VPLS PE Router on page 3851
- Example: Configuring the Interface to the VPLS CE Device on page 3852
- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841
- Understanding VPLS VLAN Encapsulation on page 3905

Example: Configuring Routing Interfaces on the VPLS PE Router

This example shows how to configure routing interfaces on the VPLS PE router.

- Requirements on page 3851
- Overview on page 3851
- Configuration on page 3851
- Verification on page 3852

Requirements

Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.

Overview

In this example, you configure the PE1 router loopback interface and the interface to the PE2 router ge-2/0/1.

Configuration

Step-by-Step Procedure

To configure the routing interface on the VPLS PE router:

1. Configure the loopback interface.

   ```
   [edit]
   user@host# set interfaces lo0 unit 0 family inet address 10.255.7.168/32 primary
   ```

2. Configure the IP address on the MPLS core interface.

   ```
   [edit]
   user@host# set interfaces ge-3/0/2 unit 0 family inet address 100.1.1.1/30
   ```

3. Configure the MPLS family.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set interfaces ge-3/0/2 unit 0 family mpls
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- *Interfaces Feature Guide for Security Devices*
- VPLS Configuration Overview on page 3846
- Understanding VPLS Interfaces on page 3849

Example: Configuring the Interface to the VPLS CE Device

This example shows how to configure the router interface that is connected to the CE device to include VPLS encapsulation.

- Requirements on page 3852
- Overview on page 3852
- Configuration on page 3852
- Verification on page 3853

Requirements

Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.

Overview

In this example, you configure the router interface ge-1/2/1 that is connected to the CE device to include VPLS encapsulation.

Configuration

Step-by-Step Procedure

To configure the interface to the VPLS CE device:

1. Configure VPLS encapsulation for the interface facing the CE router.

```
[edit]
user@host# set interfaces ge-1/2/1 encapsulation ethernet-vpls
```

2. Configure the interface for the VPLS family group.

```
[edit]
user@host# set interfaces ge-1/2/1 unit 0 family vpls
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show interfaces ge-1/2/1** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- Understanding VPLS Interfaces on page 3849

VPLS Filters and Policers Overview

This feature permits users to configure both firewall filters and policers for virtual private LAN service (VPLS). Firewall filters enable you to filter packets based on their components and perform an action on packets that match the filter. Policers enable you to limit the amount of traffic that passes into or out of an interface.

This feature can be enabled by configuring VPLS filters, policers, and accounting through various CLI commands. VPLS filters and policers act on a Layer 2 frame that includes the media access control (MAC) header (after any VLAN rewrite or other rules are applied), but that does not include the cyclical redundancy check (CRC) field.

NOTE: **You can apply VPLS filters and policers on the PE routers only to customer-facing (PE-CE) interfaces.**

Related Documentation

- Example: Configuring VPLS Policers on page 3855
- Example: Configuring VPLS Filters on page 3853

Example: Configuring VPLS Filters

This example shows how to configure VPLS filters.

- Requirements on page 3853
- Overview on page 3854
- Configuration on page 3854
- Verification on page 3855

Requirements

Before you begin:

- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.

Copyright © 2016, Juniper Networks, Inc.

- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE on the PE routers. See "Example: Configuring RSVP on the VPLS PE Router" on page 3868.

Overview

This example describes how to configure filtering and accounting for VPLS.

CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your device to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device.

Configuration

CLI Quick Configuration

To quickly configure VPLS filters, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set firewall family vpls filter blue term term1 from interface ge-3/0/0.512
set firewall family vpls filter blue term term1 from interface fe-5/0/0.512
set firewall family vpls filter blue term term1 then count count1
set firewall family vpls filter blue accounting-profile fw_profile
set accounting-options file fw_acc size 500k
set accounting-options file fw_acc transfer-interval 5
set accounting-options filter-profile fw_profile file fw_acc
set accounting-options filter-profile fw_profile interval 1
set accounting-options filter-profile fw_profile counters count1
set interfaces ge-0/0/1 unit 512 family vpls filter input blue
```

Step-by-Step Procedure

To configure filters for VPLS:

1. Configure a filter with a GE interface as the match condition and count as the action.

   ```
   [edit ]
   user@host# set firewall family vpls filter blue term term1 from interface ge-3/0/0.512
   ```

2. Configure a filter with an FE interface as the match condition and count as the action.

   ```
   [edit ]
   user@host# set firewall family vpls filter blue term term1 from interface fe-5/0/0.512
   ```

3. Configure the count.

   ```
   [edit ]
   user@host# set firewall family vpls filter blue term term1 then count count1
   ```

4. Configure the accounting profile to refer it to the counter.

Copyright © 2016, Juniper Networks, Inc.

```
[edit ]
user@host# set firewall family vpls filter blue accounting-profile fw_profile
```

5. Configure the account file size.

```
[edit ]
user@host# set accounting-options file fw_acc size 500k
```

6. Configure the account transfer interval.

```
[edit ]
user@host# set accounting-options file fw_acc transfer-interval 5
```

7. Configure the filter for the accounting profile.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile file fw_acc
```

8. Configure the filter for the interval.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile interval 1
```

9. Configure the counter.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile counters count1
```

10. Apply the filter to the interface.

```
[edit ]
user@host# set interfaces ge-0/0/1 unit 512 family vpls filter input blue
```

11. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show firewall** and **show accounting records** commands.

Related Documentation

- VPLS Filters and Policers Overview on page 3853
- VPLS Configuration Overview on page 3846
- Example: Configuring VPLS Policers on page 3855

Example: Configuring VPLS Policers

This example shows how to configure VPLS policers.

- Requirements on page 3856
- Overview on page 3856
- Configuration on page 3856
- Verification on page 3857

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE on the PE routers. See "Example: Configuring RSVP on the VPLS PE Router" on page 3868.

Overview

This example describes how to configure policing and apply it on the interface for VPLS.

CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your device to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device.

Configuration

CLI Quick Configuration

To quickly configure VPLS policers, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set firewall policer police2 if-exceeding bandwidth-percent 10
set firewall policer police2 if-exceeding burst-size-limit 1500
set firewall policer police2 then discard
set interfaces ge-0/0/1 unit 512 family vpls policer input police2
```

Step-by-Step Procedure

To configure filters for VPLS:

1. Configure bandwidth percentage.

   ```
   [edit ]
   user@host# set firewall policer police2 if-exceeding bandwidth-percent 10
   ```

2. Configure the burst size limit.

   ```
   [edit ]
   user@host# set firewall policer police2 if-exceeding burst-size-limit 1500
   ```

3. Configure the terminal action on the packet.

Copyright © 2016, Juniper Networks, Inc.

```
[edit ]
user@host# set firewall policer police2 then discard
```

4. Apply the policer to the interface.

```
[edit ]
user@host# set interfaces ge-0/0/1 unit 512 family vpls policer input police2
```

5. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show firewall** command.

Related Documentation

- VPLS Filters and Policers Overview on page 3853
- VPLS Configuration Overview on page 3846
- Example: Configuring VPLS Filters on page 3853

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 170

Configuring Routing Instances

- Understanding VPLS Routing Instances on page 3859
- Example: Configuring the VPLS Routing Instance on page 3862
- Example: Configuring Automatic Site Identifiers for VPLS on page 3864

Understanding VPLS Routing Instances

To configure VPLS functionality, you must enable VPLS support on the PE router. You must also configure PE routers to distribute routing information to the other PE routers in the VPLS and configure the circuits between the PE routers and the CE devices.

You create a VPLS routing instance on each PE router that is participating in the VPLS. The routing instance has the same name on each PE router. To configure the VPLS routing instance, you specify the following:

- Route distinguisher—Helps BGP distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPLS instances. Each routing instance that you configure on a PE router must have a unique route distinguisher.
- Route target—Defines which route is part of a VPLS. A unique route target helps distinguish between different VPLS services on the same router.
- Site name—Provides unique name for the VPLS site.
- Site identifier—Provides unique numerical identifier for the VPLS site.
- Site range—Specifies total number of sites in the VPLS. The site range must be greater than the site identifier.
- Interface to the CE router—Specifies the physical interface to the CE router that carries VPLS traffic. The interface must be configured for a VPLS encapsulation type.

NOTE: In addition to the VPLS routing instance, you must configure MPLS label-switched paths (LSPs) between the PE routers, internal BGP (IBGP) sessions between the PE routers, and an interior gateway protocol (IGP) on the PE routers.

Copyright © 2016, Juniper Networks, Inc.

CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your router to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the router performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the router.

This topic contains the following sections:

- BGP Signaling on page 3860
- VPLS Routing Table on page 3860
- Trace Options on page 3861

BGP Signaling

BGP is used to signal the paths between each of the PE routers participating in the VPLS routing instance. These paths carry VPLS traffic across the service provider's network between the VPLS sites.

NOTE: LDP signaling is not supported for the VPLS routing instance.

To configure BGP signaling, you specify the following:

- VPLS site name and site identifier—When you configure BGP signaling for the VPLS routing instance, you must specify each VPLS site that has a connection to the router. For each VPLS site, you must configure a site name and site identifier (a numerical identifier between 1 to 65,534 that uniquely identifies the VPLS site).
- Site range—When you enable BGP signaling for the VPLS routing instance, you need to configure a site range. The site range specifies the total number of sites in the VPLS.

  NOTE: The site range value must be greater than the largest site identifier.

- Site preference—You can specify the preference value advertised for a particular VPLS site. The site preference value is encoded in the BGP local preference attribute. When a PE router receives multiple advertisements with the same VPLS edge (VE) device identifier, the advertisement with the highest local preference value is preferred.

VPLS Routing Table

The VPLS routing table contains MAC addresses and interface information for both physical and virtual ports. You can configure the following characteristics for the table:

- Table size—You can modify the size of the VPLS MAC address table. The default table size is 512 MAC addresses; the minimum is 16 addresses, and the maximum is 65,536 addresses.

Copyright © 2016, Juniper Networks, Inc.

If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

The interfaces affected include all of the interfaces within the VPLS routing instance, including the local interfaces and the LSI interfaces.

- Timeout interval—You can modify the timeout interval for the VPLS table. The default timeout interval is 300 seconds; the minimum is 10 seconds, and the maximum is 1,000,000 seconds. We recommend you configure longer values for small, stable VPLS networks and shorter values for large, dynamic VPLS networks. If the VPLS table does not receive any updates during the timeout interval, the router waits one additional interval before automatically clearing the MAC address entries from the VPLS table.
- Number of addresses learned from an interface—You can configure a limit on the number of MAC addresses learned by a VPLS routing instance by setting the MAC table size. The default is 512 addresses; the minimum is 16, and the maximum is 65,536 addresses. If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

  Because this limit applies to each VPLS routing instance, the MAC addresses of a single interface can consume all the available space in the table, preventing the routing instance from acquiring addresses from other interfaces. You can limit the number of MAC addresses learned from all interfaces configured for a VPLS routing instance, as well as limit the number of MAC addresses learned from a specific interface.

  The MAC limit configured for an individual interface overrides the limit configured for all interfaces for the VPLS routing instance. Also, the table limit can override the limits configured for the interfaces.

  The MAC address limit applies only to interfaces to CE devices.

Trace Options

The following trace flags display operations associated with VPLS:

- **all**—All VPLS tracing options
- **connections**—VPLS connections (events and state changes)
- **error**—Error conditions
- **nlri**—VPLS advertisements received or sent using BGP
- **route**—Trace-routing information
- **topology**—VPLS topology changes caused by reconsideration or advertisements received from other PE routers using BGP

Related Documentation

- Example: Configuring the VPLS Routing Instance on page 3862
- Example: Configuring Routing Options on the VPLS PE Router on page 3903

Copyright © 2016, Juniper Networks, Inc.

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841
- Understanding VPLS VLAN Encapsulation on page 3905

Example: Configuring the VPLS Routing Instance

This example shows how to create a VPLS routing instance on each PE router that is participating in the VPLS.

- Requirements on page 3862
- Overview on page 3862
- Configuration on page 3862
- Verification on page 3864

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.

Overview

This example describes how to create a VPLS routing instance; configure VPLS site identifier, site range, no tunnel services option, route distinguisher, and route target for the VPLS routing instance; and specify the VPLS interface to the CE router.

NOTE: You must specify no tunnel services in the VPLS routing instance configuration, because SRX Series devices do not support tunnel serial PICs.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances green instance-type vpls
set routing-instances green protocols vpls site-range 10 site R3 site-identifier 2
set routing-instances green protocols vpls no-tunnel-services
set routing-instances green route-distinguisher 10.255.7.1:1
set routing-instances green vrf-target target:11111:1
set routing-instances green instance-type vpls interface ge-1/2/1.0
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a VPLS routing instance:

1. Configure the routing instance of type VPLS.

   ```
   [edit]
   user@host# edit routing-instances green
   ```

2. Enable the VPLS instance type.

   ```
   [edit routing-instances green]
   user@host# set instance-type vpls
   ```

3. Configure the VPLS site identifier and range for the VPLS routing instance.

   ```
   [edit routing-instances green protocols vpls]
   user@host# set site-range 10 site R3 site-identifier 2
   ```

4. Configure the no-tunnel-services option for the VPLS routing instance.

   ```
   [edit routing-instances green protocols vpls]
   user@host# set no-tunnel-services
   ```

5. Configure the route distinguisher.

   ```
   [edit routing-instances green]
   user@host# set route-distinguisher 10.255.7.1:1
   ```

6. Configure the route target.

   ```
   [edit routing-instances green]
   user@host# set vrf-target target:11111:1
   ```

7. Specify the VPLS interface to the CE router.

   ```
   [edit routing-instances green]
   user@host# set instance-type vpls interface ge-1/2/1.0
   ```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances green** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances green
instance-type vpls;
interface ge-1/2/1.0;
route-distinguisher 10.255.7.1:1;
vrf-target target:11111:1;
protocols {
  vpls {
    site-range 10;
    no-tunnel-services;
    site R3 {
      site-identifier 2;
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying VPLS Routing Instance Is Configured on page 3864
- Verifying VPLS Routing Attributes Are Configured on page 3864

Verifying VPLS Routing Instance Is Configured

Purpose Verify that the VPLS routing instance is configured.

Action From operational mode, enter the **show routing-instances** command.

Verifying VPLS Routing Attributes Are Configured

Purpose Verify that attributes such as VPLS site identifier, site range, no tunnel services option, route distinguisher, and route target for the VPLS routing instance are configured.

Action From operational mode, enter the **show routing-instances green protocols vpls** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- Understanding VPLS Routing Instances on page 3859

Example: Configuring Automatic Site Identifiers for VPLS

This example shows how to configure automatic site identifiers for VPLS sites.

Requirements

Before you begin, see information on selective stateless packet-based services in *Interfaces Feature Guide for Security Devices*.

Overview

When you enable automatic site identifiers, the Junos OS automatically assigns site identifiers to VPLS sites. In this example, you configure a routing instance called vpls instance and enable automatic site identifiers for VPLS.

NOTE: **Site identifiers for VPLS sites can be different for different routing instances.**

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure automatic site identifiers:

1. Configure the routing instance of type VPLS.

   ```
   [edit]
   user@host#set routing-instances vpls-instance
   ```

2. Enable automatic site identifiers.

   ```
   [edit routing-instances vpls-instance]
   user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id
     collision-detect-time 10
   user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id
     new-site-wait-time 20
   user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id
     reclaim-wait-time minimum 5 maximum 20
   user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id
     startup-wait-time 5
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show vpls connections** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 171

Configuring Routing and Signaling Protocols

- Example: Configuring OSPF on the VPLS PE Router on page 3867
- Example: Configuring RSVP on the VPLS PE Router on page 3868
- Example: Configuring MPLS on the VPLS PE Router on page 3869
- Example: Configuring LDP on the VPLS PE Router on page 3871
- Example: Configuring VPLS over GRE with IPsec VPNs on page 3872
- Example: Configuring VPLS with BGP Signaling on page 3889
- Example: Configuring BGP on the VPLS PE Router on page 3902
- Example: Configuring Routing Options on the VPLS PE Router on page 3903

Example: Configuring OSPF on the VPLS PE Router

This example shows how to configure OSPF on the VPLS PE router.

- Requirements on page 3867
- Overview on page 3868
- Configuration on page 3868
- Verification on page 3868

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.

Copyright © 2016, Juniper Networks, Inc.

Overview

The PE routers exchange routing information using an IGP such as OSPF. In this example, you configure OSPF area 0.0.0.0 on the VPLS PE router and traffic engineering for OSPF.

Configuration

Step-by-Step Procedure

To configure OSPF on the VPLS PE router:

1. Configure the OSPF area on the VPLS PE router.

   ```
   [edit]
   user@host# set protocols ospf area 0.0.0.0 interface t1-1/0/1.0
   user@host# set protocols ospf area 0.0.0.0 interface lo0.0
   ```

2. Configure traffic engineering for OSPF.

   ```
   [edit]
   user@host# set protocols ospf traffic-engineering
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring RSVP on the VPLS PE Router

This example shows how to configure RSVP on the VPLS PE router.

- Requirements on page 3868
- Overview on page 3869
- Configuration on page 3869
- Verification on page 3869

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See"Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.

Copyright © 2016, Juniper Networks, Inc.

- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.

Overview

This example describes how to enable RSVP for all connections that participate in the LSP on the PE1 router.

Configuration

Step-by-Step Procedure

To configure RSVP on the VPLS PE router:

1. Configure the interface to the PE2 router for RSVP.

   ```
   [edit ]
   user@host# set protocols rsvp interface t1-1/0/1.0
   ```

2. Configure the loopback interface for RSVP.

   ```
   [edit ]
   user@host# set protocols rsvp interface lo0.0
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring MPLS on the VPLS PE Router

This example shows how to configure MPLS on the VPLS PE router.

- Requirements on page 3870
- Overview on page 3870
- Configuration on page 3870
- Verification on page 3871

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See"Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE on the PE routers. See "Example: Configuring RSVP on the VPLS PE Router" on page 3868.

Overview

This example shows you how to configure MPLS on the PE1 router to advertise the Layer 2 VPN interface that communicates with the PE2 router.

CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your router to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the router performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the router.

Configuration

Step-by-Step Procedure

To configure MPLS on the VPLS PE router:

1. Configure the interface to the PE2 router for MPLS.

   ```
   [edit ]
   user@host# set protocols mpls interface t1-1/0/1.0
   ```

2. Configure the loopback for MPLS.

   ```
   [edit ]
   user@host# set protocols mpls interface lo0.0
   ```

3. Configure the path to destination 10.255.7.164.

   ```
   [edit ]
   user@host# set protocols mpls label-switched-path chelsea-sagar to 10.255.7.164
   ```

4. If you are done configuring the device, commit the configuration.

   ```
   [edit ]
   user@host# commit
   ```


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show mpls** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring LDP on the VPLS PE Router

This example shows how to configure LDP on the VPLS PE router.

- Requirements on page 3871
- Overview on page 3871
- Configuration on page 3871
- Verification on page 3872

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.

Overview

This example describes how to enable LDP for all connections that participate in the LSP on the PE1 router.

Configuration

Step-by-Step Procedure

To configure LDP on the VPLS PE router:

1. Configure the interface to the PE2 router for LDP.

   ```
   [edit ]
   user@host# set protocols ldp interface ge-3/0/2
   ```

2. Configure the loopback interface for LDP.

   ```
   [edit ]
   user@host# set protocols ldp interface lo0
   ```

3. If you are done configuring the device, commit the configuration.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring VPLS over GRE with IPsec VPNs

This example demonstrates a network scenario consisting of a central office and one branch office that will use VPLS, MPLS, GRE, and IPsec to create secure Ethernet connectivity over a Layer 3 network. This configuration can be expanded to add many other branch sites.

- Requirements on page 3872
- Overview on page 3872
- Configuration on page 3877
- Verification on page 3888

Requirements

Before you begin:

- Ensure that a layer 3 network is in place for all branch offices and that there is an ingress (head-end) device at the central office configured to terminate the VPNs from each branch office.
- Obtain IDP licenses for each SRX Series device. IDP is used to reassemble GRE packets that might become fragmented.

Overview

Junos OS can selectively choose whether traffic is processed by the flow engine or packet engine using the selective stateless packet-based feature. This feature allows you to combine flow and packet-based services in a single device. In this example, we describe a deployment scenario that uses this feature to deploy large-scale VPLS over GRE. This enables branch devices to securely transport Ethernet traffic over Layer 3 networks when used in conjunction with IPsec.

In this scenario you configure a central office ingress (head-end) using an SRX650 device and one branch office using an SRX240 device. This setup is accomplished by carrying MPLS pseudowires over GRE, which in turn, is encapsulated in IPsec in order to guarantee data integrity and confidentiality. By default, SRX Series devices use secure flow forwarding. Because VPLS services are provided in packet-mode only, the configuration requires the GRE tunnel to be terminated in a packet-mode routing instance (the default routing instance).

Copyright © 2016, Juniper Networks, Inc.

NOTE: You can also use an MX Series device as the ingress (head-end) device, which is mentioned later in this topic.

To better understand this configuration, we will discuss two scenarios. The first scenario uses pseudowires to allow the creation of point-to-point circuits between two endpoints carried over the MPLS network. If we leave the signaling protocols aside (that is, there are a few ways to provision the pseudowires), these connections are just point-to-point connections. Using this approach provides an end-to-end wire between sites. This is beneficial from a traffic processing point of view because the gateways do not need to do MAC address learning, they simply forward anything they receive to the pseudowire. Because of this, it may be difficult to deploy this setup when trying to provide connectivity to multiple branch offices.

The second scenario could use VPLS to provide a Layer 2 network abstraction. With VPLS, endpoints are expected to negotiate LSPs and pseudowires with every other endpoint (that is, they are fully meshed). When a node receives an Ethernet frame from one of its LAN interfaces the source MAC address is learned, if it's not already known, and flooded using every pseudowire connecting to all other branch nodes. However, if the destination has been previously learned, then the frame is sent to the appropriate destination. When an Ethernet frame is received through one of the pseudowires (that is, from the MPLS network), source MAC address learning is performed. The next time a frame is sent to that MAC it does not need to be flooded and the frame is flooded to every single LAN interface in the node, but not over the pseudowires. In other words, the network acts as a distributed Layer 2 switch providing any-to-any Ethernet connectivity between the devices connected to the different nodes in the network.

While the advantages of this second scenario is evident (any-to-any connectivity, automated provisioning, and simple abstraction), it comes at the cost of complexity. Every PE node has to perform Layer 2 learning and flooding of traffic, which can cause problems when either multiple broadcast/multicast or frames to unknown MAC addresses are used. As an example, if you had a topology with a thousand branch offices, each office that receives a broadcast packet must replicate it 999 times, encapsulate each copy in GRE and IPsec and forward the resulting traffic. Additionally, because each node performs Layer 2 learning, there are limitations in the maximum number of MAC addresses that each node can learn, limiting the total number of nodes in the domain.

In this example, we use a hybrid approach to these two scenarios. We use a circuit cross connect (CCC) at each branch office stitched to a VPLS instance at central office (ingress). This solution makes sense if most of the traffic flows from the branch offices to central office, and the branch-to-branch office traffic is always forwarded through the hub. The use of CCCs at branch offices combined with VPLS stitching at the central office provides a scalable way to deploy large hub-and-spoke topologies where Ethernet must be transported over an IP network (with or without encryption). At the expense of configuration complexity, it is possible to use branch SRX Series devices to terminate such connections, providing a scalable and cost-effective way to deploy small-to-large networks where Ethernet traffic is carried transparently using lower cost IP connections. Figure 147 shows this topology.

Copyright © 2016, Juniper Networks, Inc.

Figure 147: VPLS Deployment Scenario

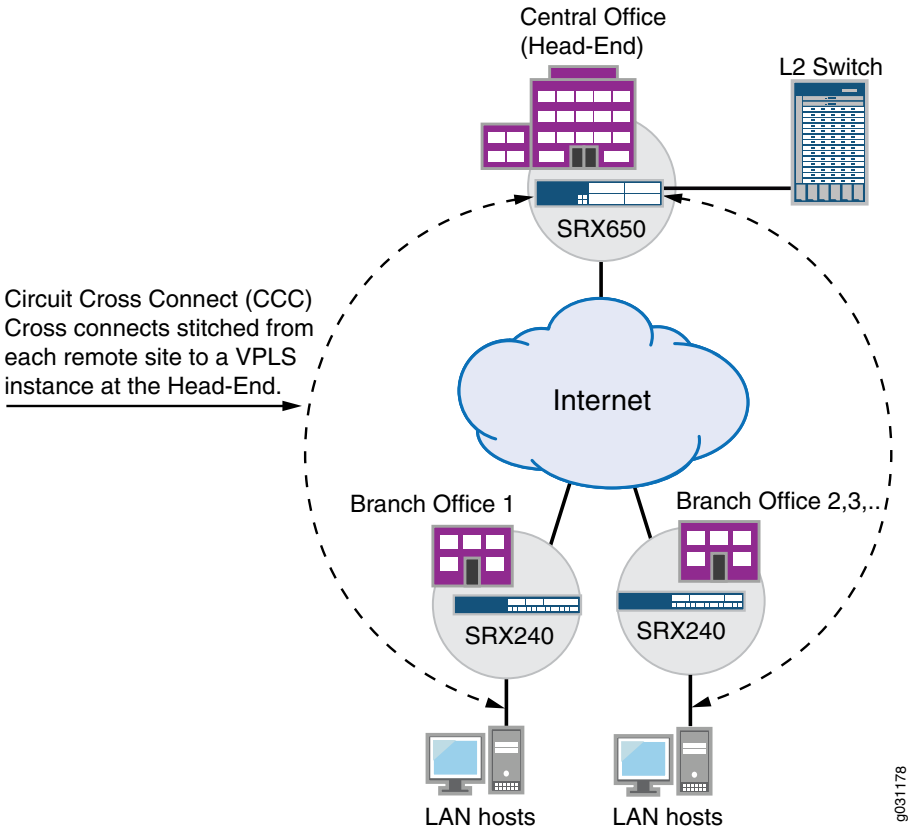

In this deployment, VPLS services are provided only in packet mode and must be configured in the default routing instance. Unfortunately, IPsec is only provided in flow mode. Hence, a flow-mode routing-instance is used that provides both GRE reassembly and IPsec termination. While the GRE termination is done in the default routing instance, a flow-mode routing instance is connected between the default routing instance and the Internet (or whatever Layer 3 network is used as a transport), and it terminates the IPsec tunnel towards the ingress device. Because it is likely that a single public IP address is available, the Internet-facing Interface is connected to the default routing instance and is used to terminate IKE; however, the tunnel interface (st0) is bound to the flow-mode routing instance. See Figure 148.

Copyright © 2016, Juniper Networks, Inc.

Figure 148: Branch Office Circuit Cross Connect Termination

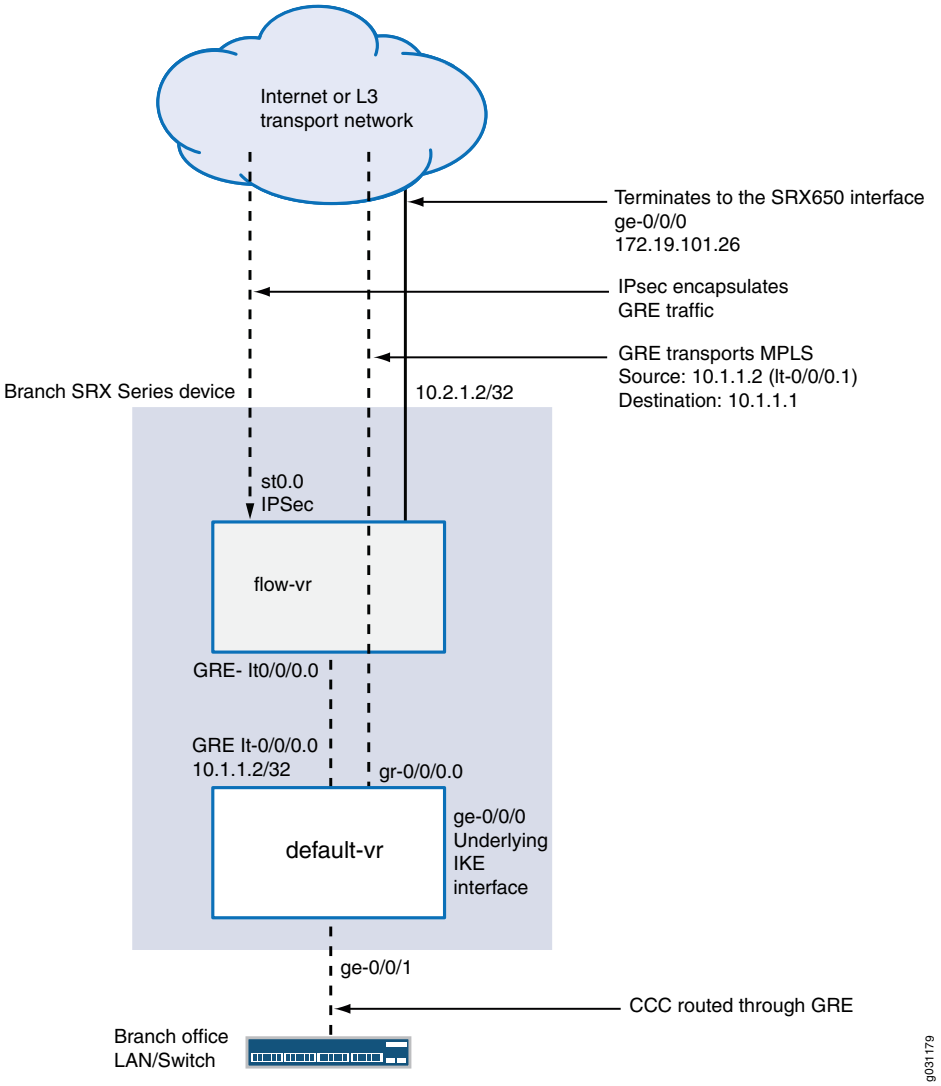

When configuring the central office SRX650, the first thing you do is terminate the IPsec tunnels, GRE, and CCC connections. Because a branch SRX Series device is used as the ingress (head-end), the configuration to terminate the CCC circuits is identical to the one used at each branch office, with the exception that instead of one tunnel, multiple tunnels (and pseudowires) are terminated.

The pseudowires are stitched to a VPLS routing instance using logical tunnel (lt) interfaces. It is possible to use an lt interface unit to terminate a CCC connection and connect this unit to a different unit that is part of a VPLS routing instance. The overall result is as if the pseudowires were terminated directly in the VPLS routing instance. Figure 149 illustrates this configuration.

Copyright © 2016, Juniper Networks, Inc.

Figure 149: Central Office Ingress (Head-End) Configuration with an SRX Series Device

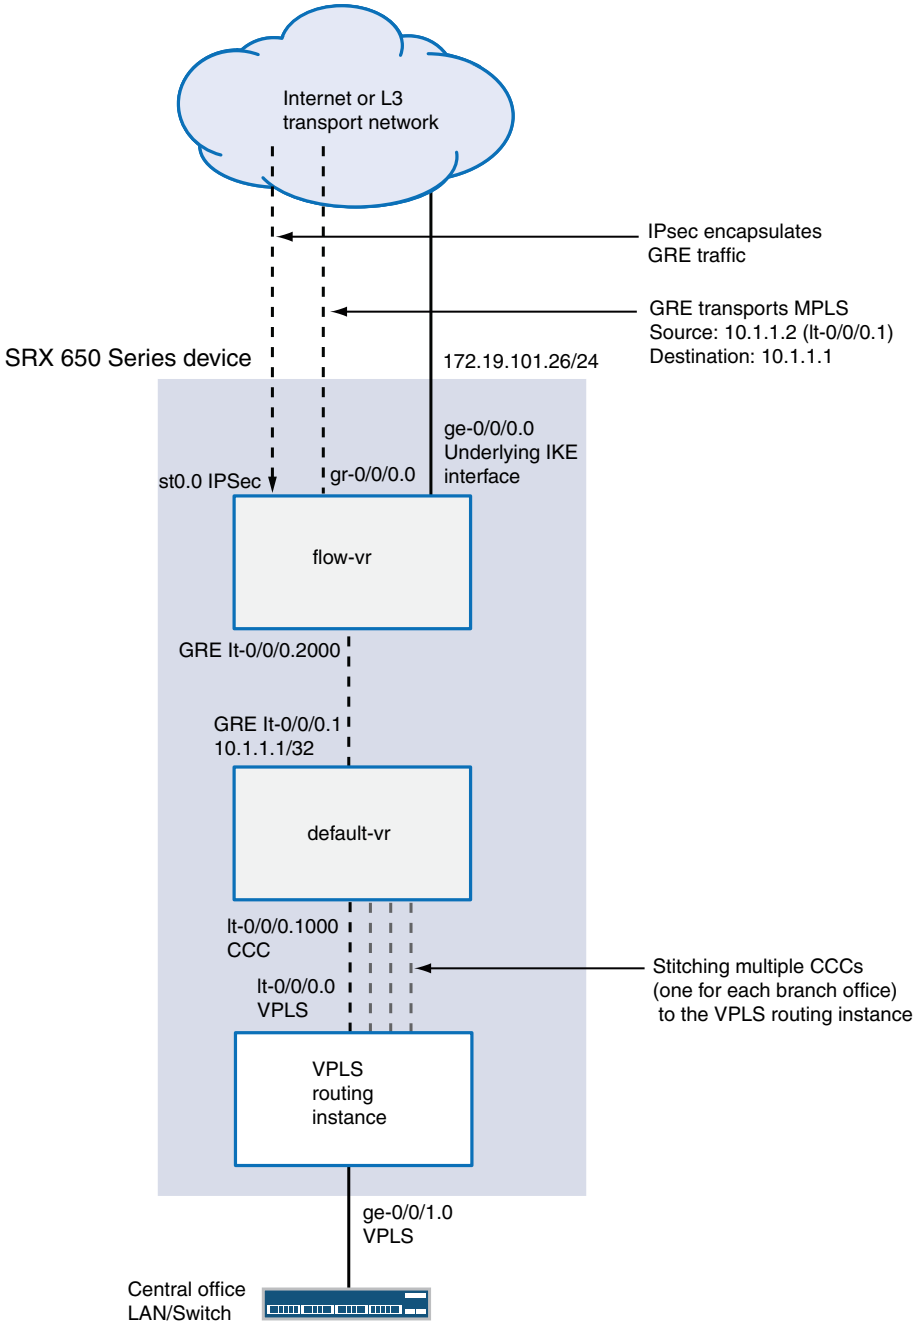

You can also use an MX Series device as the central office ingress (head-end) to terminate all branch office connections. The differences in the configuration are due to the way IPsec is configured and the fact that on MX Series devices IDP is not required to reassemble the GRE packets; MX Series devices natively support GRE reassembly. With this configuration, you still use lt interfaces to stitch the CCCs between the remote branch offices and the VPLS routing instance as shown in Figure 150.

Copyright © 2016, Juniper Networks, Inc.

Figure 150: Central Office Ingress (Head-End) Configuration with an MX Series Device

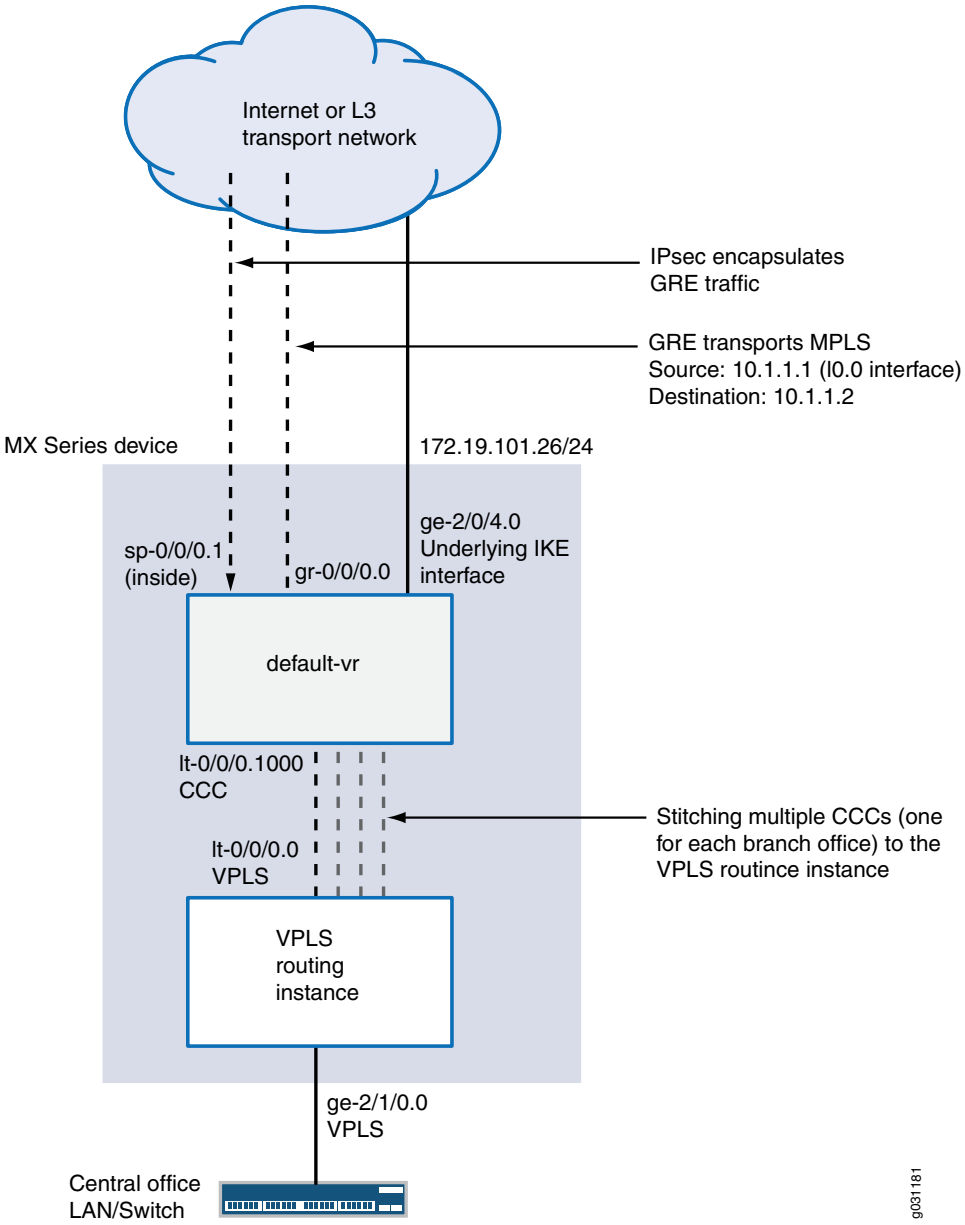

Configuration

In this example, we use SRX Series devices and the branch and ingress (head-end) sites will typically be connected to the Internet by Frame-Relay/T1-E1/xDSL/T3/E3 or even Ethernet. A provider MPLS network is not required.

- Configuring the SRX240 Device at the Branch Office on page 3878
- Configuring the SRX650 Device at the Central Office on page 3882

Copyright © 2016, Juniper Networks, Inc.

Configuring the SRX240 Device at the Branch Office

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces gr-0/0/0 description "GRE tunnel to SRX650”
set interfaces gr-0/0/0 unit 0 clear-dont-fragment-bit
set interfaces gr-0/0/0 unit 0 tunnel source 10.1.1.2
set interfaces gr-0/0/0 unit 0 tunnel destination 10.1.1.1
set interfaces gr-0/0/0 unit 0 tunnel allow-fragmentation
set interfaces gr-0/0/0 unit 0 family inet mtu 2000
set interfaces gr-0/0/0 unit 0 family inet filter input inet-packet-mode
set interfaces gr-0/0/0 unit 0 family mpls mtu 1900
set interfaces gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
set interfaces lt-0/0/0 unit 0 encapsulation frame-relay
set interfaces lt-0/0/0 unit 0 dlci 16
set interfaces lt-0/0/0 unit 0 peer-unit 1
set interfaces lt-0/0/0 unit 0 family inet
set interfaces lt-0/0/0 unit 0 description "Flow-vr Instance"
set interfaces lt-0/0/0 unit 1 encapsulation frame-relay
set interfaces lt-0/0/0 unit 1 dlci 16
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet filter input inet-packet-mode
set interfaces lt-0/0/0 unit 1 family inet address 10.1.1.2/32
set interfaces ge-0/0/1 encapsulation ethernet-ccc
set interfaces ge-0/0/1 unit 0 description "CCC Interface to customer LAN"
set interfaces ge-0/0/1 unit 0 family ccc filter input ccc-packet-mode
set interfaces ge-0/0/0 unit 0 family inet address 172.19.101.45/24
set interfaces lo0 unit 0 family inet address 10.2.1.2/32
set interfaces st0 unit 0 family inet
set routing-options static route 0.0.0.0/0 next-hop 172.19.101.1
set routing-options static route 10.1.1.1/32 next-hop lt-0/0/0.1
set routing-options static route 10.2.1.1/32 next-hop gr-0/0/0.0
set routing-options router-id 10.2.1.2
set protocols mpls interface gr-0/0/0.0
set protocols ldp interface gr-0/0/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 10.2.1.1 interface ge-0/0/1.0 virtual-circuit-id 1
set security ike policy SRX650 mode main
set security ike policy SRX650 proposal-set standard
set security ike policy SRX650 pre-shared-key ascii-text "$ABC123"
set security ike gateway SRX650 ike-policy SRX650
set security ike gateway SRX650 address 172.19.101.26
set security ike gateway SRX650 external-interface ge-0/0/0.0
set security ipsec policy SRX650 proposal-set standard
set security ipsec vpn SRX650 bind-interface st0.0
set security ipsec vpn SRX650 ike gateway SRX650
set security ipsec vpn SRX650 ike ipsec-policy SRX650
set security ipsec vpn SRX650 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces gr-0/0/0.0
set security zones security-zone untrust interfaces lo0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone untrust interfaces lt-0/0/0.1
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone vpn host-inbound-traffic system-services all
set security zones security-zone vpn host-inbound-traffic protocols all
set security zones security-zone vpn interfaces st0.0
set security zones security-zone trust-flow host-inbound-traffic system-services all
set security zones security-zone trust-flow host-inbound-traffic protocols all
set security zones security-zone trust-flow interfaces lt-0/0/0.0
set security policies from-zone trust-flow to-zone vpn policy gre match source-address
  any
set security policies from-zone trust-flow to-zone vpn policy gre match destination-address
  any
set security policies from-zone trust-flow to-zone vpn policy gre match application
  junos-gre
set security policies from-zone trust-flow to-zone vpn policy gre then permit
  application-services idp
set security idp idp-policy gre-reassembly rulebase-ips rule match-all match application
  junos-gre
set security idp idp-policy gre-reassembly rulebase-ips rule match-all then action
  ignore-connection
set security idp active-policy gre-reassembly
set firewall family inet filter inet-packet-mode term control-traffic from protocol tcp
set firewall family inet filter inet-packet-mode term control-traffic from port 22
set firewall family inet filter inet-packet-mode term control-traffic from port 80
set firewall family inet filter inet-packet-mode term control-traffic from port 8080
set firewall family inet filter inet-packet-mode term control-traffic then accept
set firewall family inet filter inet-packet-mode term packet-mode then packet-mode
set firewall family inet filter inet-packet-mode term packet-mode then accept
set firewall family mpls filter mpls-packet-mode term packet-mode then packet-mode
set firewall family mpls filter mpls-packet-mode term packet-mode then accept
set firewall family ccc filter ccc-packet-mode term all then packet-mode
set firewall family ccc filter ccc-packet-mode term all then accept
set routing-instances flow-vr instance-type virtual-router
set routing-instances flow-vr interface lt-0/0/0.0
set routing-instances flow-vr interface st0.0
set routing-instances flow-vr routing-options static route 10.1.1.1/32 next-hop st0.0
set routing-instances flow-vr routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the SRX240 at the branch office:

1. Configure a GRE tunnel to the central office.

```
[edit interfaces]
user@host# set gr-0/0/0 description "GRE tunnel to SRX650"
user@host# set gr-0/0/0 unit 0 clear-dont-fragment-bit
user@host# set gr-0/0/0 unit 0 tunnel source 10.1.1.2
user@host# set gr-0/0/0 unit 0 tunnel destination 10.1.1.1
user@host# set gr-0/0/0 unit 0 tunnel allow-fragmentation
user@host# set gr-0/0/0 unit 0 family inet mtu 2000
user@host# set gr-0/0/0 unit 0 family inet filter input inet-packet-mode
user@host# set gr-0/0/0 unit 0 family mpls mtu 1900
user@host# set gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
```

Copyright © 2016, Juniper Networks, Inc.

2. Create a logical interface that connects to the default routing instance.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation frame-relay
user@host# set lt-0/0/0 unit 0 dlci 16
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 0 family inet
user@host# set lt-0/0/0 unit 0 description "Flow-vr Instance"
```

3. Connect the logical tunnel interface to the flow mode virtual router.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation frame-relay
user@host# set lt-0/0/0 unit 1 dlci 16
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet filter input inet-packet-mode
user@host# set lt-0/0/0 unit 1 family inet address 10.1.1.2/32
```

4. Connect the CCC interface to the branch LAN.

```
[edit interfaces]
user@host# set ge-0/0/1 encapsulation ethernet-ccc
user@host# set ge-0/0/1 unit 0 description "CCC Interface to customer LAN"
user@host# set ge-0/0/1 unit 0 family ccc filter input ccc-packet-mode
```

5. Configure the interface bound to the default virtual router.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family inet address 172.19.101.45/24
```

6. Set the loopback interface to terminate the CCC connection.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.2.1.2/32
```

7. Bind the IPsec tunnel interface to the flow-mode virtual router.

```
[edit interfaces]
user@host# set st0 unit 0 family inet
```

8. Set a static route address, which will be the default gateway to the Internet.

```
[edit routing-options]
user@host# set static route 0.0.0.0/0 next-hop 172.19.101.1
```

9. Set a static route for the remote GRE tunnel endpoint.

```
[edit routing-options]
user@host# set static route 10.1.1.1/32 next-hop lt-0/0/0.1
```

10. Set a static route for the loopback interface of the SRX650 ingress (head-end) device.

```
[edit routing-options]
user@host# set static route 10.2.1.1/32 next-hop gr-0/0/0.0
```

11. Configure MPLS and the CCC using LDP as the label protocol.

```
[edit]
user@host# set routing-options router-id 10.2.1.2
user@host# set protocols mpls interface gr-0/0/0.0
user@host# set protocols ldp interface gr-0/0/0.0
user@host# set protocols ldp interface lo0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set protocols l2circuit neighbor 10.2.1.1 interface ge-0/0/1.0
  virtual-circuit-id 1
```

12. Configure the IPsec tunnel.

NOTE: The underlying IKE interface is not in the same routing instance as the tunnel interface.

```
[edit security]
user@host# set ike policy SRX650 mode main
user@host# set ike policy SRX650 proposal-set standard
user@host# set ike policy SRX650 pre-shared-key ascii-text "$ABC123"
user@host# set ike gateway SRX650 ike-policy SRX650
user@host# set ike gateway SRX650 address 172.19.101.26
user@host# set ike gateway SRX650 external-interface ge-0/0/0.0
user@host# set ipsec policy SRX650 proposal-set standard
user@host# set ipsec vpn SRX650 bind-interface st0.0
user@host# set ipsec vpn SRX650 ike gateway SRX650
user@host# set ipsec vpn SRX650 ike ipsec-policy SRX650
user@host# set ipsec vpn SRX650 establish-tunnels immediately
```

13. Configure security zones.

NOTE: In a production environment, host-inbound traffic should be restricted to only allow the necessary protocols and services.

```
[edit security]
user@host# set zones security-zone untrust host-inbound-traffic system-services
  all
user@host# set zones security-zone untrust host-inbound-traffic protocols all
user@host# set zones security-zone untrust interfaces gr-0/0/0.0
user@host# set zones security-zone untrust interfaces lo0.0
user@host# set zones security-zone untrust interfaces lt-0/0/0.1
user@host# set zones security-zone untrust interfaces ge-0/0/0.0
user@host# set zones security-zone vpn host-inbound-traffic system-services all
user@host# set zones security-zone vpn host-inbound-traffic protocols all
user@host# set zones security-zone vpn interfaces st0.0
user@host# set zones security-zone trust-flow host-inbound-traffic system-services
  all
user@host# set zones security-zone trust-flow host-inbound-traffic protocols all
user@host# set zones security-zone trust-flow interfaces lt-0/0/0.0
```

14. Configure IDP.

```
[edit security]
user@host# set policies from-zone trust-flow to-zone vpn policy gre match
  source-address any
user@host# set policies from-zone trust-flow to-zone vpn policy gre match
  destination-address any
user@host# set policies from-zone trust-flow to-zone vpn policy gre match
  application junos-gre
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policies from-zone trust-flow to-zone vpn policy gre then permit
  application-services idp
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-all match
  application junos-gre
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-all then
  action ignore-connection
user@host# set idp active-policy gre-reassembly
```

15. Configure packet-mode filters.

```
[edit firewall]
user@host# set family inet filter inet-packet-mode term control-traffic from protocol
  tcp
user@host# set family inet filter inet-packet-mode term control-traffic from port
  22
user@host# set family inet filter inet-packet-mode term control-traffic from port
  80
user@host# set family inet filter inet-packet-mode term control-traffic from port
  8080
user@host# set family inet filter inet-packet-mode term control-traffic then accept
user@host# set family inet filter inet-packet-mode term packet-mode then
  packet-mode
user@host# set family inet filter inet-packet-mode term packet-mode then accept
user@host# set family mpls filter mpls-packet-mode term packet-mode then
  packet-mode
user@host# set family mpls filter mpls-packet-mode term packet-mode then accept
user@host# set family ccc filter ccc-packet-mode term all then packet-mode
user@host# set family ccc filter ccc-packet-mode term all then accept
```

16. Configure the flow-mode virtual router.

```
[edit routing-instances]]
user@host# set flow-vr instance-type virtual-router
user@host# set flow-vr interface lt-0/0/0.0
user@host# set flow-vr interface st0.0
user@host# set flow-vr routing-options static route 10.1.1.1/32 next-hop st0.0
user@host# set flow-vr routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.0
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the SRX650 Device at the Central Office

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 172.19.101.26/24
set interfaces gr-0/0/0 unit 0 clear-dont-fragment-bit
set interfaces gr-0/0/0 unit 0 tunnel source 10.1.1.1
set interfaces gr-0/0/0 unit 0 tunnel destination 10.1.1.2
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces gr-0/0/0 unit 0 tunnel allow-fragmentation
set interfaces gr-0/0/0 unit 0 family inet mtu 1500
set interfaces gr-0/0/0 unit 0 family inet filter input inet-packet-mode
set interfaces gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
set interfaces lt-0/0/0 unit 0 description "VPLS hub port - Interconnect for CCC to
  SRX240"
set interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 0 peer-unit 1000
set interfaces lt-0/0/0 unit 1000 description "Stitch to VPLS for CCC to SRX240"
set interfaces lt-0/0/0 unit 1000 encapsulation ethernet-ccc
set interfaces lt-0/0/0 unit 1000 peer-unit 0
set interfaces lt-0/0/0 unit 1000 family ccc filter input ccc-packet-mode
set interfaces lt-0/0/0 unit 2000 encapsulation frame-relay
set interfaces lt-0/0/0 unit 2000 dlci 1
set interfaces lt-0/0/0 unit 2000 peer-unit 2001
set interfaces lt-0/0/0 unit 2000 family inet
set interfaces lt-0/0/0 unit 2001 encapsulation frame-relay
set interfaces lt-0/0/0 unit 2001 dlci 1
set interfaces lt-0/0/0 unit 2001 peer-unit 2000
set interfaces lt-0/0/0 unit 2001 family inet filter input inet-packet-mode
set interfaces lt-0/0/0 unit 2001 family inet address 10.1.1.1/32
set interfaces ge-0/0/1 unit 0
set interfaces ge-0/0/1 encapsulation ethernet-vpls
set interfaces lo0 unit 0 family inet address 10.2.1.1/32
set interfaces st0 unit 0 family inet
set routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.2001
set routing-options static route 10.2.1.2/32 next-hop gr-0/0/0.0
set protocols mpls interface gr-0/0/0.0
set protocols ldp interface gr-0/0/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 10.2.1.2 interface lt-0/0/0.1000 virtual-circuit-id 1
set security ike policy SRX mode main
set security ike policy SRX proposal-set standard
set security ike policy SRX pre-shared-key ascii-text "$ABC123"
set security ike gateway SRX240-1 ike-policy SRX
set security ike gateway SRX240-1 address 172.19.101.45
set security ike gateway SRX240-1 external-interface ge-0/0/0.0
set security ipsec policy SRX proposal-set standard
set security ipsec vpn SRX240-1 bind-interface st0.0
set security ipsec vpn SRX240-1 ike gateway SRX240-1
set security ipsec vpn SRX240-1 ike ipsec-policy SRX
set security ipsec vpn SRX240-1 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces lt-0/0/0.2001
set security zones security-zone untrust interfaces gr-0/0/0.0
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone vpn host-inbound-traffic system-services all
set security zones security-zone vpn host-inbound-traffic protocols all
set security zones security-zone vpn interfaces st0.0
set security zones security-zone trust-flow host-inbound-traffic system-services all
set security zones security-zone trust-flow host-inbound-traffic protocols all
set security zones security-zone trust-flow interfaces lt-0/0/0.2000
set security policies from-zone trust-flow to-zone vpn policy gre match source-address
  any
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust-flow to-zone vpn policy gre match destination-address
  any
set security policies from-zone trust-flow to-zone vpn policy gre match application
  junos-gre
set security policies from-zone trust-flow to-zone vpn policy gre then permit
  application-services idp
set security policies from-zone vpn to-zone trust-flow policy gre match source-address
  any
set security policies from-zone vpn to-zone trust-flow policy gre match destination-address
  any
set security policies from-zone vpn to-zone trust-flow policy gre match application
  junos-gre
set security policies from-zone vpn to-zone trust-flow policy gre then permit
  application-services idp
set security idp idp-policy gre-reassembly rulebase-ips rule match-gre match application
  junos-gre
set security idp idp-policy gre-reassembly rulebase-ips rule match-gre then action
  ignore-connection
set security idp active-policy gre-reassembly
set firewall family inet filter inet-packet-mode term control-traffic from protocol tcp
set firewall family inet filter inet-packet-mode term control-traffic from port 22
set firewall family inet filter inet-packet-mode term control-traffic from port 80
set firewall family inet filter inet-packet-mode term control-traffic from port 8080
set firewall family inet filter inet-packet-mode term control-traffic then accept
set firewall family inet filter inet-packet-mode term packet-mode then packet-mode
set firewall family inet filter inet-packet-mode term packet-mode then accept
set firewall family mpls filter mpls-packet-mode term packet-mode then packet-mode
set firewall family mpls filter mpls-packet-mode term packet-mode then accept
set firewall family ccc filter ccc-packet-mode term all then packet-mode
set firewall family ccc filter ccc-packet-mode term all then accept
set routing-instances flow-vr instance-type virtual-router
set routing-instances flow-vr interface lt-0/0/0.2000
set routing-instances flow-vr interface st0.0
set routing-instances flow-vr routing-options static route 10.1.1.1/32 next-hop
  lt-0/0/0.2000
set routing-instances flow-vr routing-options static route 10.1.1.2/32 next-hop st0.0
set routing-instances vpls-hub instance-type vpls
set routing-instances vpls-hub interface lt-0/0/0.0
set routing-instances vpls-hub interface ge-0/0/1.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the ingress (head-end) SRX650 device at the central office:

1. Configure the interface bound to the default virtual router.

   ```
   [edit interfaces]
   user@host# set ge-0/0/0 unit 0 family inet address 172.19.101.26/24
   ```

2. Create the GRE tunnel from the SRX650 to the SRX240 device.

Copyright © 2016, Juniper Networks, Inc.

NOTE: As the network expands to include multiple branch offices, you will need to add a similar GRE tunnel configuration on the SRX650 device (head-end) along with a corresponding IPsec configuration to connect to each additional branch device (SRX240).

```
[edit interfaces]
user@host# set gr-0/0/0 unit 0 clear-dont-fragment-bit
user@host# set gr-0/0/0 unit 0 tunnel source 10.1.1.1
user@host# set gr-0/0/0 unit 0 tunnel destination 10.1.1.2
user@host# set gr-0/0/0 unit 0 tunnel allow-fragmentation
user@host# set gr-0/0/0 unit 0 family inet mtu 1500
user@host# set gr-0/0/0 unit 0 family inet filter input inet-packet-mode
user@host# set gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
```

3. Configure a logical tunnel interface to stitch the CCC connection to the VPLS instance.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 0 description "VPLS hub port - Interconnect for CCC
   to SRX240"
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1000
```

4. Set unit 1000 to terminate the CCC connection.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1000 description "Stitch to VPLS for CCC to SRX240"
user@host# set lt-0/0/0 unit 1000 encapsulation ethernet-ccc
user@host# set lt-0/0/0 unit 1000 peer-unit 0
user@host# set lt-0/0/0 unit 1000 family ccc filter input ccc-packet-mode
```

5. Configure the logical tunnel interface.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 2000 encapsulation frame-relay
user@host# set lt-0/0/0 unit 2000 dlci 1
user@host# set lt-0/0/0 unit 2000 peer-unit 2001
user@host# set lt-0/0/0 unit 2000 family inet
```

6. Bind the logical tunnel interface to the default virtual router.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 2001 encapsulation frame-relay
user@host# set lt-0/0/0 unit 2001 dlci 1
user@host# set lt-0/0/0 unit 2001 peer-unit 2000
user@host# set lt-0/0/0 unit 2001 family inet filter input inet-packet-mode
user@host# set lt-0/0/0 unit 2001 family inet address 10.1.1.1/32
```

7. Set the interface to the central office LAN network.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0
user@host# set ge-0/0/1 encapsulation ethernet-vpls
```

8. Set the loopback interface to terminate the CCC connections to each branch device.

```
[edit interfaces]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set lo0 unit 0 family inet address 10.2.1.1/32
```

9. Bind the IPsec interface to the flow-mode virtual router.

```
[edit interfaces]
user@host# set st0 unit 0 family inet
```

10. Set a static route for the remote GRE tunnel endpoint.

```
[edit routing-options]
user@host# set static route 10.1.1.2/32 next-hop lt-0/0/0.2001
```

11. Set a static route for the loopback interface of the branch device.

```
[edit]
user@host# set routing-options static route 10.2.1.2/32 next-hop gr-0/0/0.0
```

12. Configure MPLS and CCC using LDP as the label protocol.

```
[edit protocols]
user@host# set mpls interface gr-0/0/0.0
user@host# set ldp interface gr-0/0/0.0
user@host# set ldp interface lo0.0
user@host# set l2circuit neighbor 10.2.1.2 interface lt-0/0/0.1000 virtual-circuit-id
  1
```

13. Configure the IPsec tunnel.

NOTE: The underlying IKE interface is not in the same routing instance as the tunnel interface.

```
[edit security]
user@host# set ike policy SRX mode main
user@host# set ike policy SRX proposal-set standard
user@host# set ike policy SRX pre-shared-key ascii-text "$ABC123"
user@host# set ike gateway SRX240-1 ike-policy SRX
user@host# set ike gateway SRX240-1 address 172.19.101.45
user@host# set ike gateway SRX240-1 external-interface ge-0/0/0.0
user@host# set ipsec policy SRX proposal-set standard
user@host# set ipsec vpn SRX240-1 bind-interface st0.0
user@host# set ipsec vpn SRX240-1 ike gateway SRX240-1
user@host# set ipsec vpn SRX240-1 ike ipsec-policy SRX
user@host# set ipsec vpn SRX240-1 establish-tunnels immediately
```

14. Configure security zones.

NOTE: In a production environment, restrict host-inbound traffic to only the necessary protocols and services.

```
[edit security]
user@host# set zones security-zone untrust host-inbound-traffic system-services
  all
user@host# set zones security-zone untrust host-inbound-traffic protocols all
user@host# set zones security-zone untrust interfaces lo0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set zones security-zone untrust interfaces lt-0/0/0.2001
user@host# set zones security-zone untrust interfaces gr-0/0/0.0
user@host# set zones security-zone untrust interfaces ge-0/0/0.0
user@host# set zones security-zone vpn host-inbound-traffic system-services all
user@host# set zones security-zone vpn host-inbound-traffic protocols all
user@host# set zones security-zone vpn interfaces st0.0
user@host# set zones security-zone trust-flow host-inbound-traffic system-services
  all
user@host# set zones security-zone trust-flow host-inbound-traffic protocols all
user@host# set zones security-zone trust-flow interfaces lt-0/0/0.2000
```

15. Configure IDP.

```
[edit security]
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match
  source-address any
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match
  destination-address any
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match
  application junos-gre
user@host# set policies from-zone trust-flow to-zone vpn policy GRE then permit
  application-services idp
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match
  source-address any
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match
  destination-address any
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match
  application junos-gre
user@host# set policies from-zone vpn to-zone trust-flow policy GRE then permit
  application-services idp
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-gre match
  application junos-gre
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-gre then
  action ignore-connection
user@host# set idp active-policy gre-reassembly
```

16. Configure packet-mode filters.

```
[edit firewall]
user@host# set family inet filter inet-packet-mode term control-traffic from protocol
  tcp
user@host# set family inet filter inet-packet-mode term control-traffic from port
  22
user@host# set family inet filter inet-packet-mode term control-traffic from port
  80
user@host# set family inet filter inet-packet-mode term control-traffic from port
  8080
user@host# set family inet filter inet-packet-mode term control-traffic then accept
user@host# set family inet filter inet-packet-mode term packet-mode then
  packet-mode
user@host# set family inet filter inet-packet-mode term packet-mode then accept
user@host# set family mpls filter mpls-packet-mode term packet-mode then
  packet-mode
user@host# set family mpls filter mpls-packet-mode term packet-mode then accept
user@host# set family ccc filter ccc-packet-mode term all then packet-mode
user@host#set family ccc filter ccc-packet-mode term all then accept
```

Copyright © 2016, Juniper Networks, Inc.

17. Configure the flow-mode virtual router.

```
[edit routing-instances]
user@host# set flow-vr instance-type virtual-router
user@host# set flow-vr interface lt-0/0/0.2000
user@host# set flow-vr interface st0.0
user@host# set flow-vr routing-options static route 10.1.1.1/32 next-hop
  lt-0/0/0.2000
user@host# set flow-vr routing-options static route 10.1.1.2/32 next-hop st0.0
```

18. Configure the VPLS instance.

```
[edit routing-instances]
user@host# set vpls-hub instance-type vpls
user@host# set vpls-hub interface lt-0/0/0.0
user@host# set vpls-hub interface ge-0/0/1.0
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Interfaces on page 3888
- Verifying an IPsec tunnel on page 3888
- Verifying GRE on page 3888
- Verifying the CCC/L2 circuit. on page 3889
- Verifying that LDP sessions are working. on page 3889

Verifying Interfaces

Purpose Verify that the interfaces are configured properly on each device in the VPLS network.

Action From configuration mode, enter **show interfaces** and verify that the IP addressing is correct for each interface, including logical tunnel (lt), loopback (lo), GRE (gr), IPsec tunnel st0, and GE interfaces.

Verifying an IPsec tunnel

Purpose Verify that an IPsec tunnel is working.

Action From operational mode, enter the **show security ipsec security associations** and the **show security ipsec statistics** command.

Verifying GRE

Purpose Verify that GRE is working.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security flow session protocol gre** command. You can also do a ping between loopback addresses.

Verifying the CCC/L2 circuit.

Purpose Verify that the CCC/L2 circuit is working.

Action From operational mode, enter the **show connections** command.

Verifying that LDP sessions are working.

Purpose Verify that LDP sessions are being created between devices.

Action From operational mode, enter the **show interfaces gr-0/0/0 detail** command.

Related Documentation

- VPLS Overview on page 3841
- Understanding VPLS Interfaces on page 3849
- Understanding Selective Stateless Packet-Based Services on page 1804
- MPLS Overview on page 3743

Example: Configuring VPLS with BGP Signaling

This example shows how to configure VPLS with BGP signaling between two devices.

- Requirements on page 3889
- Overview on page 3889
- Configuration on page 3890
- Verification on page 3900

Requirements

Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.

Overview

This example shows a minimum configuration for PE devices and CE devices to create a VPLS network with BGP signaling. The topology consists of two PE devices and two CE devices. In this example, you configure a VPLS routing instance vpls-instance between two PE devices, PE1 and PE2. You also configure the CE1 and CE2 devices that use Ethernet-based interfaces to connect VLAN 600 to their local PE devices. On the CE1 device, configure the Fast Ethernet interface that connects to the PE1 device. The VLAN identifier and IP address must match those of the CE2 device.

Figure 151 shows the topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 151: Configuring VPLS with BGP Signaling

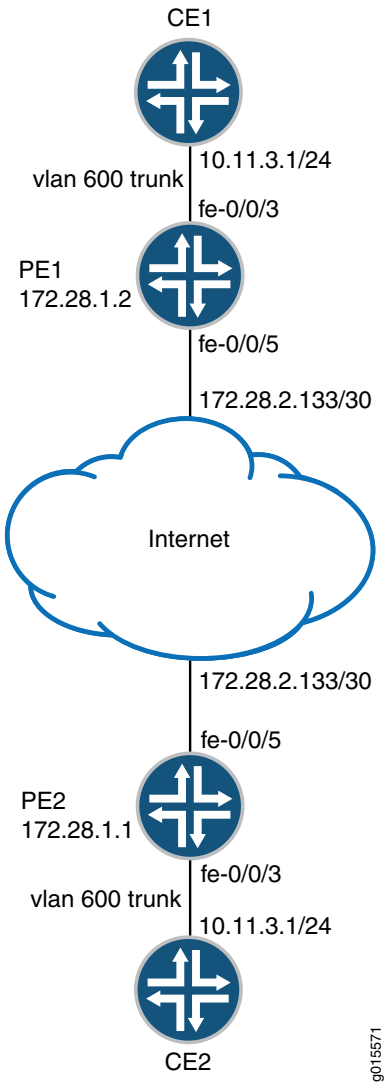

Configuration

- Configuring the CE1 Device on page 3890
- Configuring the PE1 Device on page 3891
- Configuring the PE2 Device on page 3895
- Configuring the CE2 Device on page 3899

Configuring the CE1 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/3 vlan-tagging
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family inet address 10.11.3.1/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Enable VLAN tagging on the VPLS interface.

   ```
   [edit interfaces fe-0/0/3]
   user@host# set vlan-tagging
   ```

2. Configure the VLAN ID on the logical interface.

   ```
   [edit interfaces fe-0/0/3 unit 0]
   user@host# set vlan-id 600
   ```

3. Configure the VPLS family on the logical interface.

   ```
   [edit interfaces fe-0/0/3 unit 0]
   user@host# set family inet address 10.11.3.1/24
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
fe-0/0/3 {
    vlan-tagging;
    unit 0 {
        vlan-id 600;
        family inet {
            address 10.11.3.1/24;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the PE1 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system host-name PE1
set interfaces fe-0/0/3 description "CE1 on PE1"
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family vpls
set interfaces fe-0/0/5 vlan-tagging
set interfaces fe-0/0/5 unit 37 vlan-id 37
set interfaces fe-0/0/5 unit 37 family inet address 172.28.2.133/30
set interfaces fe-0/0/5 unit 37 family mpls
```

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces lo0 unit 0 family inet address 172.28.1.2/32
set routing-options router-id 172.28.1.2
set routing-options autonomous-system 65512
set protocols rsvp interface fe-0/0/5.37
set protocols mpls label-switched-path pe1-to-pe2 to 172.28.1.1
set protocols mpls interface fe-0/0/5.37
set protocols mpls interface lo0.0
set protocols bgp group vpls-peering type internal
set protocols bgp group vpls-peering local-address 172.28.1.2
set protocols bgp group vpls-peering family l2vpn signaling
set protocols bgp group vpls-peering neighbor 172.28.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/5.37
set routing-instances vpls-instance description "Routing instance from VPLS routing"
set routing-instances vpls-instance instance-type vpls
set routing-instances vpls-instance interface fe-0/0/3.0
set routing-instances vpls-instance route-distinguisher 172.28.1.2:1
set routing-instances vpls-instance vrf-target target:65512:1
set routing-instances vpls-instance protocols vpls site-range 10
set routing-instances vpls-instance protocols vpls no-tunnel-services site site10
  automatic-site-id
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PE1:

1. Configure the hostname for the PE1 device.

```
[edit ]
user@host# set system host-name PE1
```

2. Configure VPLS VLAN encapsulation on the VPLS PE1 device.

```
[edit interfaces]
user@host# set fe-0/0/3 description "CE1 on PE1"
user@host# set fe-0/0/3 vlan-tagging
user@host# set fe-0/0/3 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 vlan-id 600
user@host# set fe-0/0/3 unit 0 family vpls
```

3. Configure the routing interface on the VPLS PE1 device.

```
[edit interfaces]
user@host# set fe-0/0/5 vlan-tagging
user@host# set fe-0/0/5 unit 37 vlan-id 37
user@host# set fe-0/0/5 unit 37 family inet address 172.28.2.133/30
user@host# set fe-0/0/5 unit 37 family mpls
user@host# set lo0 unit 0 family inet address 172.28.1.2/32
```

NOTE: For this example, it is optional to configure VLAN tagging. Remove the VLAN tagging configuration on the physical interfaces if you do not plan to configure VLAN tagging.

Copyright © 2016, Juniper Networks, Inc.

4. Configure the routing options on the VPLS PE1 device.

```
[edit routing-options]
user@host# set router-id 172.28.1.2
user@host# set autonomous-system 65512
```

5. Configure RSVP on the VPLS PE1 device.

```
[edit protocols]
user@host# set rsvp interface fe-0/0/5.37
```

6. Configure MPLS on the VPLS PE1 device.

```
[edit protocols]
user@host# set mpls label-switched-path pe1-to-pe2 to 172.28.1.1
user@host# set mpls interface fe-0/0/5.37
user@host# set mpls interface lo0.0
```

7. Configure BGP on the VPLS PE1 device.

```
[edit protocols]
user@host# set bgp group vpls-peering type internal
user@host# set bgp group vpls-peering local-address 172.28.1.2
user@host# set bgp group vpls-peering family l2vpn signaling
user@host# set bgp group vpls-peering neighbor 172.28.1.1
```

8. (Optional) Configure OSPF on the VPLS PE1 device.

NOTE: For this example, it is optional to configure OFPF. You must configure OSPF only in cases where two PE devices are not connected directly.

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
user@host# set ospf area 0.0.0.0 interface fe-0/0/5.37
```

9. Create a VPLS routing instance.

```
[edit ]
user@host# set routing-instances vpls-instance
```

10. Configure a VPLS routing instance.

```
[edit routing-instances vpls-instance]
user@host# set description "Routing instance from VPLS routing"
user@host# set instance-type vpls
user@host# set interface fe-0/0/3.0
user@host# set route-distinguisher 172.28.1.2:1
user@host# set vrf-target target:65512:1
user@host# set protocols vpls site-range 10
user@host# set protocols vpls no-tunnel-services site site10 automatic-site-id
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show system
host-name PE1;

[edit]
user@host# show interfaces
fe-0/0/5 {
   vlan-tagging;
   unit 37 {
      vlan-id 37;
      family inet {
         address 172.28.2.133/30;
      }
      family mpls;
   }
}
fe-0/0/3 {
   description "CE1 on PE1";
   vlan-tagging;
   encapsulation vlan-vpls;
   unit 0 {
      encapsulation vlan-vpls;
      vlan-id 600;
      family vpls;
   }
}
lo0 {
   unit 0 {
      family inet {
         address 172.28.1.2/32;
      }
   }
}

[edit]
user@host# show routing-options
router-id 172.28.1.2;
autonomous-system 65512;

[edit]
user@host# show protocols
rsvp {
   interface fe-0/0/5.37;
}
mpls {
   label-switched-path pe1-to-pe2 {
      to 172.28.1.1;
   }
   interface fe-0/0/5.37;
   interface lo0.0;
}
bgp {
   group vpls-peering {
      type internal;
      local-address 172.28.1.2;
      family l2vpn {
         signaling;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
      neighbor 172.28.1.1;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-0/0/5.37;
    }
  }

  [edit]
  user@host# show routing-instances
  vpls-instance {
    description "Routing instance from VPLS routing";
    instance-type vpls;
    interface fe-0/0/3.0;
    route-distinguisher 172.28.1.2:1;
    vrf-target target:65512:1;
    protocols {
      vpls {
        site-range 10;
        no-tunnel-services;
        site site10 {
          automatic-site-id;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the PE2 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system host-name PE2
set interfaces fe-0/0/3 description "CE2 on PE2"
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family vpls
set interfaces fe-0/0/5 vlan-tagging
set interfaces fe-0/0/5 unit 37 vlan-id 37
set interfaces fe-0/0/5 unit 37 family inet address 172.28.2.133/30
set interfaces fe-0/0/5 unit 37 family mpls
set interfaces lo0 unit 0 family inet address 172.28.1.1/32
set routing-options router-id 172.28.1.1
set routing-options autonomous-system 65512
set protocols rsvp interface fe-0/0/5.37
set protocols mpls label-switched-path pe2-to-pe1 to 172.28.1.2
```

Copyright © 2016, Juniper Networks, Inc.

```
set protocols mpls interface fe-0/0/5.37
set protocols mpls interface lo0.0
set protocols bgp group vpls-peering type internal
set protocols bgp group vpls-peering local-address 172.28.1.1
set protocols bgp group vpls-peering family l2vpn signaling
set protocols bgp group vpls-peering neighbor 172.28.1.2
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/5.37
set routing-instances vpls-instance description "Routing instance for VPLS routing"
set routing-instances vpls-instance instance-type vpls
set routing-instances vpls-instance interface fe-0/0/3.0
set routing-instances vpls-instance route-distinguisher 172.28.1.1:1
set routing-instances vpls-instance vrf-target target:65512:1
set routing-instances vpls-instance protocols vpls site-range 10
set routing-instances vpls-instance protocols vpls no-tunnel-services site site11
  automatic-site-id
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PE2:

1. Configure the hostname for the device.

   ```
   [edit ]
   user@host# set system host-name PE2
   ```

2. Configure VPLS VLAN encapsulation on the VPLS PE2 device.

   ```
   [edit interfaces]
   user@host# set fe-0/0/3 description "CE2 on PE2"
   user@host# set fe-0/0/3 vlan-tagging
   user@host# set fe-0/0/3 encapsulation vlan-vpls
   user@host# set fe-0/0/3 unit 0 encapsulation vlan-vpls
   user@host# set fe-0/0/3 unit 0 vlan-id 600
   user@host# set fe-0/0/3 unit 0 family vpls
   ```

3. Configure the routing interface on the VPLS PE2 device.

   ```
   [edit interfaces]
   user@host# set fe-0/0/5 vlan-tagging
   user@host# set fe-0/0/5 unit 37 vlan-id 37
   user@host# set fe-0/0/5 unit 37 family inet address 172.28.2.133/30
   user@host# set fe-0/0/5 unit 37 family mpls
   user@host# set lo0 unit 0 family inet address 172.28.1.1/32
   ```

   NOTE: For this example, it is optional to configure VLAN tagging. Remove the VLAN tagging configuration on the physical interfaces if you do not plan to configure VLAN tagging.

4. Configure the routing options on the VPLS PE2 device.

   ```
   [edit routing-options]
   user@host# set router-id 172.28.1.1
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set autonomous-system 65512
```

5. Configure RSVP on the VPLS PE2 device.

```
[edit protocols]
user@host# set rsvp interface fe-0/0/5.37
```

6. Configure MPLS on the VPLS PE2 device.

```
[edit protocols]
user@host# set mpls label-switched-path pe2-to-pe1 to 172.28.1.2
user@host# set mpls interface fe-0/0/5.37
user@host# set mpls interface lo0.0
```

7. Configure BGP on the VPLS PE2 device.

```
[edit protocols]
user@host# set bgp group vpls-peering type internal
user@host# set bgp group vpls-peering local-address 172.28.1.1
user@host# set bgp group vpls-peering family l2vpn signaling
user@host# set bgp group vpls-peering neighbor 172.28.1.2
```

8. (Optional) Configure OSPF on the VPLS PE2 device.

NOTE: For this example, it is optional to configure OFPF. You must configure OSPF only in cases where two PE devices are not connected directly.

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
user@host# set ospf area 0.0.0.0 interface fe-0/0/5.37
```

9. Create a VPLS routing instance.

```
[edit ]
user@host# set routing-instances vpls-instance
```

10. Configure a VPLS routing instance.

```
[edit routing-instances vpls-instance]
user@host# set description "Routing instance for VPLS routing"
user@host# set instance-type vpls
user@host# set interface fe-0/0/3.0
user@host# set route-distinguisher 172.28.1.1:1
user@host# set vrf-target target:65512:1
user@host# set protocols vpls site-range 10
user@host# set protocols vpls no-tunnel-services site site11 automatic-site-id
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system
host-name PE2;

[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show interfaces
fe-0/0/5 {
   vlan-tagging;
   unit 37 {
      vlan-id 37;
      family inet {
         address 172.28.2.133/30;
      }
      family mpls;
   }
}
fe-0/0/3 {
   description "CE2 on PE2";
   vlan-tagging;
   encapsulation vlan-vpls;
   unit 0 {
      encapsulation vlan-vpls;
      vlan-id 600;
      family vpls;
   }
}
lo0 {
   unit 0 {
      family inet {
         address 172.28.1.1/32;
      }
   }
}

[edit]
user@host# show routing-options
router-id 172.28.1.1;
autonomous-system 65512;

[edit]
user@host# show protocols
rsvp {
   interface fe-0/0/5.37;
}
mpls {
   label-switched-path pe2-to-pe1 {
      to 172.28.1.2;
   }
   interface fe-0/0/5.37;
   interface lo0.0;
}
bgp {
   group vpls-peering {
      type internal;
      local-address 172.28.1.1;
      family l2vpn {
         signaling;
      }
      neighbor 172.28.1.1;
   }
}
ospf {
```


Copyright © 2016, Juniper Networks, Inc.

```
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-0/0/5.37;
    }
}
```

```
[edit]
user@host# show routing-instances
vpls-instance {
  description "Routing instance from VPLS routing";
  instance-type vpls;
  interface fe-0/0/3.0;
  route-distinguisher 172.28.1.1:1;
  vrf-target target:65512:1;
  protocols {
    vpls {
      site-range 10;
      no-tunnel-services;
      site site11 {
        automatic-site-id;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the CE2 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family inet address 10.11.3.2/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Enable VLAN tagging on the VPLS interface.

   ```
   [edit interfaces fe-0/0/3]
   user@host# set vlan-tagging
   ```

2. Configure the VLAN ID on the logical interface.

   ```
   [edit interfaces fe-0/0/3 unit 0]
   user@host# set vlan-id 600
   ```

3. Configure the VPLS family on the logical interface.

   ```
   [edit interfaces fe-0/0/3 unit 0]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set family inet address 10.11.3.2/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
fe-0/0/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 600;
    family inet {
      address 10.11.3.2/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

NOTE: If VLAN trunking is not needed between the CE devices, remove the configuration on VLAN tagging on the interfaces connecting the CE and PE devices. Also, use ethernet-VPLS-encapsulation instead of vlan-vpls on the CE facing interfaces of the PE devices.

Verification

Confirm that the configuration is working properly.

- Verifying Interfaces on page 3900
- Verifying Routing Information on page 3900
- Verifying VPLS Information on page 3901
- Verifying Automatic Site Identifier Generation on page 3901

Verifying Interfaces

Purpose Verify that the interfaces are configured correctly.

Action From operational mode, enter the **show interfaces terse** command.

Verifying Routing Information

Purpose Verify that the routing information is configured correctly.

Action From operational mode, enter the following commands:

- **show route forwarding-table family mpls**
- **show route forwarding-table family vpls (destination | extensive | matching | table)**
- **show route instance (detail)**

Copyright © 2016, Juniper Networks, Inc.

Verifying VPLS Information

Purpose Verify that the VPLS is configured correctly.

Action From operational mode, enter the following commands:

- **show system statistics vpls**
- **show vpls connections**
- **show vpls statistics**

Verifying Automatic Site Identifier Generation

Purpose Verify that the automatic site identifier has been generated.

Action From operational mode, enter the **show vpls connections** command.

```
[edit]
user@host# show vpls connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down NP -- interface hardware not present
CM -- control-word mismatch -> -- only outbound connection is up
CN -- circuit not provisioned <- -- only inbound connection is up
OR -- out of range Up -- operational
OL -- no outgoing label Dn -- down
LD -- local site signaled down CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch MI -- Mesh-Group ID not available
BK -- Backup connection ST -- Standby connection
PF -- Profile parse failure PB -- Profile busy
RS -- remote site standby SN -- Static Neighbor
VM -- VLAN ID mismatch
Legend for interface status
Up -- operational
Dn -- down
Instance: customer2
Local site: airwalk (2)
connection-site Type St Time last up # Up trans
4 rmt Up Mar 1 03:26:21 2012 1
Remote PE: 200.100.100.2, Negotiated control-word: No
Incoming label: 262148, Outgoing label: 262146
Local interface: lsi.1048838, Status: Up, Encapsulation: VPLS
Description: Intf - vpls customer2 local site 2 remote site 4
Instance: customer4
Local site: airwalk (6)
connection-site Type St Time last up # Up trans
8 rmt Up Feb 21 03:27:33 2012 1
```

Copyright © 2016, Juniper Networks, Inc.

```
Remote PE: 200.200.200.2, Negotiated control-word: No
Incoming label: 262160, Outgoing label: 262174
Local interface: lsi.1048836, Status: Up, Encapsulation: VPLS
Description: Intf - vpls customer4 local site 6 remote site 8
```

Related Documentation

- VPLS Overview on page 3841
- Understanding VPLS Interfaces on page 3849
- MPLS Overview on page 3743

Example: Configuring BGP on the VPLS PE Router

This example shows how to configure BGP on the VPLS PE router.

- Requirements on page 3902
- Overview on page 3902
- Configuration on page 3903
- Verification on page 3903

Requirements

Before you begin:

- See "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE. See "Example: Configuring RSVP on the VPLS PE Router" on page 3868. Then configure MPLS LSPs on the PE routers. See "Example: Configuring MPLS on the VPLS PE Router" on page 3869. Alternatively, configure LDP on the PE routers. See "Example: Configuring LDP on the VPLS PE Router" on page 3871.
- Configure routing options on the PE router. See "Example: Configuring Routing Options on the VPLS PE Router" on page 3903.

Overview

In this example, you configure an internal BGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. The PE routers use this information to determine which labels to use for traffic destined for remote sites.

Copyright © 2016, Juniper Networks, Inc.

NOTE: On all high-end SRX Series devices, BGP-based virtual private LAN service (VPLS) works on child ports and physical interfaces, but not over aggregated Ethernet (ae) interfaces.

Configuration

Step-by-Step Procedure

To configure BGP on the VPLS PE router:

1. Configure the BGP internal group on the VPLS PE router.

   ```
   [edit ]
   user@host# set protocols bgp group ibgp type internal local-address 10.255.7.168
     neighbor 10.255.7.164
   ```

2. Configure the BGP family L2vpn and specify NLRI signaling.

   ```
   [edit ]
   user@host# set protocols bgp family L2 VPN signaling
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring Routing Options on the VPLS PE Router

This example shows how to configure the routing options on the VPLS PE router.

- Requirements on page 3903
- Overview on page 3904
- Configuration on page 3904
- Verification on page 3904

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See"Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.

Copyright © 2016, Juniper Networks, Inc.

- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867
- Configure RSVP-TE, see "Example: Configuring RSVP on the VPLS PE Router" on page 3868 and then MPLS LSPs on the PE routers, see "Example: Configuring MPLS on the VPLS PE Router" on page 3869. Alternatively configure LDP on the PE routers, see "Example: Configuring LDP on the VPLS PE Router" on page 3871.

Overview

This example describes how to specify the router ID and the AS number for each router involved in the VPLS . In this example, the routers PE1 and PE2 use the same AS number (100).

Configuration

Step-by-Step Procedure

To configure the routing options on the VPLS PE router:

1. Configure the router ID on the VPLS PE router.

   ```
   [edit]
   user@host# set routing-options router-id 10.255.7.168
   ```

2. Configure the AS number on the VPLS PE router.

   ```
   [edit]
   user@host# set routing-options autonomous-system 100
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show routing-options** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 172

Configuring Encapsulation

- Understanding VPLS VLAN Encapsulation on page 3905
- Understanding VPLS VLAN Encapsulation on a Logical Interface on page 3906
- Example: Configuring VPLS VLAN Encapsulation on page 3906
- Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces on page 3909
- Example: Configuring Extended VLAN VPLS Encapsulation on page 3910

Understanding VPLS VLAN Encapsulation

Gigabit Ethernet IQ, Gigabit Ethernet PIMs with small form-factor pluggable optics (SFPs), SRX Series devices with Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use flexible Ethernet services, VLAN virtual private LAN service (VPLS) encapsulation.

NOTE: **VLAN encapsulation is not supported on SRX100 devices because there is no Gigabit Ethernet port.**

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the **[edit interfaces interface-name]** hierarchy level, specifying **vlan-ccc** or **vlan-vpls**:

[edit interfaces *interface-name*] encapsulation (vlan-ccc | vlan-vpls);

To configure the encapsulation on an aggregated Ethernet interface, include the encapsulation statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **ethernet-vpls** or **vlan-vpls**:

[edit interfaces interface-name] encapsulation (ethernet-vpls | vlan-vpls);

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512

Copyright © 2016, Juniper Networks, Inc.

through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.

Related Documentation

- Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces on page 3909
- Understanding VPLS VLAN Encapsulation on a Logical Interface on page 3906
- Example: Configuring Extended VLAN VPLS Encapsulation on page 3910
- Example: Configuring VPLS VLAN Encapsulation on page 3906
- VPLS Overview on page 3841

Understanding VPLS VLAN Encapsulation on a Logical Interface

You cannot configure a logical interface with VLAN VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In VPLS mode, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLAN. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLAN.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.

Related Documentation

- VPLS Configuration Overview on page 3846
- Understanding VPLS VLAN Encapsulation on page 3905

Example: Configuring VPLS VLAN Encapsulation

This example shows how to configure VPLS VLAN encapsulation and enable it on the physical and the logical interfaces.

- Requirements on page 3907
- Overview on page 3907
- Configuration on page 3907
- Verification on page 3908

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE, see "Example: Configuring RSVP on the VPLS PE Router" on page 3868 and then MPLS LSPs on the PE routers, see "Example: Configuring MPLS on the VPLS PE Router" on page 3869. Alternatively configure LDP on the PE routers, see "Example: Configuring LDP on the VPLS PE Router" on page 3871.
- Configure routing options on the PE router. See "Example: Configuring Routing Options on the VPLS PE Router" on page 3903.
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See "Example: Configuring BGP on the VPLS PE Router" on page 3902.

Overview

This example describes how to enable VLAN tagging on VPLS interface ge-3/0/6, configure the encapsulation type on the physical and logical interfaces, and configure the VPLS family on the logical interface.

NOTE: Perform the following CLI quick configuration and procedures on all of the PE interfaces (CE facing).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-3/0/6 vlan-tagging
set interfaces ge-3/0/6 encapsulation vlan-vpls
set interfaces ge-3/0/6 unit 0 encapsulation vlan-vpls
set interfaces ge-3/0/6 unit 0 vlan-id 512
set interfaces ge-3/0/6 unit 0 family vpls
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure VPLS VLAN encapsulation:

1. Enable VLAN tagging on the VPLS interface.

   ```
   [edit interfaces ge-3/0/6]
   user@host# set vlan-tagging
   ```

2. Configure the encapsulation type on the physical interface.

   ```
   [edit interfaces ge-3/0/6]
   user@host# set interfaces ge-3/0/6 encapsulation vlan-vpls
   ```

3. Configure the encapsulation type on the logical interface.

   ```
   [edit interfaces ge-3/0/6 unit 0]
   user@host# set encapsulation vlan-vpls
   ```

4. Configure the VLAN ID on the logical interface.

   ```
   [edit interfaces ge-3/0/6 unit 0]
   user@host# set vlan-id 512
   ```

5. Configure the family VPLS on the logical interface.

   ```
   [edit interfaces ge-3/0/6 unit 0]
   user@host# set family vpls
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-3/0/6** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces ge-3/0/6
vlan-tagging;
encapsulation vlan-vpls;
unit 0 {
   encapsulation vlan-vpls;
   vlan-id 512;
   family vpls;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying VPLS VLAN Encapsulation on page 3908
- Verifying VPLS VLAN Encapsulation for Logical Interfaces on page 3909

Verifying VPLS VLAN Encapsulation

Purpose

Verify that the VPLS VLAN encapsulation is enabled at the interfaces.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show interfaces** command.

Verifying VPLS VLAN Encapsulation for Logical Interfaces

Purpose Verify that the VPLS VLAN encapsulation is enabled at the logical interface.

Action From operational mode, enter the **show interfaces ge-3/0/6 unit 0** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces

This example shows how to configure the VPLS VLAN encapsulation on either a Gigabit Ethernet IQ or Gigabit Ethernet physical interface.

- Requirements on page 3909
- Overview on page 3910
- Configuration on page 3910
- Verification on page 3910

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE, see "Example: Configuring RSVP on the VPLS PE Router" on page 3868 and then MPLS LSPs on the PE routers, see "Example: Configuring MPLS on the VPLS PE Router" on page 3869. Alternatively configure LDP on the PE routers, see "Example: Configuring LDP on the VPLS PE Router" on page 3871.
- Configure routing options on the PE router. See "Example: Configuring Routing Options on the VPLS PE Router" on page 3903.
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See "Example: Configuring BGP on the VPLS PE Router" on page 3902

Copyright © 2016, Juniper Networks, Inc.

Overview

This example describes how to configure Ethernet VPLS encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface and enable the VPLS family on the interface.

Configuration

Step-by-Step Procedure

To configure VPLS VLAN encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface:

1. Configure the ethernet-vpls encapsulation on the interface.

   ```
   [edit ]
   user@host# set interfaces ge-3/0/6 encapsulation ethernet-vpls
   ```

2. Enable the VPLS family on the interface.

   ```
   [edit ]
   user@host# set interfaces ge-3/0/6 unit 0 family vpls
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Example: Configuring Extended VLAN VPLS Encapsulation

This example shows how to configure extended VLAN VPLS encapsulation and enable it on the physical and the logical interfaces.

- Requirements on page 3911
- Overview on page 3911
- Configuration on page 3911
- Verification on page 3912

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Before you begin, see "Understanding Selective Stateless Packet-Based Services" on page 1804.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See "Example: Configuring Routing Interfaces on the VPLS PE Router" on page 3851 and "Example: Configuring the Interface to the VPLS CE Device" on page 3852.
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See "Example: Configuring the VPLS Routing Instance" on page 3862.
- Configure an IGP on the PE routers to exchange routing information. See "Example: Configuring OSPF on the VPLS PE Router" on page 3867.
- Configure RSVP-TE, see "Example: Configuring RSVP on the VPLS PE Router" on page 3868 and then MPLS LSPs on the PE routers, see "Example: Configuring MPLS on the VPLS PE Router" on page 3869. Alternatively configure LDP on the PE routers, see "Example: Configuring LDP on the VPLS PE Router" on page 3871.
- Configure routing options on the PE router. See "Example: Configuring Routing Options on the VPLS PE Router" on page 3903.
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See "Example: Configuring BGP on the VPLS PE Router" on page 3902.

Overview

This example describes how to enable VLAN tagging on the VPLS interface ge-3/0/6, configure the extended-vlan-vpls type on the physical and logical interfaces, and configure the VPLS family on the logical interface.

NOTE: Perform the following CLI quick configurations and procedures on all PE interfaces (CE facing).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-3/0/6 vlan-tagging
set interfaces ge-3/0/6 encapsulation extended-vlan-vpls
set interfaces ge-3/0/6 unit 0 vlan-id 100
set interfaces ge-3/0/6 unit 0 family vpls
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure extended VPLS VLAN encapsulation:

1. Enable VLAN tagging on the VPLS interface as it will receive tagged packets from CE.

   ```
   [edit interfaces ge-3/0/6]
   user@host# set vlan-tagging
   ```

2. Configure the encapsulation type on the physical interface.

   ```
   [edit interfaces ge-3/0/6]
   user@host# set interfaces ge-3/0/6 encapsulation vlan-vpls
   ```

3. Configure the VLAN ID on the logical interface.

   ```
   [edit interfaces ge-3/0/6 unit 0]
   user@host# set encapsulation vlan-vpls vlan-id 100
   ```

4. Configure the VPLS family on the logical interface.

   ```
   [edit interfaces ge-3/0/6 unit 0]
   user@host# set family vpls
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-3/0/6** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces ge-3/0/6
vlan-tagging;
encapsulation extended-vlan-vpls;
unit 0 {
   encapsulation vlan-vpls;
   vlan-id 100;
   family vpls;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Extended VLAN VPLS Encapsulation on page 3912
- Verifying Extended VLAN VPLS Encapsulation for Logical Interfaces on page 3913

Verifying Extended VLAN VPLS Encapsulation

Purpose

Verify that the extended VLAN VPLS encapsulation is enabled at the interfaces.

Action

From operational mode, enter the **show interfaces** command.

Copyright © 2016, Juniper Networks, Inc.

Verifying Extended VLAN VPLS Encapsulation for Logical Interfaces

Purpose Verify that the extended VLAN VPLS encapsulation is enabled at the logical interface.

Action From operational mode, enter the **show interfaces ge-3/0/6 unit 0** command.

Related Documentation

- VPLS Configuration Overview on page 3846
- VPLS Overview on page 3841

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 58

Configuration Statements and Operational Commands

- Configuration Statements on page 3917
- Operational Commands on page 3985

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 173

Configuration Statements

- Accounting-Options Configuration Statement Hierarchy on page 3917
- Firewall Configuration Statement Hierarchy on page 3919
- Forwarding-Options Configuration Statement Hierarchy on page 3929
- Interfaces Configuration Statement Hierarchy on page 3941
- Policy-Options Configuration Statement Hierarchy on page 3957
- [edit security address-book] Hierarchy Level on page 3958
- [edit security forwarding-options] Hierarchy Level on page 3959
- [edit security ike] Hierarchy Level on page 3959
- [edit security ipsec] Hierarchy Level on page 3961
- [edit security policies] Hierarchy Level on page 3963
- condition (Policy Options) on page 3967
- family (Security Forwarding Options) on page 3968
- flow-server (Forwarding Options) on page 3969
- forwarding-options (Security) on page 3971
- fragment on page 3972
- hash-key (Forwarding Options) on page 3973
- iso (Security Forwarding Options) on page 3974
- mpls (Security Forwarding Options) on page 3975
- multicast-scope on page 3976
- policer (Firewall) on page 3977
- simple-filter (Firewall) on page 3978
- template (Flow Monitoring) on page 3980
- traceoptions (Security Flow) on page 3981
- version9 (Flow Server) on page 3984

Accounting-Options Configuration Statement Hierarchy

Use the statements in the **accounting-options** configuration hierarchy to collect and log data about basic system operations and services on the device.

Copyright © 2016, Juniper Networks, Inc.

```
accounting-options {
  class-usage-profile profile-name {
    destination-classes {
      destination-class-name;
    }
    file filename;
    interval minutes;
    source-classes {
      source-class-name;
    }
  }
  file filename {
    archive-sites url {
      password password ;
    }
    files number;
    nonpersistent;
    size bytes;
    start-time yyyy-mm-dd.hh:mm;
    transfer-interval minutes;
  }
  filter-profile profile-name {
    counters {
      counter-name;
    }
    file filename;
    interval minutes;
  }
  interface-profile profile-name {
    fields {
      field-name;
    }
    file filename;
    interval minutes;
  }
  mib-profile profile-name {
    file filename;
    interval minutes;
    object-names {
      mib-object-name;
    }
    operation (get | get-next | walk) ;
  }
  periodic-refresh disable;
  routing-engine-profile profile-name {
    fields {
      field-name ;
    }
    file filename;
    interval minutes;
  }
}
```

Related Documentation

- *Administration Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

- Configuration Statements at the [edit accounting-options] Hierarchy Level

Firewall Configuration Statement Hierarchy

Use the statements in the **firewall** configuration hierarchy to configure stateless firewall filters—also known as access control lists (ACLs)—on the device.

```
firewall {
  family {
    any {
      filter filter-name {
        term term-name {
          from {
            forwarding-class [forwarding-class-name];
            forwarding-class-except [forwarding-class-name];
            interface interface-name;
            interface-set interface-set-name;
            packet-length [range];
            packet-length-except [range];
          }
          then {
            accept;
            count value;
            discard;
            forwarding-class forwarding-class-name;
            loss-priority (high |low |medium-high |medium-low);
            next term;
            policer policer-name;
          }
        }
      }
    }
    ethernet-switching {
      filter filter-name {
        accounting-profile [accounting-profile-name];
        interface-specific;
        term term-name {
          filter filter-name;
          from {
            forwarding-class [forwarding-class-name];
            forwarding-class-except [forwarding-class-name];
            interface interface-name;
            interface-set interface-set-name;
            packet-length [range];
            packet-length-except [range];
          }
          then {
            accept;
            count value;
            discard;
            forwarding-class forwarding-class-name;
            loss-priority (high |low |medium-high |medium-low);
            next term;
            policer policer-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
  ccc {
    filter filter-name {
      accounting-profile [accounting-profile-name];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          (interface-group [ group-names ] | interface-group-except [ group-names ]);
        }
        then {
          accept;
          count value;
          discard;
          forwarding-class forwarding-class-name;
          loss-priority (high |low |medium-high |medium-low);
          next term;
          packet mode;
          policer policer-name;
        }
      }
    }
  }
  inet {
    dialer-filter filter-name {
      accounting-profile [accounting-profile-name];
      term term-name {
        from {
          address {
            ip-prefix</prefix-length> <except>;
          }
          destination-address {
            ip-prefix</prefix-length> <except>;
          }
          (destination-port [ port-names ] | destination-port-except [ port-names ]);
          destination-prefix-list {
            list-name <except>;
          }
          (dscp [ code-point-values ] | dscp-except [ code-point-values ]);
          (esp-spi [ values ] | esp-spi-except [ values ]);
          first-fragment;
          fragment-flags flag;
          (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
          (icmp-code [ codes ] | icmp-code-except [ codes ]);
          (icmp-type [ types ] | icmp-type-except [ types ]);
          (ip-options [ option-names ] | ip-options-except [ option-names ]);
          is-fragment;
          (packet-length [ values ] | packet-length-except [ values ]);
          (port [ port-names ] | port-except [ port-names ]);
          (precedence [ precedence-names ] | precedence-except [ precedence-names ]);
          prefix-list {
            list-name <except>;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        source-address {
          ip-prefix</prefix-length> <except>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
          list-name <except>;
        }
        tcp-established;
        tcp-flags flag;
        tcp-initial;
        (ttl [ ttl-values ] | ttl-except [ ttl-values ]);
      }
      then {
        (ignore | note);
        log;
        sample;
        syslog;
      }
    }
  }
  filter filter-name {
    accounting-profile [accounting-profile-name];
    interface-specific;
    term term-name {
      from {
        address {
          ip-prefix</prefix-length> <except>;
        }
        destination-address {
          ip-prefix</prefix-length> <except>;
        }
        (destination-port [ port-names ] | destination-port-except [ port-names ]);
        destination-prefix-list {
        list-name <except>;
        }
        (dscp [ code-point-values ] | dscp-except [ code-point-values ]);
        (esp-spi [ values ] | esp-spi-except [ values ]);
        first-fragment;
        (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
        fragment-flags flag;
        (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
        (icmp-code [ codes ] | icmp-code-except [ codes ]);
        (icmp-type [ types ] | icmp-type-except [ types ]);
        (interface-group [ group-names ] | interface-group-except [ group-names ]);
        interface-set set-name;
        (ip-options [ option-names ] | ip-options-except [ option-names ]);
        is-fragment;
        (packet-length [ values ] | packet-length-except [ values ]);
        (port [ port-names ] | port-except [ port-names ]);
        (precedence [ precedence-names ] | precedence-except [ precedence-names ]);
        prefix-list {
          list-name <except>;
        }
        (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
```

```
            service-filter-hit;
            source-address {
              ip-prefix</prefix-length> <except>;
            }
            (source-port [ port-names ] | source-port-except [ port-names ]);
            source-prefix-list {
              list-name <except>;
            }
            tcp-established;
            tcp-flags flag;
            tcp-initial;
          }
          then {
            (accept | discard | reject);
            count counter-name;
            forwarding-class class-name;
            log;
            loss-priority (high | low);
            next;
            packet-mode;
            policer policer-name;
            port-mirror;
            routing-instance routing-instance-name> <topology topology-name>;
            sample;
            service-accounting;
            service-filter-hit;
            syslog;
            topology topology-name;
            virtual-channel;
          }
        }
      }
      prefix-action prefix-action-name {
        count;
        destination-prefix-length length;
        filter-specific;
        policer policer-name;
        source-prefix-length length;
        subnet-prefix-length length;
      }
      service-filter filter-name {
        term term-name {
          from {
            address {
              ip-prefix</prefix-length> <except>;
            }
            destination-address {
              ip-prefix</prefix-length> <except>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
              list-name <except>;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (packet-length [ values ] | packet-length-except [ values ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
              list-name <except>;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
            source-address {
              ip-prefix</prefix-length> <except>;
            }
            (source-port [ port-names ] | source-port-except [ port-names ]);
            source-prefix-list {
              list-name <except>;
            }
            tcp-flags flag;
          }
          then {
            count counter-name;
            log;
            port-mirror;
            sample;
            (service | skip);
          }
        }
      }
      simple-filter filter-name {
        term term-name {
          from {
            destination-address ip-prefix</prefix-length>;
            destination-port port-name;
            protocol protocol-name;
            source-address ip-prefix</prefix-length>;
            source-port port-name;
          }
          then {
            (accept | discard);
            forwarding-class class-name;
            policer policer-name;
            three-color-policer policer-name {
              (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
            }
          }
        }
      }
    }
    inet6 {
      dialer-filter filter-name {
        accounting-profile [accounting-profile-name];
        term term-name {
          from {
            address {
              ip-prefix</prefix-length> <except>;
            }
```

```
        destination-address {
          ip-prefix</prefix-length> <except>;
        }
        (destination-port [ port-names ] | destination-port-except [ port-names ]);
        destination-prefix-list {
          list-name <except>;
        }
        (icmp-code [ codes ] | icmp-code-except [ codes ]);
        (icmp-type [ types ] | icmp-type-except [ types ]);
        (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
        (packet-length [ values ] | packet-length-except [ values ]);
        (port [ port-names ] | port-except [ port-names ]);
        (precedence [ precedence-names ] | precedence-except [ precedence-names ]);
        prefix-list {
          list-name <except>;
        }
        source-address {
          ip-prefix</prefix-length> <except>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
          list-name <except>;
        }
      }
      then {
        (ignore | note);
        log;
        sample;
        syslog;
      }
    }
  }
  filter filter-name {
    accounting-profile [accounting-profile-name];
    interface-specific;
    term term-name {
      filter filter-name;
      from {
        address {
          ip-prefix</prefix-length> <except>;
        }
        destination-address {
          ip-prefix</prefix-length> <except>;
        }
        (destination-port [ port-names ] | destination-port-except [ port-names ]);
        destination-prefix-list {
          list-name <except>;
        }
        (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
        (icmp-code [ codes ] | icmp-code-except [ codes ]);
        (icmp-type [ types ] | icmp-type-except [ types ]);
        interface interface-name;
        (interface-group [ group-names ] | interface-group-except [ group-names ]);
        interface-set set-name;
        (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
        (packet-length [ values ] | packet-length-except [ values ]);
```


Copyright © 2016, Juniper Networks, Inc.

```
        (port [ port-names ] | port-except [ port-names ]);
        prefix-list {
          list-name <except>;
        }
        service-filter-hit;
        source-address {
          ip-prefix</prefix-length> <except>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
          list-name <except>;
        }
        tcp-established;
        tcp-flags flag;
        tcp-initial;
        (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
      }
      then {
        then {
          (accept | discard | reject);
          count counter-name;
          forwarding-class class-name;
          log;
          loss-priority (high | low);
          next;
          packet-mode;
          policer policer-name;
          routing-instance routing-instance-name> <topology topology-name>;
          sample;
          service-accounting;
          service-filter-hit;
          syslog;
          topology topology-name;
        }
      }
    }
  }
  service-filter filter-name {
    term term-name {
      from {
        address {
          ip-prefix</prefix-length> <except>;
        }
        destination-address {
          ip-prefix</prefix-length> <except>;
        }
        (destination-port [ port-names ] | destination-port-except [ port-names ]);
        destination-prefix-list {
          list-name <except>;
        }
        (esp-spi [ values ] | esp-spi-except [ values ]);
        (interface-group [ group-names ] | interface-group-except [ group-names ]);
        (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
        (port [ port-names ] | port-except [ port-names ]);
        prefix-list {
          list-name <except>;
```

```
            }
            source-address {
              ip-prefix</prefix-length> <except>;
            }
            (source-port [ port-names ] | source-port-except [ port-names ]);
            source-prefix-list {
              list-name <except>;
            }
            tcp-flags flag;
          }
          then {
            count counter-name;
            log;
            port-mirror;
            sample;
            (service | skip);
          }
        }
      }
    }
    mpls {
      dialer-filter filter-name {
        accounting-profile [accounting-profile-name];
        term term-name {
          from {
            (exp [ exp-bits ] | exp-except [ exp-bits ]);
          }
          then {
            (ignore | note);
            log;
            sample;
            syslog;
          }
        }
      }
      filter filter-name {
        accounting-profile [accounting-profile-name];
        interface-specific;
        term term-name {
          filter filter-name;
          from {
            (exp [ exp-bits ] | exp-except [ exp-bits ]);
            (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
            interface interface-name;
            interface-set set-name;
          }
          then {
            (accept | discard);
            count value;
            forwarding-class forwarding-class-name;
            loss-priority (high |low |medium-high |medium-low);
            next term;
            policer policer-name;
            sample;
            three-color-policer policer-name {
              (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
```

 Copyright © 2016, Juniper Networks, Inc.

```
            }
          }
        }
      }
    }
    vpls {
      filter filter-name {
        accounting-profile [accounting-profile-name];
        interface-specific;
        term term-name {
          filter filter-name;
          from {
            destination-mac-address {
              mac-address;
            }
            (ether-type [ protocol-types ] | ether-type-except [ protocol-types ]);
            (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
            interface interface-name;
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            interface-set interface-set-name;
            source-mac-address {
              mac-address <except>;
            }
            (vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ]);
          }
          then {
            (accept | discard);
            count value;
            forwarding-class forwarding-class-name;
            loss-priority (high |low |medium-high |medium-low);
            next term;
            policer policer-name;
          }
        }
      }
    }
  }
  filter filter-name {
    accounting-profile [accounting-profile-name];
    interface-specific;
    term term-name {
      filter filter-name;
      from {
        address {
          ip-prefix</prefix-length> <except>;
        }
        destination-address {
          ip-prefix</prefix-length> <except>;
        }
        (destination-port [ port-names ] | destination-port-except [ port-names ]);
        destination-prefix-list {
          list-name <except>;
        }
        (dscp [ code-point-values ] | dscp-except [ code-point-values ]);
        (esp-spi [ values ] | esp-spi-except [ values ]);
        first-fragment;
```

Copyright © 2016, Juniper Networks, Inc.

```
            (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (icmp-code [ codes ] | icmp-code-except [ codes ]);
            (icmp-type [ types ] | icmp-type-except [ types ]);
            interface interface-name;
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            interface-set set-name;
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (packet-length [ values ] | packet-length-except [ values ]);
            (port [ port-names ] | port-except [ port-names ]);
            (precedence [ precedence-names ] | precedence-except [ precedence-names ]);
            prefix-list {
              list-name <except>;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
            service-filter-hit;
            source-address {
              ip-prefix</prefix-length> <except>;
            }
            (source-port [ port-names ] | source-port-except [ port-names ]);
            source-prefix-list {
              list-name <except>;
            }
            tcp-established;
            tcp-flags flag;
            tcp-initial;
          }
          then {
            (accept | discard | reject);
            count counter-name;
            forwarding-class class-name;
            log;
            loss-priority (high | low);
            next;
            packet-mode;
            policer policer-name;
            port-mirror;
            routing-instance routing-instance-name> <topology topology-name>;
            sample;
            service-accounting;
            service-filter-hit;
            syslog;
            topology topology-name;
            virtual-channel;
          }
        }
      }
      interface-set interface-set-name {
        interface-name;
      }
      policer policer-name {
        filter-specific;
        if-exceeding {
          (bandwidth-limit bps | bandwidth-percent percentage);
```


Copyright © 2016, Juniper Networks, Inc.

```
        burst-size-limit bytes;
      }
      logical-interface-policer;
      then {
        discard;
        forwarding-class forwarding-class-name;
        loss-priority (high |low |medium-high |medium-low);
        out-of-profile;
      }
    }
    three-color-policer policer-name {
      filter-specific;
      single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
      }
      two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
      }
    }
  }
```

Related Documentation

- MPLS Overview on page 3743

Forwarding-Options Configuration Statement Hierarchy

Use the statements in the **forwarding-options** configuration hierarchy to configure forwarding options including protocol family, flow monitoring, accounting properties, load-balancing, port mirroring, traffic sampling, and packet capture.

```
forwarding-options {
  accounting group-name {
    output {
      aggregate-export-interval seconds;
      cflowd hostname {
        aggregation {
          autonomous-system;
          destination-prefix;
          protocol-port;
          source-destination-prefix {
            caida-compliant;
          }
          source-prefix;
        }
        autonomous-system-type (origin | peer);
        port port-number;
        version (5 | 8);
      }
```

```
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
        source-address ip-address;
      }
    }
  }
  dhcp-relay {
    active-server-group server-group-name;
    dhcpv6 {
      active-server-group server-group-name;
      authentication {
        password password;
        username-include {
          circuit-type;
          client-id;
          delimiter delimiter;
          domain-name domain-name;
          interface-name;
          logical-system-name;
          relay-agent-interface-id;
          relay-agent-remote-id;
          relay-agent-subscriber-id;
          routing-instance-name;
          user-prefix user-prefix;
        }
      }
      dynamic-profile (dynamic-profile-name | junos-default-profile) {
        aggregate-clients (merge | replace);
        use-primary (dynamic-profile-name | junos-default-profile);
      }
      group group-name {
        active-server-group server-group-name;
        authentication {
          password password;
          username-include {
            circuit-type;
            client-id;
            delimiter delimiter;
            domain-name domain-name;
            interface-name;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix;
          }
        }
        dynamic-profile (dynamic-profile-name | junos-default-profile) {
          aggregate-clients (merge | replace);
          use-primary (dynamic-profile-name | junos-default-profile);
        }
        interface interface-name  {
```


Copyright © 2016, Juniper Networks, Inc.

```
    dynamic-profile (dynamic-profile-name | junos-default-profile) {
      aggregate-clients (merge | replace);
      use-primary (dynamic-profile-name | junos-default-profile);
    }
    exclude;
    overrides {
      (allow-snooped-clients | no-allow-snooped-clients);
      interface-client-limit number;
      no-bind-on-request;
      send-release-on-delete;
    }
    service-profile service-profile;
    trace;
    upto interface-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        detection-time {
          threshold milliseconds;
        }
        holdown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        session-mode (automatic | multihop | single-hop);
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        version (0 | 1 | automatic);
      }
    }
  }
  overrides {
    (allow-snooped-clients | no-allow-snooped-clients);
    interface-client-limit number;
    no-bind-on-request;
    send-release-on-delete;
  }
  relay-agent-interface-id {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
  }
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
```

```
            detection-time {
              threshold milliseconds;
            }
            holdown-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
          }
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
              }
            version (0 | 1 | automatic);
        }
      overrides {
        (allow-snooped-clients | no-allow-snooped-clients);
        interface-client-limit number;
        no-bind-on-request;
        send-release-on-delete;
      }
      relay-agent-interface-id {
        prefix {
          host-name;
          logical-system-name;
          routing-instance-name;
        }
        use-interface-description (device | logical);
      }
      server-group server-group-name {
        ip-address;
      }
      service-profile dynamic-profile-name;
    }
    group group-name {
      active-server-group server-group-name;
      interface interface-name {
        exclude;
        upto interface-name;
      }
      relay-option-60 {
        vendor-option {
          default-local-server-group local-server-group-name;
          default-relay-server-group relay-server-group-name;
          drop;
          equals {
            ascii ascii-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group relay-server-group-name;
            }
            hexadecimal hexadecimal-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group relay-server-group-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
          }
          starts-with {
            ascii ascii-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group relay-server-group-name;
            }
            hexadecimal hexadecimal-name {
              drop;
              local-server-group local-server-group-name;
              relay-server-group relay-server-group-name;
            }
          }
        }
      }
      relay-option-82 {
        circuit-id {
          prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
          }
          use-interface-description (device | logical);
        }
      }
    }
    relay-option-60 {
      vendor-option {
        default-local-server-group local-server-group-name;
        default-relay-server-group relay-server-group-name;
        drop;
        equals {
          ascii ascii-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group relay-server-group-name;
          }
          hexadecimal hexadecimal-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group relay-server-group-name;
          }
        }
        starts-with {
          ascii ascii-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group relay-server-group-name;
          }
          hexadecimal hexadecima-name {
            drop;
            local-server-group local-server-group-name;
            relay-server-group relay-server-group-name;
          }
        }
```

```
    }
  }
  relay-option-82 {
    circuit-id {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
    }
  }
  server-group server-group-name {
    ip-address;
  }
}
family inet {
  filter {
    input filter-name;
    output filter-name;
  }
}
family inet6 {
  filter {
    input filter-name;
    output filter-name;
  }
}
family mpls {
  filter {
    input filter-name;
    output filter-name;
  }
}
family vpls {
  filter {
    input filter-name;
  }
  flood {
    input filter-name;
  }
}
hash-key {
  family inet {
    layer-3;
    layer-4;
    session-id;
  }
  family mpls {
    label-1;
    label 2;
    label-3;
    no-labels;
    payload {
      ip {
        layer-3-only;
```


Copyright © 2016, Juniper Networks, Inc.

```
            port-data {
              destination-lsb;
              destination-msb;
              source-lsb;
              source-msb;
            }
          }
        }
      }
      family multiservice {
        destination-mac;
        source-mac;
      }
    }
    helpers {
      bootp {
        client-response-ttl number;
        description text;
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          disable;
          remote-id {
            prefix (hostname | mac | none);
            use-interface-description;
            use-string text;
          }
          vendor-id {
            text;
          }
        }
        interface interface-name {
          broadcast;
          client-response-ttl number;
          description text;
          dhcp-option82 {
            circuit-id {
              prefix hostname;
              use-interface-description;
              use-vlan-id;
            }
            disable;
            remote-id {
              prefix (hostname | mac | none);
              use-interface-description;
              use-string text;
            }
            vendor-id {
              text;
            }
          }
          maximum-hop-count number;
          minimum-wait-time seconds;
```

```
      no-listen;
      server name-or-address {
        logical-system logical-system-name;
        routing-instance (default | routing-instance-name);
      }
      vpn;
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server name-or-address {
      logical-system logical-system-name;
      routing-instance (default | routing-instance-name);
    }
    vpn;
  }
  domain {
    description text;
    interface interface-name {
      broadcast;
      description text;
      no-listen;
      server name-or-address {
        logical-system logical-system-name;
        routing-instance (default | routing-instance-name);
      }
    }
    server name-or-address {
      logical-system logical-system-name;
      routing-instance (default | routing-instance-name);
    }
  }
  port port-number  {
    description text;
    interface interface-name {
      broadcast;
      description text;
      no-listen;
      server name-or-address {
        logical-system logical-system-name;
        routing-instance (default | routing-instance-name);
      }
    }
    server name-or-address {
      logical-system logical-system-name;
      routing-instance (default | routing-instance-name);
    }
  }
  rtsdb-client-traceoptions {
    if-rtsdb {
      flag all <disable>;
      flag init <disable>;
      flag map <disable>;
      flag routing-socket <disable>;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    tftp {
      description text;
      interface interface-name {
        broadcast;
        description text;
        no-listen;
        server name-or-address {
          logical-system logical-system-name;
          routing-instance (default | routing-instance-name);
        }
      }
      server name-or-address {
        logical-system logical-system-name;
        routing-instance (default | routing-instance-name);
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  load-balance {
    indexed-load-balance;
    per-prefix {
      hash-seed number;
    }
  }
  packet-capture {
    disable;
    file {
      filename;
      files number;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    maximum-capture-size size;
  }
  port-mirroring {
    disable;
    family any {
      output {
        interface interface-name;
      }
    }
    family inet {
      output {
        interface interface-name {
          next-hop ip-address;
```

```
            }
          no-filter-check;
        }
      }
      input {
        rate number;
        run-length number;
      }
      instance instance-name {
        family any {
          output {
            interface interface-name;
          }
        }
        input {
          rate number;
          run-length number;
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        no-remote-trace;
      }
    }
    sampling {
      disable;
      family inet {
        disable;
        output {
          aggregate-export-interval seconds;
          extension-service service-name;
          file {
            disable;
            filename filename;
            files number;
            size maximum-file-size;
            (stamp | no-stamp);
            (world-readable | no-world-readable);
          }
          flow-active-timeout seconds;
          flow-inactive-timeout seconds;
          flow-server ip-address-or-host-name {
            aggregation {
              autonomous-system;
              destination-prefix;
              protocol-port;
              source-destination-prefix {
                caida-compliant;
              }
              source-prefix;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address ip-address;
            version (5 | 500 |8);
            version9 {
              template template-name;
            }
          }
          inline-jflow {
            flow-export-rate number;
            source-address ip-address;
          }
        }
      }
      family inet6 {
        disable;
        output {
          aggregate-export-interval seconds;
          extension-service service-name;
          flow-active-timeout seconds;
          flow-inactive-timeout seconds;
          flow-server host-name {
            aggregation {
              autonomous-system;
              destination-prefix;
              protocol-port;
              source-destination-prefix {
                caida-compliant;
              }
              source-prefix;
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address ip-address;
            version9 {
              template template-name;
            }
          }
        }
      }
      input {
        max-packets-per-second number;
        maximum-packet-length number;
        rate number;
        run-length number;
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
```

```
    }
    no-remote-trace;
  }
  instance instance-name {
    disable;
    family inet {
      disable;
      output {
        aggregate-export-interval seconds;
        extension-service service-name;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-server ip-address-or-host-name {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
            }
            source-prefix;
          }
          autonomous-system-type (origin | peer);
          (local-dump | no-local-dump);
          port port-number;
          source-address ip-address;
          version (5 | 8);
          version-ipfix {
            template template-name;
          }
          version9 {
            template template-name;
          }
        }
        inline-jflow {
          flow-export-rate number;
          source-address ip-address;
        }
      }
    }
    family inet6 {
      disable;
      output {
        aggregate-export-interval seconds;
        extension-service service-name;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-server ip-address-or-host-name {
          aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
              caida-compliant;
            }
            source-prefix;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address ip-address;
            version-ipfix {
              template template-name;
            }
            version9 {
              template template-name;
            }
          }
          inline-jflow {
            flow-export-rate number;
            source-address ip-address;
          }
        }
      }
      family mpls {
        disable;
        output {
          aggregate-export-interval seconds;
          flow-active-timeout seconds;
          flow-inactive-timeout seconds;
          flow-server ip-address-or-host-name {
            aggregation {
              autonomous-system;
              destination-prefix;
              protocol-port;
              source-destination-prefix {
                caida-compliant;
              }
              source-prefix;
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address ip-address;
          }
        }
      }
    }
  }
}
```

Related Documentation

- [edit security forwarding-options] Hierarchy Level on page 3006
- *Interfaces Feature Guide for Security Devices*
- *Junos OS Monitoring and Troubleshooting Library for Security Devices*

Interfaces Configuration Statement Hierarchy

Use the statements in the **interfaces** configuration hierarchy to configure interfaces on the device.

Copyright © 2016, Juniper Networks, Inc.

NOTE: For a gigabit ethernet interface, the **gigether-options** and **ether-options** are identical, but only the **gigether-options** are documented.

```
interfaces {
  interface-name {
    accounting-profile name;
    clocking (external | internal);
    dce;
    description text;
    disable;
    e1-options {
      bert-algorithm algorithm;
      bert-error-rate rate;
      bert-period seconds;
      fcs (16 | 32);
      framing (g704 | g704-no-crc4 | unframed);
      idle-cycle-flag (flags | ones);
      invert-data data;
      loopback (local | remote);
      start-end-flag (shared | filler);
      timeslots time-slot-range;
    }
    e3-options {
      bert-algorithm algorithm;
      bert-error-rate rate;
      bert-period seconds;
      compatibility-mode {
        digital-link {
          subrate value;
        }
        kentrox {
          subrate value;
        }
        larscom;
      }
      fcs (16 | 32);
      framing (g.751 | g.832);
      idle-cycle-flag value;
      invert-data;
      loopback (local | remote);
      (no-payload-scrambler | payload-scrambler);
      (no-unframed | -unframed);
      start-end-flag (filler | shared);
    }
    encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc |
      ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc |
      extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls
      | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
    fastether-options {
      802.3ad interface-name {
        (backup | primary);
        lacp {
          port-priority port-number;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    (auto-negotiation | no-auto-negotiation);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
}
flexible-vlan-tagging;
gigether-options {
    802.3ad interface-name {
        (backup | primary);
        lacp {
            port-priority port-number;
        }
    }
    (auto-negotiation <remote-fault> (local-interface-offline | local-interface-online)
        | no-auto-negotiation);
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth [number];
        }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
}
gratuitous-arp-reply;
hierarchical-scheduler {
    maximum-hierarchy-levels 2;
}
hold-time {
    down milliseconds;
    up milliseconds;
}
keepalives {
    down-count number;
    interval number;
    up-count number;
}
link-mode (full-duplex | half-duplex);
lmi {
    lmi-type (ansi | c-lmi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
```

```
    t391dte number;
    t392dce number;
  }
  logical-tunnel-options {
    per-unit-mac-disable;
  }
  mac mac-address;
  mtu bytes;
  native-vlan-idvlan-id;
  no-gratuitous-arp-request;
  no-keepalives;
  optics-options {
    alarm {
      low-light-alarm (link-down | syslog);
    }
    warning {
      low-light-warning (link-down | syslog);
    }
    wavelength wavelength-options;
  }
  otn-options {
    bytes {
      transmit-payload-type number];
    }
    fec (efec | gfec | none);
    (laser-enable | no-laser-enable);
    (line-loopback | no-line-loopback);
    rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
    trigger {
      oc-lof {
        hold-time {
          down milliseconds;
          up milliseconds;
          }
        ignore;
      }
      oc-lom {
        hold-time {
          down milliseconds;
          up milliseconds;
          }
        ignore;
      }
      oc-los {
        hold-time {
          down milliseconds;
          up milliseconds;
          }
        ignore;
      }
      oc-wavelength-lock {
        hold-time {
          down milliseconds;
          up milliseconds;
          }
        ignore;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
odu-ais {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-bdi {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-lck {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-oci {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-sd {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-tca-bbe {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-tca-es {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-tca-ses {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
```

Copyright © 2016, Juniper Networks, Inc.

```
}
odu-tca-uas {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
odu-ttim {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
opu-ptim {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-ais {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-bdi {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-bdi {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-fec-deg {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-fec-deg {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
otu-fec-exe {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-iae {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-sd {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-tca-bbe {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-tca-es {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-tca-ses {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-tca-uas {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
}
otu-ttim {
  hold-time {
    down milliseconds;
    up milliseconds;
    }
  ignore;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi |
        otu-dapi |otu-expected-receive-dapi | otu-expected-receive-sapi |otu-sapi);
}
passive-monitor-mode;
(per-unit-scheduler | no-per-unit-schedule);
port-mirror-instance;
ppp-options {
    chap {
        access-profile name;;
        default-chap-secret secret;
        local-name name;
        no-rfc2486;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile (dynamic-profile | junos-default-profile);
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
        access-profile name;
        default-password password;
        local-name name;
        local-password password;
        no-rfc2486;
        passive;
    }
}
promiscuous-mode;
receive-bucket {
    overflow {
        discard;
        tag;
    }
    rate number;
    threshold number;
}
redundant-pseudo-interface-options {
    redundancy-group number;
}
satop-options {
    excessive-packet-loss-rate {
        sample-period milliseconds;
        threshold percentage;
    }
    idle-pattern number;
    (jitter-buffer-auto-adjust | jitter-buffer-latency milliseconds | jitter-buffer-packets
        number;
```


Copyright © 2016, Juniper Networks, Inc.

```
    payload-size number;
}
speed (100m |10m | 1g);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        cascade-port;
        link-mode (full-duplex | half-duplex);
        speed (100m |10m | 1g);
        vlan-id number;
    }
}
t1-options {
    alarm-compliance {
        accunet-t1-5-service;
    }
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flags (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    bert-algorithm algorithm ;
    bert-error-rate rate ;
    bert-period seconds ;
    (cbit-parity | no-cbit-parity);
    compatibility-mode {
        adtran {
            subrate value;
        }
        digital-link {
            subrate value;
        }
        kentrox {
            subrate value;
        }
        larscom;
            subrate value;
        }
        verilink;
            subrate value;
        }
    }
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
```

```
    idle-cycle-flag (flags | ones);
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (no-payload-scrambler | payload-scrambler);
    (no-unframed | unframed);
    start-end-flag value (filler | shared);
}
traceoptions {
    flag (all | event | ipc | media);
}
transmit-bucket {
    overflow {
        discard;
    }
    rate number;
    threshold number;
}
(traps | no-traps);
unit unit-number {
    accept-source-mac {
        mac-address mac-address;
    }
    accounting-profile name;
    arp-resp (restricted | unrestricted);
    backup-options {
        interface interface-name;
    }
    bandwidth bandwidth;
    description text;
    disable;
    encapsulation (dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge |
        vlan-ccc | vlan-vpls |vlan-tcc);
    family {
        ccc {
            filter {
                group number;
                input filter-name;
                input-list [filter-name];
                output filter-name;
                output-list [filter-name];
            }
            policer {
                input input-policer-name;
                output output-policer-name;
            }
        }
        ethernet-switching {
            bridge-domain-type (svlan| bvlan);
            filter {
                group number;
                input filter-name;
                input-list [filter-name];
                output filter-name;
                output-list [filter-name];
            }
```


Copyright © 2016, Juniper Networks, Inc.

```
    interface-mode (access | trunk);
    native-vlan-id native-vlan-id;
    policer {
      input input-policer-name;
      output outputpolicer-name;
    }
    port-mode (access | tagged-access | trunk);
    reflective-relay;
    vlan-id vlan-id;
    vlan members [vlan-id];
    vlan-rewrite {
      translate {
        from-vlan-id;
        to-vlan-id ;
      }
    }
  }
  inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
    address (source–address/prefix) {
      arp destination-address {
        (mac mac-address | multicast-mac multicast-mac-address);
        publish publish-address;
      }
      broadcast address;
      preferred;
      primary;
      vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key-value;
        authentication-type (md5 | simple);
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds
        (preempt <hold-timeseconds> | no-preempt );
        priority value;
        track {
          interface interface-name {
            bandwidth-threshold bandwidth;
            priority-cost value;
          }
          priority-hold-time seconds;
          route route-address{
            routing-instance routing-instance;
            priority-cost value;
          }
        }
        virtual-address [address];
        virtual-link-local-address address;
```

Copyright © 2016, Juniper Networks, Inc.

```
      vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
      }
    }
    web-authentication {
      http;
      https;
      redirect-to-https;
    }
  }
  dhcp {
    client-identifier {
      (ascii string | hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  dhcp-client {
    client-identifier {
      prefix {
        host-name;
        logical-system-name;
        routing-instance-name;
      }
      use-interface-description (device | logical);
      user-id (ascii string| hexadecimal string);
    }
    lease-time (length | infinite);
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
  }
  filter {
    group number;
    input filter-name;
    input-list [filter-name];
    output filter-name;
    output-list [filter-name];
  }
  mtu value;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp arp-name;
    input input-name;
    output output-name;
  }
  primary;
  rpf-check {
```


Copyright © 2016, Juniper Networks, Inc.

```
        fail-filter filter-name;
        mode {
          loose;
        }
      }
      sampling {
        input;
        output;
        simple-filter;
      }
      targeted-broadcast {
        (forward-and-send-to-re |forward-only);
      }
      unnumbered-address {
        interface-name;
        preferred-source-address preferred-source-address;
      }
    }
    inet6 {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          ouput;
        }
      }
      address source–address/prefix {
        eui-64;
        ndp address {
          (mac mac-address | multicast-mac multicast-mac-address);
          publish;
        }
        preferred;
        primary;
        vrrp-inet6-group group_id {
          (accept-data | no-accept-data);
          advertisements-threshold number;
          authentication-key value;
          authentication-type (md5 | simple);
          fast-interval milliseconds;
          inet6-advertise-interval milliseconds;
          (preempt <hold-time seconds>| no-preempt );
          priority value;
          track {
            interface interface-name {
              bandwidth-threshold value;
              priority-cost value;
            }
            priority-hold-time seconds;
            route route-address{
              routing-instance routing-instance;
            }
          }
          virtual-inet6-address [address];
          virtual-link-local-address address;
          vrrp-inherit-from {
```

```
            active-group value;
            active-interface interface-name;
          }
        }
        web-authentication {
          http;
          https;
          redirect-to-https;
        }
      }
      (dad-disable | no-dad-disable);
      dhcpv6-client {
        client-ia-type (ia-na | ia-pd);
        client-identifier duid-type (duid-ll | duid-llt | vendor);
        client-type (autoconfig | stateful);
        rapid-commit;
        req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server |
          sip-domain | sip-server |time-zone | vendor-spec);
        retransmission-attempt number;
        update-router-advertisement {
          interface interface-name;
        }
        update-server;
      }
      filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
      }
      mtu value;
      nd6-stale-time seconds;
      no-neighbor-learn;
      policer {
        input input-name;
        output output-name;
      }
      rpf-check {
        fail-filter filter-name;
        mode {
          loose;
        }
      }
      sampling {
        input;
        output;
      }
      unnumbered-address {
        interface-name;
        preferred-source-address preferred-source-address;
      }
    }
    iso {
      address source-address;
      mtu value;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    mlfr-end-to-end {
      bundle bundle-name;
    }
    mlfr-uni-nni {
      bundle bundle-name;
    }
    mlppp {
      bundle bundle-name;
    }
    mpls {
      filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
      }
      mtu mtu-value;
      policer {
        input input-name;
        output output-name;
      }
    }
    tcc {
      policer {
        input input-name;
        output output-name;
      }
      proxy {
        inet-address inet-address;
      }
      remote {
        inet-address inet-address;
        mac-address mac-address;
      }
    }
    vpls {
      filter {
        group number;
        input filter-name;
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
      }
      policer {
        input input-name;
        output output-name;
      }
    }
  }
  input-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number ;
    (pop | push | swap);
    tag-protocol-id tpid;
```

```
    vlan-id number;
}
interface-shared-with {
    psd-name;
}
native-inner-vlan-id value;
(no-traps | traps);
output-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        no-rfc2486;
        passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
        access-profile name;
        default-password password;
        local-name name;
        local-password password;
        no-rfc2486;
        passive;
    }
}
proxy-arp (restricted | unrestricted);
radio-router {
    bandwidth number;
    credit {
        interval number;
    }
    data-rate number;
    latency number;
    quality number;
    resource number;
    threshold number;
}
swap-by-poppush;
traps;
vlan-id vlan-id;
vlan-id-range vlan-id-range;
vlan members [vlan-id];
vlan-id-range vlan-id1-vlan-id2;
```


Copyright © 2016, Juniper Networks, Inc.

```
        vlan-tags {
          (inner vlan-id | inner-range vlan-id1-vlan-id2);
          inner-list [vlan-id];
          outer vlan-id;
        }
      }
      vlan-tagging;
    }
  }
```

Related Documentation

- Understanding Interfaces on page 2441

Policy-Options Configuration Statement Hierarchy

Use the statements in the **policy-options** configuration hierarchy to configure routing policies that control information from routing protocols that the device imports into its routing table and exports to its neighbors.

```
policy-options {
  application-maps application-map-name{
    application application-name;
      code-points [value ];
    }
  }
  as-path aspath-name {
    as-path-regular-expression;
    dynamic-db ;
  }
  as-path-group group-name {
    as-path as-path-name {
      as-path-regular-expression;;
    }
    dynamic-db;
  }
  community community-name {
    dynamic-db;
    invert-match;
    members [value];
  }
  condition condition-name {
    dynamic-db;
    if-route-exists {
      address;
      table table-name;
    }
    route-active-on (node0 | node1);
  }
  damping name {
    disable;
    half-life minutes;
    max-supress minutes;
    reuse number;
    suppress number;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    policy-statement policy-name {
      dynamic-db;
      from match-conditions;
      term term-name;
      then action;
      to match-conditions;
    }
    prefix-list name {
      address-prefix;
      apply-path;
      dynamic-db;
    }
    vsi-policy name {
      from {
        vsi-manager vsi-manager-id {
          vsi-instance vsi-instance;
          vsi-type number;
          vsi-version number;
        }
      }
      then {
        filter filter name;
      }
    }
}
```

Related Documentation

- MPLS Overview on page 3743

[edit security address-book] Hierarchy Level

```
security {
  address-book (book-name | global) {
    address address-name {
      ip-prefix {
        description text;
      }
      description text;
      dns-name domain-name {
        ipv4-only;
        ipv6-only;
      }
      range-address lower-limit to upper-limit;
      wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
      address address-name;
      address-set address-set-name;
      description text;
    }
    attach {
      zone zone-name;
    }
    description text;
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
```

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- Understanding Address Books on page 1069

[edit security forwarding-options] Hierarchy Level

```
security {
  forwarding-options {
    family {
      inet6 {
        mode (drop | flow-based | packet-based);
      }
      iso {
        mode packet-based;
      }
      mpls {
        mode packet-based;
      }
    }
  }
  mirror-filter filter-name {
    destination-port port-number;
    destination-prefix destination-prefix;
    interface-in interface-name;
    interface-out interface-name;
    output {
      destination-mac mac-address;
      interface interface-name;
    }
    protocol protocol;
    source-port port-number;
    source-prefix source-prefix;
  }
  secure-wire secure-wire-name interface [interface-name-1 interface-name-2];
}
```

Related Documentation
- Security Configuration Statement Hierarchy on page 603

[edit security ike] Hierarchy Level

```
security {
  ike {
    gateway gateway-name {
      address [ip-address-or-hostname];
      dead-peer-detection {
        (always-send | optimized | probe-idle-tunnel);
        interval seconds;
        threshold number;
      }
      dynamic {
        connections-limit number;
```

```
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
      ike-user-type (group-ike-id | shared-ike-id);
    }
    external-interface external-interface-name;
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
      (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address
        | user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    remote-identity {
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    version (v1-only | v2-only);
    xauth {
      access-profile profile-name;
    }
}
policy policy-name {
    certificate {
      local-certificate certificate-id;
      peer-certificate-type (pkcs7 | x509-signature);
    }
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
}
proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1);
    authentication-method (dsa-signatures | ecdsa-signatures-256 |
      ecdsa-signatures-384 | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
}
respond-bad-spi <max-responses>;
traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
      size maximum-file-size;
    }
    flag flag;
    no-remote-trace;
```


Copyright © 2016, Juniper Networks, Inc.

```
        rate-limit messages-per-second;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041

[edit security ipsec] Hierarchy Level

```
security {
  ipsec {
    internal {
      security-association {
        manual encryption {
          iked_encryption enabled;
          algorithm 3des-cbc;
          key ascii-text key;
          }
      }
    }
    policy policy-name {
      description description;
      perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24
        | group5);
      proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
      proposals [proposal-name];
    }
    proposal proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
      description description;
      encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc |
        aes-192-gcm | aes-256-cbc | aes-256-gcm | des-cbc);
      lifetime-kilobytes kilobytes;
      lifetime-seconds seconds;
      protocol (ah | esp);
    }
    security-association sa-name {
      manual {
        direction bidirectional {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key {
              ascii-text key;
              hexadecimal key;
            }
          }
          auxiliary-spi auxiliary-spi-value;
          encryption {
            algorithm (3des-cbc | des-cbc | null);
            key {
              ascii-text key;
              hexadecimal key;
            }
```

Copyright © 2016, Juniper Networks, Inc.

```
            }
            protocol (ah | esp);
            spi spi-value;
          }
        }
        mode transport;
      }
      traceoptions {
        flag flag;
      }
      vpn vpn-name {
        bind-interface interface-name;
        copy-outer-dscp;
        df-bit (clear | copy | set);
        establish-tunnels (immediately | on-traffic);
        ike {
          gateway gateway-name;
          idle-time seconds;
          install-interval seconds;
          ipsec-policy ipsec-policy-name;
          no-anti-replay;
          proxy-identity {
            local ip-prefix;
            remote ip-prefix;
            service (any | service-name);
          }
        }
        manual {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          encryption {
            algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
            key (ascii-text key | hexadecimal key);
          }
          external-interface external-interface-name;
          gateway ip-address;
          protocol (ah | esp);
          spi spi-value;
        }
        traffic-selector traffic-selector-name {
          local-ip ip-address/netmask;
          remote-ip ip-address/netmask;
        }
        vpn-monitor {
          destination-ip ip-address;
          optimized;
          source-interface interface-name;
        }
      }
      vpn-monitor-options {
        interval seconds;
        threshold number;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          destination-address-excluded;
          source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          source-address-excluded;
          source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
          }
        }
        scheduler-name scheduler-name;
        then {
          count {
            alarm {
              per-minute-threshold number;
              per-second-threshold number;
            }
          }
          deny;
          log {
            session-close;
            session-init;
```

```
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
global {
   policy policy-name {
      description description;
      match {
         application {
            [application];
            any;
         }
         destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
         }
         from-zone {
            [zone-name];
            any;
         }
         source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
         }
         source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
         }
         to-zone {
            [zone-name];
            any;
         }
      }
      scheduler-name scheduler-name;
      then {
         count {
            alarm {
               per-minute-threshold number;
               per-second-threshold number;
            }
         }
         deny;
         log {
            session-close;
            session-init;
         }
         permit {
            application-services {
               application-firewall {
                  rule-set rule-set-name;
               }
               application-traffic-control {
```

```
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
          redirect-wx | reverse-redirect-wx;
          ssl-proxy {
            profile-name profile-name;
          }
          uac-policy {
            captive-portal captive-portal;
          }
          utm-policy policy-name;
        }
        destination-address {
          drop-translated;
          drop-untranslated;
        }
        firewall-authentication {
          pass-through {
            access-profile profile-name;
            client-match user-or-group-name;
            ssl-termination-profile profile-name;
            web-redirect;
            web-redirect-to-https;
          }
          user-firewall {
            access-profile profile-name;
            domain domain-name
            ssl-termination-profile profile-name;
          }
          web-authentication {
            client-match user-or-group-name;
          }
        }
        services-offload;
        tcp-options {
          initial-tcp-mss mss-value;
          reverse-tcp-mss mss-value;
          sequence-check-required;
          syn-check-required;
        }
      }
      reject;
    }
  }
}
policy-rematch;
policy-stats {
  system-wide (disable | enable);
}
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
```


Copyright © 2016, Juniper Networks, Inc.

```
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

condition (Policy Options)

Syntax

```
condition condition-name {
  dynamic-db;
  if-route-exists {
    address;
    table table-name;
  }
  route-active-on (node0 | node1);
}
```

Hierarchy Level [edit policy-options]

Release Information Statement introduced in Junos OS Release 9.0.

Description For chassis cluster configurations, specify the match condition for use in routing to a redundant Ethernet (**reth**) interface.

Options ***condition-name*** —Name of the routing policy match condition.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743

family (Security Forwarding Options)

Syntax

```
family {
  inet6 {
    mode (drop | flow-based | packet-based);
  }
  iso {
    mode packet-based;
  }
  mpls {
    mode packet-based;
  }
}
```

Hierarchy Level [edit security forwarding-options]

Release Information Statement introduced in Junos OS Release 8.5 .

Description Determine the protocol family to be used for packet forwarding.

NOTE: Packet-based processing is not supported on the following SRX Series devices: SRX5400, SRX5600, and SRX5800.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743

Copyright © 2016, Juniper Networks, Inc.

flow-server (Forwarding Options)

Syntax

```
flow-server ip-address-or-host-name {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  source-address ip-address;
  version (5 | 500 |8);
  version9 {
    template template-name;
  }
}
```

Hierarchy Level

[edit forwarding-options sampling family inet output]
[edit forwarding-options sampling family inet6 output]

Release Information

Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.

Description

Configure sending traffic aggregates in cflowd format.

Options

- **aggregation (version 8 only**—Aggregations to perform for exported flows.
- **autonomous-system-type**—Type of autonomous system number to export.
- **local-dump**—Dump cflowd records to log file before exporting.
- **no-local-dump**—Do not dump cflowd records to log file before exporting.
- **port**—UDP port number on host collecting cflowd packets.
- **source-address**—Source IPv4/IPv6 address for cflowd packets.

NOTE: The flow server may be IPv4 or IPv6 when you configure the collector IP address using the **set forwarding-options sampling family inet flow server** command, but only IPv4 sampling is achieved.

The flow server may be IPv4 or IPv6 when you configure the collector IP address using the **set forwarding-options sampling family inet6 flow server** command, but only IPv6 sampling is achieved.

- **version**—Format of exported cflowd aggregates.

Copyright © 2016, Juniper Networks, Inc.

- **version9**—Exported data in version 9 format.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Configuring the inet6 IPv6 Protocol Family on page 2459
- Enabling Flow-Based Processing for IPv6 Traffic on page 2460

Copyright © 2016, Juniper Networks, Inc.

forwarding-options (Security)

Syntax

```
forwarding-options {
  family {
    inet6 {
      mode (drop | flow-based | packet-based);
    }
    iso {
      mode packet-based;
    }
    mpls {
      mode packet-based;
    }
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 8.5 .

Description Determine how the **inet6, iso**, and **mpls** protocol families manage security forwarding options.

NOTE:

- Packet-based processing is not supported on the following SRX Series devices: SRX5400, SRX5600, and SRX5800.
- On SRX Series devices, the default mode for processing traffic is flow mode. To configure an SRX Series device as a border router, you must change the mode from flow-based processing to packet-based processing. Use the **set security forwarding-options family mpls mode packet-based** statement to configure the SRX device to packet mode. You must reboot the device for the configuration to take effect.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743
- *Understanding External BGP Peering Sessions*
- Understanding Packet-Based Processing on page 1803
- Juniper Networks Devices Processing Overview on page 1665

Copyright © 2016, Juniper Networks, Inc.

fragment

Syntax fragment;

Hierarchy Level [edit security screen ids-option *screen-name* icmp]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the device to detect and drop any ICMP frame with the More Fragments flag set or with an offset indicated in the **offset** field.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743

Copyright © 2016, Juniper Networks, Inc.

hash-key (Forwarding Options)

Syntax

```
hash-key {
  family inet {
    layer-3;
    layer-4;
    session-id;
  }
  family mpls {
    label-1;
    label 2;
    label-3;
    no-labels;
    payload {
      ip {
        layer-3-only;
        port-data {
          destination-lsb;
          destination-msb;
          source-lsb;
          source-msb;
        }
      }
    }
  }
  family multiservice {
    destination-mac;
    source-mac;
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement modified in Junos OS Release 10.2.

Description Select which packet header data to use for per-flow load balancing.

Options

- **inet**—IPv4 protocol family.
- **mpls**—MPLS protocol family.
- **layer-3**—Incorporate Layer 3 data into the hash key.
- **layer-4**—Incorporate Layer 4 data into the hash key.
- **session-id**—Incorporate session ID data into the hash key (SRX3000 and SRX5000 lines only). The session ID data has higher precedence than the Layer 3 or 4 information.
- **label-1**—Incorporate the first MPLS label into the hash key.
- **label-2**—Incorporate the second MPLS label into the hash key.
- **label-3**—Incorporate the third MPLS label into the hash key.
- **no-labels**—Include no MPLS labels into the hash key.

Copyright © 2016, Juniper Networks, Inc.

- **payload**—Incorporate payload data into the hash key.
- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3–only**—Include only Layer 3 IP information.
- **port-data**—Include the source and destination port field information.
- **source-msb**—Include the most significant byte of the source port.
- **source-lsb**—Include the least significant byte of the source port.
- **destination-msb**—Include the most significant byte of the destination port.
- **destinatione-lsb**—Include the least significant byte of the destination port.
- **source-mac**—Include source MAC address in hash key.
- **destinatione-mac**—Include destination MAC address in hash key.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- MPLS Overview on page 3743

iso (Security Forwarding Options)

Syntax

```
iso {
  mode packet-based;
}
```

Hierarchy Level
[edit security forwarding-options family]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Enable the forwarding of IS-IS traffic. By default, the device drops IS-IS traffic.

NOTE: Junos OS security processing is not applied to IS-IS packets forwarded by the device.

NOTE: Packet-based processing is not supported on the following SRX Series devices: SRX1500, SRX5600, and SRX5800.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security forwarding-options] Hierarchy Level on page 3006

Copyright © 2016, Juniper Networks, Inc.

mpls (Security Forwarding Options)

Syntax

```
mpls {
  mode packet-based;
}
```

Hierarchy Level [edit security forwarding-options family]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable the forwarding of MPLS traffic. By default, the device drops MPLS traffic.

CAUTION: Because MPLS operates in packet mode, security services are not available.

NOTE: Packet-based processing is not supported on the following SRX Series devices: SRX1500, SRX5600, and SRX5800.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743

multicast-scope

Syntax multicast-scope (*scope-value* | global | link-local | node-local | organization-local | orhigher | orlower | site-local);

Hierarchy Level [edit policy-options policy-statement *policy-name* from]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multicast scoping to match the routing policy.

Options

- ***scope-value* <orhigher | orlower>**—The multicast-scope value is a number from 0 through 15.
- **global <orhigher | orlower>**—Global multicast scope
- **link-local <orhigher | orlower>**—Link-local scope
- **node-local <orhigher | orlower>**—Node-local scope
- **organization-local <orhigher | orlower>**—Organizational-local scope
- **orhigher**—Match on numerically higher scopes
- **orlower**—Match on numerically lower scopes
- **site-local <orhigher | orlower>**—Site-local values

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- MPLS Overview on page 3743

 Copyright © 2016, Juniper Networks, Inc.

policer (Firewall)

Syntax

```
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-interface-policer;
  then {
    discard;
    forwarding-class forwarding-class-name;
    loss-priority (high | low |medium-high |medium-low);
    out-of-profile;
  }
}
```

Hierarchy Level [edit firewall]

Release Information Command introduced in Junos OS Release 9.5.

Description Configure policer rate limits and actions. To activate a policer, you must include the policer action modifier in the **then** statement in a firewall filter term or on an interface.

Options

- ***policer-name***—Name of the policer to evaluate when packets are received on the interface
- **bandwidth-limit *bps***—Specify the bandwidth limit as a number of bits per second
- **bandwidth-percent *percentage***—Specify the bandwidth limit in percentage value
- **burst-size-limit *bytes***—Specify the burst size limit as a number of bytes
- **filter-specific**— Specify that policer is filter-specific
- **logical-interface-policer**— Specify that policer is logical interface policer
- **discard**—Always discard non conforming red packets.
- **forwarding-class *classname***—Specify the particular forwarding class
- **loss-priority**—Set the loss priority to high or low
- **out-of-profile**— Discard packets only if both congested and over threshold

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Firewall Configuration Statement Hierarchy on page 3919

Copyright © 2016, Juniper Networks, Inc.

simple-filter (Firewall)

Syntax

```
simple-filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      (accept | discard);
      forwarding-class class-name;
      policer policer-name;
      three-color-policer policer-name {
        (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
      }
    }
  }
}
```

Hierarchy Level [edit firewall family *family-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Define a simple filter. Simple filters are recommended for metropolitan Ethernet applications.

Options

- **from**—Match packet fields to values. If the **from** option is not included, all packets are considered to match and the actions and action modifiers in the **then** statement are taken.
- **match-conditions**—One or more conditions to use to make a match.
- ***term-name***—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include space in the name, enclose it in quotation marks (“ ”).
- **then**—Actions to take on matching packets. If the then option is not included and a packet matches all the conditions in the from statement, the packet is accepted.

NOTE: On SRX1400, SRX3400, and SRX3600 devices, the Forwarding class as match condition feature is not supported by a simple filter.

NOTE: SRX3400 and SRX3600 devices have the following limitations of a simple filter:

- The forwarding class is the match condition.
- In the packet processor on an IOC, up to 400 logical interfaces can be applied with simple filters.

Copyright © 2016, Juniper Networks, Inc.

- In the packet processor on an IOC, the maximum number of terms of all simple filters is 2000.
- In the packet processor on an IOC, the maximum number of policers is 2000
- In the packet processor on an IOC, the maximum number of three-color-policers is 2000
- The maximum burst size of a policer or three-color-policer is 16 MB.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Firewall Configuration Statement Hierarchy on page 3919

Copyright © 2016, Juniper Networks, Inc.

template (Flow Monitoring)

Syntax

```
template template-name {
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  ipv4-template;
  ipv6-template;
  option-refresh-rate {
    packets packets;
    seconds seconds;
  }
  template-refresh-rate {
    packets packets;
    seconds seconds;
  }
}
```

Hierarchy Level [edit services flow-monitoring version9]

Release Information Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.

Description Specify one or more version 9 templates.

Options

- **flow-active-timeout**—Interval after which active flow is exported. The range is from 10 through 600. The default value is 60.
- **flow-inactive-timeout**—Period of inactivity that marks a flow inactive. The range is from 10 through 600. The default value is 60.
- **ipv4-template**—IPv4 template configuration.
- **ipv6-template**—IPv6 template configuration.
- **option-refresh-rate**—Rate at which the device sends options. The range is from 1 through 480,000. The default value is 4800.
  - **packets**—Specify the number of packets. The range is from 1 through 480,000.
  - **seconds**—Specify the number of seconds. The range is from 10 through 600.
- **template-refresh-rate**—Rate at which the device sends template definitions. The range is from 1 through 480,000. The default value is 4800.
  - **packets**—Specify the number of packets. The range is from 1 through 480,000.
  - **seconds**—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- version9 (Flow Server) on page 3984

 Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Flow)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
  packet-filter filter-name {
    destination-port port-identifier;
    destination-prefix address;
    interface interface-name;
    protocol protocol-identifier;
    source-port port-identifier;
    source-prefix address;
  }
  rate-limit messages-per-second;
  trace-level (brief | detail | error);
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 12.1X46-D10 with the **trace-level** option and additional flags.

Description Configure flow tracing options.

Options **file**—Configure the trace file options.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

files *number*—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10 files

match *regular-expression*—Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size,

Copyright © 2016, Juniper Networks, Inc.

trace-file*.0** is renamed ***trace-file*.1** and ***trace-file is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

Range: 0 KB through 1 GB

Default: 128 KB

Copyright © 2016, Juniper Networks, Inc.

world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

all—Trace with all flags enabled

basic-datapath—Trace basic packet flow activity

fragmentation—Trace IP fragmentation and reassembly events

high-availability—Trace flow high-availability information

host-traffic—Trace flow host traffic information

multicast—Trace multicast flow information

route—Trace route lookup information

session—Trace session creation and deletion events

session-scan—Trace session scan information

tcp-basic—Trace TCP packet flow information

tunnel—Trace tunnel information

no-remote-trace—Set remote tracing as disabled.

packet-filter ***filter-name***—Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port ***port-identifier***—Match TCP/UDP destination port

destination-prefix ***address***—Destination IP address prefix

interface ***interface-name***—Logical interface

protocol ***protocol-identifier***—Match IP protocol type

source-port ***port-identifier***—Match TCP/UDP source port

source-prefix ***address***—Source IP address prefix

rate-limit ***messages-per-second***—Limit the incoming rate of trace messages.

Copyright © 2016, Juniper Networks, Inc.

trace-level—Set the level for trace logging. This option is available only when the flag is set.

brief—Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.

detail—Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.

error—Trace error information, such as system failure, unknown message type, and packet drop.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

version9 (Flow Server)

Syntax

```
version9 {
  template template-name;
}
```

Hierarchy Level
[edit forwarding-options sampling family inet output flow-server *ip-address*]
[edit forwarding-options sampling family inet6 output flow-server *ip-address*]

Release Information
Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.

Description
Export data in version 9 format.

Options
- **apply-groups**—Groups from which to inherit configuration data.
- **apply-groups-except**—Do not inherit configuration data from these groups.
- **template**—Template configuration.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- forwarding-options (Security) on page 3971

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 174

Operational Commands

- show bgp neighbor (View)
- show interfaces flow-statistics
- show interfaces statistics (View)
- show security flow status
- show security ipsec security-associations
- show security ipsec statistics

Copyright © 2016, Juniper Networks, Inc.

show bgp neighbor (View)

Syntax **show bgp neighbor**
< *neighbor-address*>
<**instance** *instance* >

Release Information Command modified in Junos OS Release 8.5.

Description Display the state of the specified Border Gateway Protocol (BGP) neighbor. If a peer is forced to Idle state because of license check failure, the output displays the state and the reason—**LicenseCheckFailed**.

Options
- **instance** ***instance***—(Optional) Display peer information for a particular routing instance.
- ***neighbor-address***—(Optional) Display information for only the BGP peer at the specified IP address.

Required Privilege Level view

Related Documentation
- MPLS Overview on page 3743

List of Sample Output
show bgp neighbor 5.5.5.2 on page 3988
show bgp neighbor instance master on page 3989

Output Fields Table 374 lists the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 374: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor's port number.
AS	AS number of the peer.
Local	Address of the local device. The address is followed by the peer's port number.
Type	Type of peer: Internal or External.
State	Current state of the BGP session: • **Active**—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message. • **Connect**—BGP is waiting for the transport protocol connection to complete. • **Established**—The BGP session has been established, and the peers are exchanging update messages. • **Idle**—Either the BGP license check failed, or this is the first stage of a connection and BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer.

Copyright © 2016, Juniper Networks, Inc.

Table 374: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Flags	Internal BGP flags: • **Aggregate Label**—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • **CleanUp**—The peer session is being shut down. • **Delete**—This peer has been deleted. • **Idled**—This peer has been permanently idled. • **Initializing**—The peer session is initializing. • **SendRtn**—Messages are being sent to the peer. • **Sync**—This peer is synchronized with the rest of the peer group. • **TryConnect**—Another attempt is being made to connect to the peer. • **Unconfigured**—This peer is not configured. • **WriteFailed**—An attempt to write to this peer failed. • **Last State**—Previous state of the BGP session.
Last State	Previous state of the BGP session.
Last Event	Last activity that occurred in the BGP session: • **Closed**—The BGP session closed. • **ConnectRetry**—The transport protocol connection failed, and BGP is trying again to connect. • **HoldTime**—The session ended because the hold timer expired. • **KeepAlive**—The local device sent a BGP keepalive message to the peer. • **Open**—The local device sent a BGP open message to the peer. • **OpenFail**—The local device did not receive an acknowledgment of a BGP open message from the peer. • **RecvKeepAlive**—The local device received a BGP keepalive message from the peer. • **RecvNotify**—The local device received a BGP notification message from the peer. • **RecvOpen**—The local device received a BGP open message from the peer. • **RecvUpdate**—The local device received a BGP update message from the peer. • **Start**—The peering session started. • **Stop**—The peering session stopped. • **TransportError**—A TCP error occurred.
Last Error	Last error that occurred in the BGP session: • **Cease**—An error occurred, such as a version mismatch, that caused the session to close. • **Finite State Machine Error**—In setting up the session, BGP received a message that it did not understand. • **Hold Time Expired**—The session's hold time expired. • **Message Header Error**—The header of a BGP message was malformed. • **Open Message Error**—A BGP open message contained an error. • **None**—No errors occurred in the BGP session. • **Update Message Error**—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Copyright © 2016, Juniper Networks, Inc.

Table 374: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Options	Configured BGP options: • **AddressFamily**—Configured address family: inet or inet-vpn. • **GracefulRestart**—Graceful restart is configured. • **HoldTime**—Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent. • **Local Address**—Address configured with the **local-address** statement. • **NLRI**—Configured multicast BGP state for the BGP group: multicast, unicast, or both if you have configured **nlri any**. • **Peer AS**—Configured peer autonomous system (AS). • **Preference**—Preference value configured with the **preference** statement. • **Refresh**—Configured to refresh automatically when the policy changes. • **Rib-group**—Configured routing table group.
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local device.
Holdtime	Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent.
Preference	Preference value configured with the **preference** statement.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Trace file	Name of the file to receive the output of the tracing operation.

Sample Output

show bgp neighbor 5.5.5.2

```
user@host> show bgp neighbor 5.5.5.2
  Type: Internal    State: Idle (LicenseCheckFailed)   Flags: <ImportEval Sync>
Peer: 5.5.5.2 AS 200          Local: unspecified AS 200
  Type: Internal    State: Idle (LIcenseCheckFailed)       (route reflector
client)Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Options: <Preference LogUpDown Cluster AddressFamily PeerAS Rib-group Refresh>

  Address families configured: inet-unicast inet-vpn-unicast l2vpn-signaling
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Trace options:  all
  Trace file: /var/log/bgp size 131072 files 10
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show bgp neighbor instance master

```
user@host> show bgp neighbor instance master
Peer: 5.5.5.1 AS 200           Local: 5.5.5.2 AS 200
  Type: Internal    State: Idle (LicenseCheckFailed)         Flags: <>
  Last State: Idle          Last Event: Start
  Last Error: Cease
  Export: [ static ]
  Options: <Preference LocalAddress LogUpDown AddressFamily PeerAS Rib-group
Refresh>
  Address families configured: inet-unicast inet-vpn-unicast
  Local Address: 5.5.5.2 Holdtime: 90 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Update Message Error' Sent: 3 Recv: 0
  Error: 'Cease' Sent: 2 Recv: 0
  Trace file: /var/log/bgp size 131072 files 10
```

Copyright © 2016, Juniper Networks, Inc.

show interfaces flow-statistics

Syntax **show interfaces flow-statistics** *<interface-name>*

Release Information Command introduced in Junos OS Release 9.2.

Description Display interfaces flow statistics.

Options *Interface-name*—(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the "Interface Naming Conventions" on page 2446.

- **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
- **br-*pim*/0/*port***—Basic Rate Interface for establishing ISDN connections.
- **ce1-*pim*/0/*port***—Channelized E1 interface.
- **ct1-*pim*/0/*port***—Channelized T1 interface.
- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
- **e1-*pim*/0/*port***—E1 interface.
- **e3-*pim*/0/*port***—E3 interface.
- **fe-*pim*/0/ *port***—Fast Ethernet interface.
- **ge-*pim*/0/*port***—Gigabit Ethernet interface.
- **se-*pim*/0/*port***—Serial interface.
- **t1-*pim*/0/*port***—T1 (also called DS1) interface.
- **t3-*pim*/0/ *port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665
- Understanding Interfaces on page 2441

List of Sample Output **show interfaces flow-statistics (Gigabit Ethernet) on page 3993**

Output Fields Table 196 lists the output fields for the **show interfaces flow-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 375: show interfaces flow-statistics Output Fields

Field Name	Field Description
Traffic statistics	Number of packets and bytes transmitted and received on the physical interface.

Copyright © 2016, Juniper Networks, Inc.

Table 375: show interfaces flow-statistics Output Fields *(continued)*

Field Name	Field Description
Local statistics	Number of packets and bytes transmitted and received on the physical interface.
Transit statistics	Number of packets and bytes transiting the physical interface.
Flow input statistics	Statistics on packets received by flow module.
Flow output statistics	Statistics on packets sent by flow module.
Flow error statistics	Packet drop statistics for the flow module. For further details, see Table 197.

Table 376: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.
Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.
User authentication errors	Packet was dropped because policy requires authentication; however: • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.

Copyright © 2016, Juniper Networks, Inc.

Table 376: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Flow:	
No one interested in self packets	This counter is incremented for one of the following reasons: • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	The packet was dropped because a valid route was not available to forward the packet. For new sessions, the counter is incremented for one of the following reasons: • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because: • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found
No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.
Policy denied	The error counter is incremented for one of the following reasons: • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.
TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-

Copyright © 2016, Juniper Networks, Inc.

Table 376: Flow Error Statistics (Packet Drop Statistics for the Flow Module) *(continued)*

Invalid zone received packet	-
No NAT gate	-

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp igmp ldp msdp nhrp
ospf pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     2564
      Bytes permitted by policy :       3478
      Connections established :         1
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 0
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
```

Copyright © 2016, Juniper Networks, Inc.

```
Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```


Copyright © 2016, Juniper Networks, Inc.

show interfaces statistics (View)

Syntax show interfaces statistics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level view

Related Documentation
- Understanding Interfaces on page 2441

List of Sample Output show interfaces statistics on page 3995

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
     Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
     Input packets : 2743333
     Output packets: 6790470992
     Security: Zone: untrust
     Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset
 http https ike netconf ping reverse-telnet
    reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl lsping ntp sip
     Protocol inet, MTU: 9192
       Addresses, Flags: Is-Preferred Is-Primary
         Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

show security flow status

Syntax **show security flow status**

Release Information Command introduced in Junos OS Release 10.2; session distribution mode option added in Junos OS Release 12.1X44-D10; enhanced route scaling mode option added in Junos OS Release 12.1X45-D10. GTP-U distribution option added in Junos OS Release 15.1X49-D40.
Starting with Junos OS Release 15.1X49-D10, SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) are introduced for SRX5400, SRX5600, and SRX5800 devices that perform hash-based data path packet forwarding to interconnect with all existing IOC and SPC cards using the XL chip (packet-processing chip).
The IOC3 XL chip uses a hash-based method to distribute ingress traffic to a pool of SPUs by default. Selection of hash keys depends on application protocols.

Description Display the flow processing modes and logging status.

Required Privilege Level view

Related Documentation
- Juniper Networks Devices Processing Overview on page 1665

List of Sample Output
show security flow status on page 3997
show security flow status (IPsec Performance Acceleration) on page 3997
show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3) on page 3998

Output Fields Table 234 lists the output fields for the **show security flow status** command. Output fields are listed in the approximate order in which they appear.

Table 377: show security flow status Output Fields

Field Name	Field Description
Flow forwarding mode	Flow processing mode. • Inet forwarding mode • Inet6 forwarding mode • MPLS forwarding mode • ISO forwarding mode • Session distribution mode • Enhanced route scaling mode
Flow trace status	Flow logging status. • Flow tracing status • Flow tracing options

Copyright © 2016, Juniper Networks, Inc.

Table 377: show security flow status Output Fields *(continued)*

Field Name	Field Description
flow session distribution	SPU load distribution mode. • RR-based • Hash-based GTP-U distribution • Enabled
Flow packet ordering	packet-ordering mode. • Hardware • Software
Flow ipsec performance acceleration	IPsec VPN performance acceleration status.

Sample Output

show security flow status

```
root> show security flow status
Flow forwarding mode:
Inet forwarding mode: flow based
Inet6 forwarding mode: flow based
MPLS forwarding mode: drop
ISO forwarding mode: drop
Enhanced route scaling mode: Enabled (reboot needed to disable)
Flow trace status
Flow tracing status: on
Flow tracing options: all
Flow session distribution
Distribution mode: Hash-based
GTP-U distribution: Enabled
Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
```

show security flow status (IPsec Performance Acceleration)

```
root> show security flow status
Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
Flow trace status
    Flow tracing status: off
Flow session distribution
    Distribution mode: RR-based
    GTP-U distribution: Enabled Flow packet ordering
Ordering mode: Software (reboot needed to change to software)
Flow ipsec performance acceleration: on
```

Copyright © 2016, Juniper Networks, Inc.

show security flow status (for hash-based datapath forwarding using SRX5K-MPC3-40G10G (IOC3) and SRX5K-MPC3-100G10G (IOC3)

```
root> show security flow status
node0:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware

node1:
--------------------------------------------------------------------------
  Flow forwarding mode:
    Inet forwarding mode: flow based
    Inet6 forwarding mode: drop
    MPLS forwarding mode: drop
    ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
    GTP-U distribution: Enabled
  Flow ipsec performance acceleration: off
  Flow packet ordering
    Ordering mode: Hardware
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec security-associations

Syntax

show security ipsec security-associations
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number>*
sa-type shortcut
vpn-name *vpn-name* <traffic-selector *traffic-selector-name*>

Release Information

Command introduced in Junos OS Release 8.5. Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for the **vpn-name** option added in Junos OS Release 11.4R3. Support for the **traffic-selector** option and traffic selector field added in Junos OS Release 12.1X46-D10. Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10.

Description

Display information about the IPsec security associations (SAs).

Options

- none—Display information about all SAs.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **family**—(Optional) Display SAs by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc** ***slot-number***—(Optional) Display information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index** ***SA-index-number***—(Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- **kmd-instance**—(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number***—(Optional) Display information about existing IPsec SAs in this PIC slot. This option is used to filter the output.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.
- **vpn-name** ***vpn-name***—Name of the VPN. If configured, **traffic-selector** ***traffic-selector-name*** can optionally be specified.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- clear security ipsec security-associations on page 6986
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

List of Sample Output
show security ipsec security-associations (IPv4) on page 4003
show security ipsec security-associations (IPv6) on page 4003
show security ipsec security-associations index 131073 on page 4003
show security ipsec security-associations brief on page 4004
show security ipsec security-associations detail on page 4004
show security ipsec security-associations family inet6 on page 4005
show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 4005
show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel) on page 4005
show security ike sa index 222075191 detail on page 4006
show security ipsec security-associations detail (ADVPN Partner, Static Tunnel) on page 4007
show security ike sa index 788674 detail on page 4008
show security ipsec security-associations sa-type shortcut (ADVPN) on page 4008
show security ipsec security-associations sa-type shortcut detail (ADVPN) on page 4008
show security ipsec security-associations family inet detail on page 4009

Output Fields Table 363 lists the output fields for the **show security ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 378: show security ipsec security-associations

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.
ID	Index number of the SA. You can use this number to get additional information about the SA.
VPN name	IPsec name for VPN.
Gateway	IP address of the remote gateway.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are **hmac-md5-95**, **hmac-sha1-96,** or **ESP**. • An encryption algorithm used to encrypt data traffic. Options are **3des-cbc, aes-128-cbc, aes-192-cbc**, **aes-256-cbc**, or **des-cbc.**

Copyright © 2016, Juniper Networks, Inc.

Table 378: show security ipsec security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
Sta	State has two options, **Installed** and **Not Installed**. • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.
Mon	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays **U** (up) or **D** (down). A hyphen (-) means VPN monitoring is not enabled for this SA.
vsys or Virtual-system	The root system.
Tunnel index	Numeric identifier of the specific IPsec tunnel for the SA.
Local gateway	Gateway address of the local system.
Remote gateway	Gateway address of the remote system.
Traffic selector	Name of the traffic selector.
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote identity	IP address of the destination peer gateway.
DF-bit	State of the don't fragment bit: **set** or **cleared**.
Policy-name	Name of the applicable policy.
Location	**FPC**—Flexible PIC Concentrator (FPC) slot number. **PIC**—PIC slot number. **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance.
Tunnel events	Tunnel event and the number of times the event has occurred. See "Tunnel Events" on page 6797 for descriptions of tunnel events and the action you can take.
Direction	Direction of the SA; it can be inbound or outbound.

Copyright © 2016, Juniper Networks, Inc.

Table 378: show security ipsec security-associations *(continued)*

Field Name	Field Description
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is **AH** or **ESP**, **AUX-SPI** is always 0. • When the value is **AH+ESP**, **AUX-SPI** is always a positive integer.
Mode	Mode of the SA: • **transport**—Protects host-to-host connections. • **tunnel—**Protects connections between security gateways.
Type	Type of the SA: • **manual**—Security parameters require no negotiation. They are static and are configured by the user. • **dynamic**—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.
State	State of the SA: • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.

Protocol	Protocol supported.
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  ◦ **Authentication**—Type of authentication used.
	  ◦ **Encryption**—Type of encryption used.

Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in seconds**—Number of seconds left until the SA expires.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • **Expires in seconds**—Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • **Expires in kilobytes**—Number of kilobytes left until the SA expires.
Anti-replay service	State of the service that prevents packets from being replayed. It can be **Enabled** or **Disabled**.

Copyright © 2016, Juniper Networks, Inc.

Table 378: show security ipsec security-associations *(continued)*

Field Name	Field Description
Replay window size	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Copy-Outer-DSCP	Indicates if copying outer IP header DSCP and ECN to inner IP header is enabled or disabled.

Sample Output

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys
  131075 11.0.28.241    500   ESP:3des/sha1   86758ff0 6918/ unlim   -   0
  131075 11.0.28.241    500   ESP:3des/sha1   3183ff26 6918/ unlim   -   0
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  131074 ESP:3des/sha1 14caf1d9 3597/ unlim   -   root 500   1212::1112
  131074 ESP:3des/sha1 9a4db486 3597/ unlim   -   root 500   1212::1112
```

show security ipsec security-associations index 131073

```
user@host> show security ipsec security-associations index 131073
  ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
```

Copyright © 2016, Juniper Networks, Inc.

```
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations brief

```
user@host> show security ipsec security-associations brief
Total active tunnels: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<16384 1.1.1.1 500 ESP:3des/sha1 af88baa 28795/unlim D 0
>16384 1.1.1.1 500 ESP:3des/sha1 f4e3e5f4 28795/unlim D 0
```

show security ipsec security-associations detail

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Version: IKEv1
DF-bit: clear
, Copy-Outer-DSCP Enabled
Bind-interface: st0.99

Port: 500, Nego#: 8, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
Tunnel events:
Mon Oct 26 2015 22:27:50 -0700: IPSec SA rekey successfully completed (7 times)
Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:41:07 -0700: IKE SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information
 updated (1 times)
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: inbound, SPI: 81ed9998, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled
```


Copyright © 2016, Juniper Networks, Inc.

```
, Replay window size: 64
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: outbound, SPI: 80565248, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```

show security ipsec security-associations family inet6

```
user@host> show security ipsec security-associations family inet6
  Virtual-system: root
  Local Gateway: 1212::1111, Remote Gateway: 1212::1112
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
  Total active tunnels: 1

ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys

<2    1.1.1.2          500   ESP:3des/sha1   67a7d25d 28280/unlim   -   0

>2    1.1.1.2          500   ESP:3des/sha1   a23cbcdc 28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear
  Bind-interface: st0.1

  Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
```

Copyright © 2016, Juniper Networks, Inc.

```
  Tunnel events:
  Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
  Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding
IKE/IPSec SAs are deleted (1 times)
  Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ike sa index 222075191 detail

```
user@host> show security ike sa index 222075191 detail
node0:
--------------------------------------------------------------------------
IKE peer 31.1.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                20474
   Output bytes  :                21091
   Input  packets:                  237
   Output packets:                  237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 31.1.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local
certificate. Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate
 not found. Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

Copyright © 2016, Juniper Networks, Inc.

show security ike sa index 788674 detail

```
user@host> show security ike sa index 788674 detail
IKE peer 11.1.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                22535
   Output bytes  :                21918
   Input  packets:                  256
   Output packets:                  256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 31.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes-cbc-256/sha1 6f164ee0 3580/ unlim - root 500 23.0.0.111
>268173318 ESP:aes-cbc-256/sha1 e6f29cb0 3580/ unlim - root 500 23.0.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
```


Copyright © 2016, Juniper Networks, Inc.

```
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
```

Copyright © 2016, Juniper Networks, Inc.

```
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1090 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec statistics

Syntax

show security ipsec statistics
<**fpc** *slot-number* >
<**index** *SA-index-number* >
<**kmd-instance** *kmd-instance-name* >
pic *slot-number*

Release Information

Command introduced in Junos OS Release 8.5. **fpc** and **pic** options added in Junos OS Release 9.3. **kmd-instance** option added in Junos OS Release 10.4.

Description

Display standard IPsec statistics.

Options

- none—Display statistics about all IPsec security associations (SAs).
- **fpc** ***slot-number*** —Specific to SRX Series devices. Display statistics about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index** ***SA-index-number*** —(Optional) Display statistics for the SA with this index number.
- **kmd-instance** ***kmd-instance-name*** —Specific to SRX Series devices. Display information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number*** —Specific to SRX Series devices. Display statistics about existing IPsec SAs in this PIC slot. This option is used to filter the output.

Required Privilege Level

view

Related Documentation

- clear security ipsec statistics on page 6987

List of Sample Output

show security ipsec statistics on page 4012
show security ipsec statistics index 5 on page 4013
show security ipsec statistics fpc 6 pic 1 (SRX Series devices) on page 4013

Output Fields

Table 379 lists the output fields for the **show security ipsec statistics** command. Output fields are listed in the approximate order in which they appear.

Table 379: show security ipsec statistics Output Fields

Field Name	Field Description
Virtual-system	The root system.

Table 379: show security ipsec statistics Output Fields *(continued)*

Field Name	Field Description
ESP Statistics	• **Encrypted bytes**—Total number of bytes encrypted by the local system across the IPsec tunnel. • **Decrypted bytes**—Total number of bytes decrypted by the local system across the IPsec tunnel. • **Encrypted packets**—Total number of packets encrypted by the local system across the IPsec tunnel. • **Decrypted packets**—Total number of packets decrypted by the local system across the IPsec tunnel.
AH Statistics	• **Input bytes**—Total number of bytes received by the local system across the IPsec tunnel. • **Output bytes**—Total number of bytes transmitted by the local system across the IPsec tunnel. • **Input packets**—Total number of packets received by the local system across the IPsec tunnel. • **Output packets**—Total number of packets transmitted by the local system across the IPsec tunnel.
Errors	• **AH authentication failures**—Total number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • **Replay errors**—Total number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • **ESP authentication failures**—Total number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • **ESP decryption failures**—total number of ESP decryption errors. • **Bad headers**—Total number of invalid headers detected. • **Bad trailers**—Total number of invalid trailers detected.

Sample Output

show security ipsec statistics

```
user@host> show security ipsec statistics
Virtual-system: Root
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security ipsec statistics index 5

```
user@host> show security ipsec statistics index 5
Virtual-system: Root
SA index: 5
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

Sample Output

show security ipsec statistics fpc 6 pic 1 (SRX Series devices)

```
user@host> show security ipsec statistics fpc 6 pic 1
ESP Statistics:
Encrypted bytes:           536408
Decrypted bytes:           696696
Encrypted packets:           1246
Decrypted packets:            888
AH Statistics:
Input bytes:                    0
Output bytes:                   0
Input packets:                  0
Output packets:                 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Multicast Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 59

Overview

- Introduction to Multicast on page 4019

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 175

Introduction to Multicast

- Multicast Overview on page 4019
- Understanding Layer 3 Multicast Functionality on the SRX5K-MPC on page 4031
- Supported IP Multicast Protocol Standards on page 4032

Multicast Overview

IP has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.

NOTE: On all SRX Series devices, reordering is not supported for multicast fragments. Reordering of unicast fragments is supported.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time. A host does not have to belong to a group to send packets to members of a group.

Routers use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

Copyright © 2016, Juniper Networks, Inc.

Comparing Multicast to Unicast

The Junos® operating system (Junos OS) routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routing devices not only for *unicast* traffic streams sent between one pair of clients and servers, but also for *multicast* traffic streams containing video, audio, or both, between a single server source and many client receivers. The routing protocols used for multicast differ in many key ways from unicast routing protocols.

Information is delivered over a network by three basic methods: unicast, broadcast, and multicast.

The differences among unicast, broadcast, and multicast can be summarized as follows:

- Unicast: One-to-one, from one source to one destination.
- Broadcast: One-to-all, from one source to all possible destinations.
- Multicast: One-to-many, from one source to multiple destinations expressing an interest in receiving the traffic.

NOTE: This list does not include a special category for many-to-many applications, such as online gaming or videoconferencing, where there are many sources for the same receiver and where receivers often double as sources. Many-to-many is a service model that repeatedly employs one-to-many multicast and therefore requires no unique protocol. The original multicast specification, RFC 1112, supports both the any-source multicast (ASM) many-to-many model and the source-specific multicast (SSM) one-to-many model.

With unicast traffic, many streams of IP packets that travel across networks flow from a single source, such as a website server, to a single destination such as a client PC. Unicast traffic is still the most common form of information transfer on networks.

Broadcast traffic flows from a single source to all possible destinations reachable on the network, which is usually a LAN. Broadcasting is the easiest way to make sure traffic reaches its destinations.

Television networks use broadcasting to distribute video and audio. Even if the television network is a cable television (CATV) system, the source signal reaches all possible destinations, which is the main reason that some channels' content is scrambled. Broadcasting is not feasible on the Internet because of the enormous amount of unnecessary information that would constantly arrive at each end user's device, the complexities and impact of scrambling, and related privacy issues.

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

Copyright © 2016, Juniper Networks, Inc.

On an IP network, because destinations (clients) do not often communicate directly with sources (servers), the routing devices between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. Multicast routing devices replicate packets received on one input interface and send the copies out on multiple output interfaces.

In IP multicast, the source and destination are almost always hosts and not routing devices. Multicast routing devices distribute the multicast traffic across the network from source to destinations. The multicast routing device must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities, but some router architectures cannot send multiple copies of packets and so do not support multicasting directly.

IP Multicast Uses

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast, originally defined as a host extension in RFC 1112 in 1989, provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

Unicast traffic is not strictly limited to data applications. Telephone conversations, wireless or not, contain digital audio samples and might contain digital photographs or even video and still flow from a single source to a single destination. In the same way, multicast traffic is not strictly limited to multimedia applications. In some data applications, the flow of traffic is from a single source to many destinations that require the packets, as in a news or stock ticker service delivered to many PCs. For this reason, the term *receiver* is preferred to *listener* for multicast destinations, although both terms are common.

Network applications that can function with unicast but are better suited for multicast include collaborative groupware, teleconferencing, periodic or “push” data delivery (stock quotes, sports scores, magazines, newspapers, and advertisements), server or website replication, and distributed interactive simulation (DIS) such as war simulations or virtual reality. Any IP network concerned with reducing network resource overhead for one-to-many or many-to-many data or multimedia applications with multiple receivers benefits from multicast.

If unicast were employed by radio or news ticker services, each radio or PC would have to have a separate traffic session for each listener or viewer at a PC (this is actually the method for some Web-based services). The processing load and bandwidth consumed by the server would increase linearly as more people “tune in” to the server. This is extremely inefficient when dealing with the global scale of the Internet. Unicast places the burden of packet duplication on the server and consumes more and more backbone bandwidth as the number of users grows.

If broadcast were employed instead, the source could generate a single IP packet stream using a broadcast destination address. Although broadcast eliminates the server packet duplication issue, this is not a good solution for IP because IP broadcasts can be sent only to a single subnetwork, and IP routing devices normally isolate IP subnetworks on

Copyright © 2016, Juniper Networks, Inc.

separate interfaces. Even if an IP packet stream could be addressed to literally go everywhere, and there were no need to “tune” to any source at all, broadcast would be extremely inefficient because of the bandwidth strain and need for uninterested hosts to discard large numbers of packets. Broadcast places the burden of packet rejection on each host and consumes the maximum amount of backbone bandwidth.

For radio station or news ticker traffic, multicast provides the most efficient and effective outcome, with none of the drawbacks and all of the advantages of the other methods. A single source of multicast packets finds its way to every *interested* receiver. As with broadcast, the transmitting host generates only a single stream of IP packets, so the load remains constant whether there is one receiver or one million. The network routing devices replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routing devices. The links leading to subnets consisting of entirely uninterested receivers carry no multicast traffic. Multicast minimizes the burden placed on sender, network, and receiver.

IP Multicast Terminology

Multicast has its own particular set of terms and acronyms that apply to IP multicast routing devices and networks. Figure 152 depicts some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *routing device*, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a *multicast routing protocol* to build a *distribution tree* that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to *sources*. In multicast terminology, the distribution tree is *rooted at the source* (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the *upstream* interface, although the less precise terms *incoming* or *inbound* interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the *downstream* interface, although the less precise terms *outgoing* or *outbound* interface are used as well. There can be 0 to *N*–1 downstream interfaces on a routing device, where *N* is the number of logical interfaces on the routing device. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.

Copyright © 2016, Juniper Networks, Inc.

Figure 152: Multicast Terminology in an IP Network

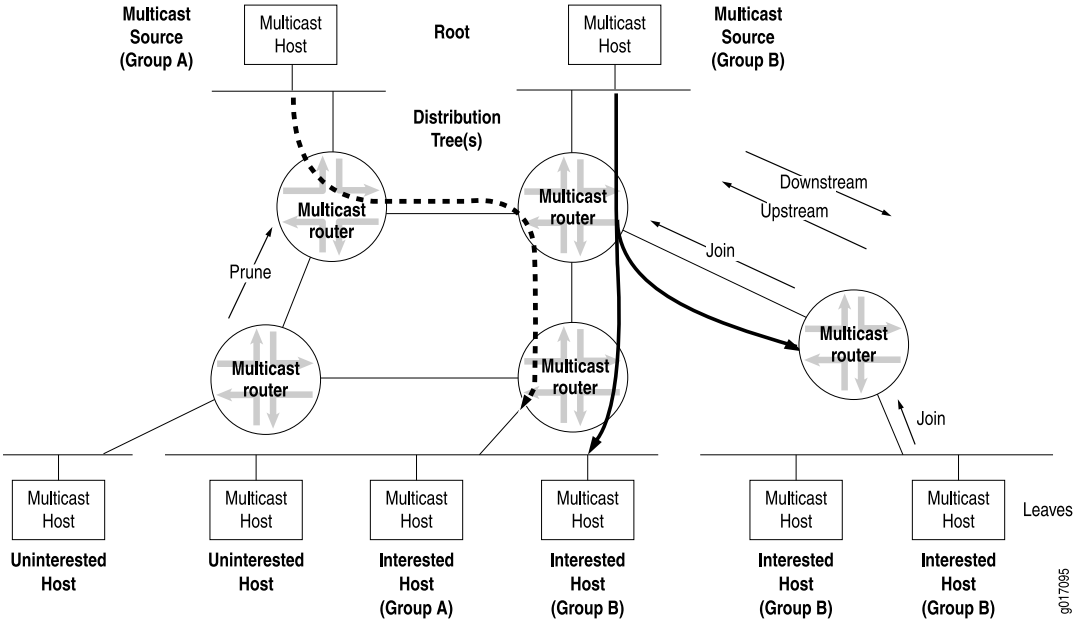

Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols.

Reverse-Path Forwarding for Loop Prevention

The routing device's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In RPF, every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the routing device verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. routing devices can build and maintain separate tables for RPF purposes.

Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast routing device operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

Copyright © 2016, Juniper Networks, Inc.

Administrative Scoping for Loop Prevention

Scoping limits the routing devices and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. routing devices at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

Multicast Leaf and Branch Terminology

Each subnetwork with hosts on the routing device that has at least one interested receiver is a *leaf* on the distribution tree. routing devices can have multiple leaves on different interfaces and must send a copy of the IP multicast packet out on each interface with a leaf. When a new leaf subnetwork is added to the tree (that is, the interface to the host subnetwork previously received no copies of the multicast packets), a new *branch* is built, the leaf is joined to the tree, and replicated packets are sent out on the interface. The number of leaves on a particular interface does not affect the routing device. The action is the same for one leaf or a hundred.

NOTE: On Juniper Networks security devices, if the maximum number of leaves on a multicast distribution tree is exceeded, multicast sessions are created up to the maximum number of leaves, and any multicast sessions that exceed the maximum number of leaves are ignored. The maximum number of leaves on a multicast distribution tree is device specific.

When a branch contains no leaves because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is *pruned* from the distribution tree, and no multicast packets are sent out that interface. Packets are replicated and sent out multiple interfaces only where the distribution tree branches at a routing device, and no link ever carries a duplicate flow of packets.

Collections of hosts all receiving the same stream of IP packets, usually from the same multicast source, are called *groups*. In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast address, or *group address*. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

IP Multicast Addressing

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Class D addresses are commonly referred to as *multicast addresses* because the entire classful address concept is obsolete. Multicast addresses can never appear as the source address in an IP packet and can only be the destination of a packet.

Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices. Blocks of multicast addresses can still be described in terms of prefix length in traditional notation, but only for convenience. For example, the multicast

Copyright © 2016, Juniper Networks, Inc.

address range from 232.0.0.0 through 232.255.255.255 can be written as 232.0.0.0/8 or 232/8.

Internet service providers (ISPs) do not typically allocate multicast addresses to their customers because multicast addresses relate to content, not to physical devices. Receivers are not assigned their own multicast addresses, but need to know the multicast address of the content. Sources need to be assigned multicast addresses only to produce the content, not to identify their place in the network. Every source and receiver still needs an ordinary, unicast IP address.

Multicast addressing most often references the receivers, and the source of multicast content is usually not even a member of the multicast group for which it produces content. If the source needs to monitor the packets it produces, monitoring can be done locally, and there is no need to make the packets traverse the network.

Many applications have been assigned a range of multicast addresses for their own use. These applications assign multicast addresses to sessions created by that application. You do not usually need to statically assign a multicast address, but you can do so.

Multicast Addresses

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address 224.0.0.0 is reserved and cannot be assigned to any group. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range from 239.0.0.0 through 239.255.255.255 is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

Layer 2 Frames and IPv4 Multicast Addresses

Multicasting on a LAN is a good place to start an investigation of multicasting at Layer 2. At Layer 2, multicast deals with media access control (MAC) frames and addresses instead of IPv4 or IPv6 packets and addresses. Consider a single LAN, without routing devices, with a multicast source sending to a certain group. The rest of the hosts are receivers interested in the multicast group's content. So the multicast source host generates packets with its unicast IP address as the source, and the multicast group address as the destination.

Which MAC addresses are used on the frame containing this packet? The packet source address—the unicast IP address of the host originating the multicast content—translates easily and directly to the MAC address of the source. But what about the packet's destination address? This is the IP multicast group address. Which destination MAC address for the frame corresponds to the packet's multicast group address?

Copyright © 2016, Juniper Networks, Inc.

One option is for LANs simply to use the LAN broadcast MAC address, which guarantees that the frame is processed by every station on the LAN. However, this procedure defeats the whole purpose of multicast, which is to limit the circulation of packets and frames to interested hosts. Also, hosts might have access to many multicast groups, which multiplies the amount of traffic to noninterested destinations. Broadcasting frames at the LAN level to support multicast groups makes no sense.

However, there is an easy way to effectively use Layer 2 frames for multicast purposes. The MAC address has a bit that is set to 0 for unicast (the LAN term is *individual address*) and set to 1 to indicate that this is a multicast address. Some of these addresses are reserved for multicast groups of specific vendors or MAC-level protocols. Internet multicast applications use the range 0x01-00-5E-00-00-00 to 0x01-00-5E-FF-FF-FF. Multicast receivers (hosts running TCP/IP) listen for frames with one of these addresses when the application joins a multicast group. The host stops listening when the application terminates or the host leaves the group at the packet layer (Layer 3).

This means that 3 bytes, or 24 bits, are available to map IPv4 multicast addresses at Layer 3 to MAC multicast addresses at Layer 2. However, all IPv4 addresses, including multicast addresses, are 32 bits long, leaving 8 IP address bits left over. Which method of mapping IPv4 multicast addresses to MAC multicast addresses minimizes the chance of “collisions” (that is, two different IP multicast groups at the packet layer mapping to the same MAC multicast address at the frame layer)?

First, it is important to realize that all IPv4 multicast addresses begin with the same 4 bits (**1110**), so there are really only 4 bits of concern, not 8. A LAN must not drop the last bits of the IPv4 address because these are almost guaranteed to be host bits, depending on the subnet mask. But the high-order bits, the leftmost address bits, are almost always network bits, and there is only one LAN (for now).

One other bit of the remaining 24 MAC address bits is reserved (an initial **0** indicates an Internet multicast address), so the 5 bits following the initial 1110 in the IPv4 address are dropped. The 23 remaining bits are mapped, one for one, into the last 23 bits of the MAC address. An example of this process is shown in Figure 153.

Copyright © 2016, Juniper Networks, Inc.

Figure 153: Converting MAC Addresses to Multicast Addresses

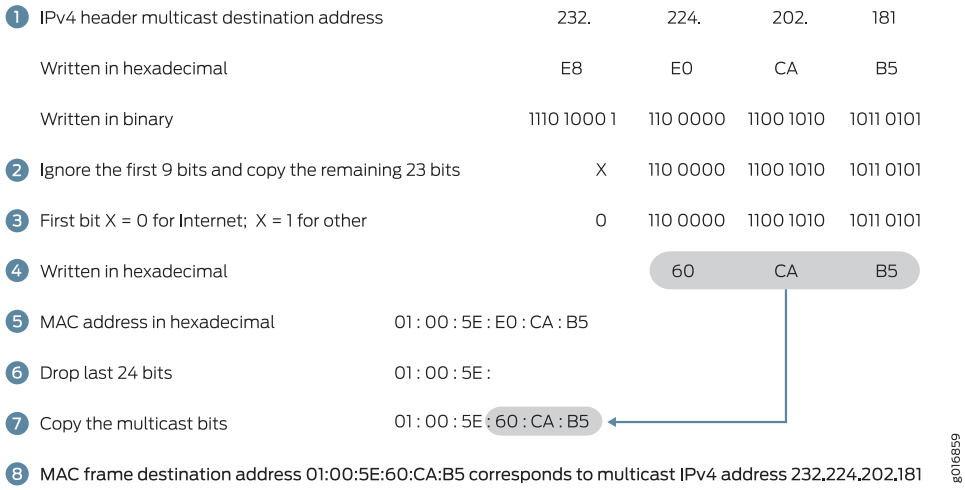

Note that this process means that there are 32 (2^5) IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses 224.8.7.6 and 229.136.7.6 translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could be interested in frames sent to both of the those multicast groups, the IP software must reject one or the other.

NOTE: This “collision” problem does not exist in IPv6 because of the way IPv6 handles multicast groups, but it is always a concern in IPv4. The procedure for placing IPv6 multicast packets inside multicast frames is nearly identical to that for IPv4, except for the MAC destination address 0x3333 prefix (and the lack of “collisions”).

Once the MAC address for the multicast group is determined, the host's operating system essentially orders the LAN interface card to join or leave the multicast group. Once joined to a multicast group, the host accepts frames sent to the multicast address as well as the host’s unicast address and ignores other multicast group’s frames. It is possible for a host to join and receive multicast content from more than one group at the same time, of course.

Multicast Interface Lists

To avoid multicast routing loops, every multicast routing device must always be aware of the interface that leads to the source of that multicast group content by the shortest path. This is the upstream (incoming) interface, and packets are never to be forwarded back toward a multicast source. All other interfaces are potential downstream (outgoing) interfaces, depending on the number of branches on the distribution tree.

routing devices closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A routing device with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's

Copyright © 2016, Juniper Networks, Inc.

content. Interfaces on the routing device's outgoing interface list send copies of the group's packets received on the incoming interface list for that group. The incoming and outgoing interface lists might be different for different multicast groups.

The multicast forwarding state in a routing device is usually written in either (S,G) or (*,G) notation. These are pronounced “ess comma gee” and “star comma gee,” respectively. In (S,G), the S refers to the unicast IP address of the source for the multicast traffic, and the G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

The asterisk (*) in the (*,G) notation is a wildcard indicating that the state applies to any multicast application source sending to group G. So, if two sources are originating exactly the same content for multicast group 224.1.1.2, a routing device could use (*,224.1.1.2) to represent the state of a routing device forwarding traffic from both sources to the group.

Multicast Routing Protocols

Multicast routing protocols enable a collection of multicast routing devices to build (join) distribution trees when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group, prune branches, locate sources and groups, and prevent routing loops.

There are several multicast routing protocols:

- Distance Vector Multicast Routing Protocol (DVMRP)—The first of the multicast routing protocols and hampered by a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G), and builds its own multicast routing tables for RPF checks.
- Multicast OSPF (MOSPF)—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routing devices do not have to flood their entire domain with multicast traffic from every source. MOSPF uses source-based distribution trees in the form (S,G).
- *Bidirectional PIM mode*—A variation of PIM. Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees as in PIM-SM and is therefore optimized for routing state size instead of path length. This means that the end-to-end latency might be longer compared to PIM sparse mode. Bidirectional PIM routes are always wildcard-source (*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM (*,G) routes forward traffic from all sources and the RP. Bidirectional PIM routing devices must have the ability to accept traffic on many potential incoming interfaces. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Bidirectional PIM is recommended in deployments with many dispersed sources and many dispersed receivers.

Copyright © 2016, Juniper Networks, Inc.

- PIM *dense mode*—In this mode of PIM, the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is *flooded* with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). This is the *dense mode* of multicast operation. LANs are appropriate networks for dense-mode operation. Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the Internet. In contrast to DVMRP and MOSPF, PIM dense mode allows a routing device to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routing devices use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), as do all dense-mode protocols. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special notation for that operational mode. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.
- PIM *sparse mode*—In this mode of PIM, the assumption is that very few of the possible receivers want packets from each source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) an interest in the traffic. This multicast protocol allows a routing device to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode has an *explicit* join message, so routing devices determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to the rendezvous point (RP). PIM sparse mode uses an RP routing device as the initial source of multicast group traffic and therefore builds distribution trees in the form (*,G), as do all sparse-mode protocols. PIM sparse mode migrates to an (S,G) source-based tree if that path is shorter than through the RP for a particular multicast group's traffic. WANs are appropriate networks for sparse-mode operation, and indeed a common multicast guideline is not to run dense mode on a WAN under any circumstances.
- Core Based Trees (CBT)—Shares all of the characteristics of PIM sparse mode (sparse mode, explicit join, and shared (*,G) trees), but is said to be more efficient at finding sources than PIM sparse mode. CBT is rarely encountered outside academic discussions. There are no large-scale deployments of CBT, commercial or otherwise.
- PIM source-specific multicast (SSM)—Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of an RP. Used with IGMPv3 to create a shortest-path tree between receiver and source.
- IGMPv1—The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the routing device, but uses a timeout to determine when hosts leave a group. Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routing devices.
- IGMPv2—Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.
- IGMPv3—Defined in RFC 3376, *Internet Group Management Protocol, Version 3.* Among other features, IGMPv3 optimizes support for a single source of content for a multicast

Copyright © 2016, Juniper Networks, Inc.

group, or source-specific multicast (SSM). Used with PIM SSM to create a shortest-path tree between receiver and source.

- Bootstrap Router (BSR) and Auto-Rendezvous Point (RP)—Allow sparse-mode routing protocols to find RPs within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.
- Multicast Source Discovery Protocol (MSDP)—Allows groups located in one multicast routing domain to find RPs in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain. Typically runs on the same routing device as PIM sparse mode RP. Not appropriate if all receivers and sources are located in the same routing domain.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.
- Pragmatic General Multicast (PGM)—Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.

The differences among the multicast routing protocols are summarized in Table 380.

Table 380: Multicast Routing Protocols Compared

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(*,G) Shared Tree
DVMRP	Yes	No	Yes	No	Yes	No
MOSPF	Yes	No	No	Yes	Yes	No
PIM dense mode	Yes	No	Yes	No	Yes	No
PIM sparse mode	No	Yes	No	Yes	Yes, maybe	Yes, initially
Bidirectional PIM	No	No	No	Yes	No	Yes
CBT	No	Yes	No	Yes	No	Yes
SSM	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv1	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv2	No	Yes	No	Yes	Yes, maybe	Yes, initially
IGMPv3	No	Yes	No	Yes	Yes, maybe	Yes, initially

Copyright © 2016, Juniper Networks, Inc.

Table 380: Multicast Routing Protocols Compared *(continued)*

Multicast Routing Protocol	Dense Mode	Sparse Mode	Implicit Join	Explicit Join	(S,G) SBT	(*,G) Shared Tree
BSR and Auto-RP	No	Yes	No	Yes	Yes, maybe	Yes, initially
MSDP	No	Yes	No	Yes	Yes, maybe	Yes, initially

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded routing device can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by the Transmission Control Protocol (TCP) (but TCP always resends missing segments), or the simple drop-and-continue strategy of the User Datagram Protocol (UDP) datagram service (but reordering can become an issue). Modern multicast uses UDP almost exclusively.

T Series Router Multicast Performance

The Juniper Networks T Series Core Routers handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when 1 input port and 63 output ports need a copy of the packet, the T Series routing platform copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T Series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.

Understanding Layer 3 Multicast Functionality on the SRX5K-MPC

Multicast is a “one source, many destinations” method of traffic distribution, meaning that only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

In the data plane of the SRX Series chassis, the SRX5000 line Module Port Concentrator (SRX5K-MPC) forwards Layer 3 IP multicast data packets, which include multicast protocol packets (for example, MLD, IGMP and PIM packets), and the data packets.

In incoming direction, the MPC receives multicast packets from an interface and forwards them to the central point or to a Services Processing Unit (SPU). The SPU performs multicast route lookup, flow-based security check, and packet replication.

In outgoing direction, the MPC receives copies of a multicast packet or Layer 3 multicast control protocol packets from SPU, and transmits them to either multicast capable routers or to hosts in a multicast group.

In the SRX Series chassis, the SPU perform multicast route lookup, if available, to forward an incoming multicast packet and replicates it for each multicast outgoing interface. After receiving replicated multicast packets and their corresponding outgoing interface information from the SPU, the MPC transmits these packets to next hops.

Copyright © 2016, Juniper Networks, Inc.

NOTE: On all high-end SRX Series devices, during RG1 failover with multicast traffic and high number of multicast sessions, the failover delay is from 90 through 120 seconds for traffic to resume on the secondary node. The delay of 90 through 120 seconds is only for the first failover. For subsequent failovers, the traffic resumes within 8 through 18 seconds.

Related Documentation

- Enabling PIM Sparse Mode on page 4118

Supported IP Multicast Protocol Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for IP multicast protocols, including the Distance Vector Multicast Routing Protocol (DVMRP), Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Multicast Source Discovery Protocol (MSDP), Pragmatic General Multicast (PGM), Protocol Independent Multicast (PIM), Session Announcement Protocol (SAP), and Session Description Protocol (SDP):

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3376, *Internet Group Management Protocol, Version 3*
- RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*
- RFC 4601, *Protocol Independent Multicast – Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*
- *Using IGMPv3 and MLDv2 for Source-Specific Multicast*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-pim-sm-bsr-05.txt, *Bootstrap Router (BSR) Mechanism for PIM*

  The scoping mechanism is not supported.
- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs* (expires December 2004)

Copyright © 2016, Juniper Networks, Inc.

The following RFCs and Internet drafts do not define standards, but provide information about multicast protocols and related technologies. The IETF classifies them variously as “Best Current Practice,” “Experimental,” or “Informational.”

- RFC 1075, *Distance Vector Multicast Routing Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3446, *Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- RFC 3973, *Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification (Revised)*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol*
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8*
- Internet draft draft-ietf-mmusic-sap-00.txt, *SAP: Session Announcement Protocol*
- Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*

  Only section 7, “Data MDT: Optimizing flooding,” is supported.

Related Documentation

- *Accessing Standards Documents on the Internet*

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 60

Managing Group Membership

- Configuring IGMP on page 4037
- Examples: Configuring MLD on page 4063

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 176

Configuring IGMP

- Understanding Group Membership Protocols on page 4037
- Understanding IGMP on page 4038
- Configuring IGMP on page 4040
- Enabling IGMP on page 4041
- Modifying the IGMP Host-Query Message Interval on page 4042
- Modifying the IGMP Query Response Interval on page 4043
- Specifying Immediate-Leave Host Removal for IGMP on page 4044
- Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4045
- Accepting IGMP Messages from Remote Subnetworks on page 4046
- Modifying the IGMP Last-Member Query Interval on page 4047
- Modifying the IGMP Robustness Variable on page 4047
- Limiting the Maximum IGMP Message Rate on page 4049
- Changing the IGMP Version on page 4049
- Enabling IGMP Static Group Membership on page 4050
- Recording IGMP Join and Leave Events on page 4056
- Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4057
- Tracing IGMP Protocol Traffic on page 4059
- Disabling IGMP on page 4060
- IGMP and Nonstop Active Routing on page 4060

Understanding Group Membership Protocols

There is a big difference between the multicast protocols used between host and routing device and between the multicast routing devices themselves. Hosts on a given subnetwork need to inform their routing device only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routing devices only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routing devices of their participation in a multicast group. Between adjacent routing devices, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology

Copyright © 2016, Juniper Networks, Inc.

and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a routing device to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the routing device sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the routing device that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routing devices:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the routing device, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the routing device, especially on older or smaller routing devices.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routing devices can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast* (*SSM*).
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a routing device to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any routing device attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the routing device.

Related Documentation

- Examples: Configuring MLD on page 4063

Understanding IGMP

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routing devices in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routing devices. Multicast routing devices use IGMP to learn, for each of their attached physical networks, which groups have members.

Copyright © 2016, Juniper Networks, Inc.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

A routing device receives explicit join and prune messages from those neighboring routing devices that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The routing device then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routing devices are automatically or statically designated as the RP, and all routing devices must explicitly join through the RP.
4. Each routing device along the path toward the RP builds a wildcard (any-source) state for the group and sends join and prune messages toward the RP.

   The term *route entry* is used to refer to the state maintained in a routing device to represent the distribution tree.

   A route entry can include such fields as:

   - source address
   - group address
   - incoming interface from which packets are accepted
   - list of outgoing interfaces to which packets are sent
   - timers
   - flag bits

   The wildcard route entry's incoming interface points toward the RP.

   The outgoing interfaces point to the neighboring downstream routing devices that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.
5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is an integral part of IP and must be enabled on all routing devices and hosts that need to receive IP multicast traffic.

For each attached network, a multicast routing device can be either a querier or a nonquerier. The querier routing device periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

Copyright © 2016, Juniper Networks, Inc.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a routing device can specify particular routing devices from which it accepts or rejects traffic. With IGMPv3, a multicast routing device can learn which sources are of interest to neighboring routing devices.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routing devices, IGMPv3 routing devices must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

Related Documentation

- Supported IP Multicast Protocol Standards on page 4032
- Configuring IGMP on page 4037

Configuring IGMP

Before you begin:

1. Determine whether the routing device is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the routing device is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
    accounting;
    interface interface-name {
        disable;
        (accounting | no-accounting);
        group-policy [ policy-names ];
        immediate-leave;
        oif-map map-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
        promiscuous-mode;
        ssm-map ssm-map-name;
        static {
          group multicast-group-address {
            exclude;
            group-count number;
            group-increment increment;
            source ip-address {
              source-count number;
              source-increment increment;
            }
          }
        }
        version version;
      }
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
      traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
      }
    }
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).

NOTE: You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one "pseudo PIM interface" is created to represent all IGMP downstream (IGMP-only) interfaces on the router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical

Copyright © 2016, Juniper Networks, Inc.

networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/0/0.0
   ```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

   ```
   [edit protocols igmp]
   user@host# show
   interface fe-0/0/0.0;
   interface ge-1/0/0.0 {
     disable;
   }
   ```

3. Enable IGMP on the interface by deleting the **disable** statement.

   ```
   [edit protocols igmp]
   delete interface ge-1/0/0.0 disable
   ```

4. Verify the configuration.

   ```
   [edit protocols igmp]
   user@host# show
   interface fe-0/0/0.0;
   interface ge-1/0/0.0;
   ```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

Related Documentation

- Understanding IGMP on page 4038
- Disabling IGMP on page 4060
- **show igmp interface on page 4678**

Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

Copyright © 2016, Juniper Networks, Inc.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

   ```
   [edit protocols igmp]
   user@host# set query-interval 200
   ```

   The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.

3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 4038
- Modifying the IGMP Query Response Interval on page 4043
- Modifying the IGMP Robustness Variable on page 4047
- **show igmp interface on page 4678**
- **show igmp statistics on page 4692**

Modifying the IGMP Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the

Copyright © 2016, Juniper Networks, Inc.

group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

   ```
   [edit protocols igmp]
   user@host# set query-response-interval 0.4
   ```

2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 4038
- Modifying the IGMP Host-Query Message Interval on page 4042
- Modifying the IGMP Robustness Variable on page 4047
- **show igmp interface on page 4678**
- **show igmp statistics on page 4692**

Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

   ```
   [edit protocols IGMP]
   user@host# set interface ge-0/0/0.1 immediate-leave
   ```

2. Verify the configuration by checking the Immediate Leave field in the output of the **show igmp interface** command.

Related Documentation

- Understanding IGMP on page 4038
- show igmp interface on page 4678

Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a routing device running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

   ```
   [edit policy-statement reject_policy_v2]
   user@host# set from route-filter 224.1.1.1/32 exact
   user@host# set from route-filter 239.0.0.0/8 orlonger
   user@host# set then reject
   ```

2. Configure an IGMPv3 policy.

Copyright © 2016, Juniper Networks, Inc.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

Related Documentation

- Understanding IGMP on page 4038
- show igmp statistics on page 4692

Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.

NOTE: When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.

NOTE: When enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Understanding IGMP on page 4038
- show igmp interface on page 4678
- show igmp statistics on page 4692

Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

   ```
   [edit protocols igmp]
   user@host# set query-last-member-interval 0.1
   ```

2. Verify the configuration by checking the **IGMP Last Member Query Interval** field in the output of the **show igmp interfaces** command.

NOTE: You can configure the last-member query count by configuring the robustness variable. The two are always equal.

Related Documentation
- Modifying the IGMP Robustness Variable on page 4047

Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and

Copyright © 2016, Juniper Networks, Inc.

IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

   When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

   ```
   [edit protocols igmp]
   user@host# set robust-count 5
   ```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

Related Documentation

- Modifying the IGMP Host-Query Message Interval on page 4042
- Modifying the IGMP Query Response Interval on page 4043
- Modifying the IGMP Last-Member Query Interval on page 4047

Copyright © 2016, Juniper Networks, Inc.

- show pim interfaces on page 4763
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the routing device.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a routing device with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the routing device can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

Related Documentation

- maximum-transmit-rate (Protocols IGMP) on page 4501

Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

   ```
   [edit protocols igmp]
   user@host# set interface ge-0/0/0 version 3
   ```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding IGMP on page 4038
- show pim interfaces on page 4763
- show igmp statistics on page 4692

Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp

   interface fe-0/1/2.0 {
     static {
       group 225.1.1.1 ;
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

   ```
   user@host> show igmp group
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

NOTE: When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

Copyright © 2016, Juniper Networks, Inc.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

   ```
   [edit protocols igmp]
   user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
   ```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@host> show configuration protocol igmp
   ```

   ```
   interface fe-0/1/2.0 {
     static {
       group 225.1.1.1 {
         group-count 3;
       }
     }
   }
   ```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

   ```
   user@host> show igmp group
   Interface: fe-0/1/2
        Group: 225.1.1.1
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.2
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
        Group: 225.1.1.3
            Source: 10.0.0.2
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented

Copyright © 2016, Juniper Networks, Inc.

for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment
  0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      group-increment 0.0.0.2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.3
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.5
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

 Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
  3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2 {
        source-count 3;
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.1
         Source: 10.0.0.3
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.1
         Source: 10.0.0.4
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
  3 source-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2 {
        source-count 3;
        source-increment 0.0.0.2;
      }
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
     Group: 225.1.1.1
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.1
         Source: 10.0.0.4
         Last reported by: Local
         Timeout: 0 Type: Static
     Group: 225.1.1.1
         Source: 10.0.0.6
         Last reported by: Local
         Timeout: 0 Type: Static
```

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      exclude;
      source 10.0.0.2;
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
Interface: fe-0/1/2
     Group: 225.1.1.1
         Group mode: Exclude
         Source: 10.0.0.2
         Last reported by: Local
         Timeout: 0 Type: Static
```

Related Documentation

- Enabling MLD Static Group Membership on page 4075
- group (Protocols IGMP) on page 4454
- group-count on page 4458
- group-increment (Protocols IGMP) on page 4459
- source-count (Protocols IGMP) on page 4578
- source-increment (Protocols IGMP) on page 4579
- static (Protocols IGMP) on page 4587

Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

Table 381 describes the recordable IGMP events.

Table 381: IGMP Event Messages

ERRMSG Tag	Definition
RPD_IGMP_JOIN	Records IGMP join events.
RPD_IGMP_LEAVE	Records IGMP leave events.
RPD_IGMP_ACCOUNTING_ON	Records when IGMP accounting is enabled on an IGMP interface.
RPD_IGMP_ACCOUNTING_OFF	Records when IGMP accounting is disabled on an IGMP interface.
RPD_IGMP_MEMBERSHIP_TIMEOUT	Records IGMP membership timeout events.

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
  .*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

   This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
  "anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the ***monitor start*** and ***monitor stop*** commands.

```
user@host> monitor start igmp-events
```

```
*** igmp-events ***
Apr 16 13:08:23  host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

Related Documentation

- Understanding IGMP on page 4038

Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (*,G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.

Copyright © 2016, Juniper Networks, Inc.

- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

   [edit]
   user@host# **edit protocols igmp interface** ***interface-name***

2. Specify the group limit for the interface.

   [edit protocols igmp interface *interface-name*]
   user@host# **set group-limit** ***limit***

3. (Optional) Configure the amount of time between log messages.

   [edit protocols igmp interface *interface-name*]
   user@host# **set log-interval** ***seconds***

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

Copyright © 2016, Juniper Networks, Inc.

Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.
general	Trace general flow.
group	Trace group operations.
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
mtrace	Trace mtrace packets. Use the **mtrace** command to troubleshoot the software.
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages, including general and group-specific queries.
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

   [edit routing-options traceoptions]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

Related Documentation

- Understanding IGMP on page 4038

Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols igmp interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]**

Related Documentation

- Enabling IGMP on page 4041

IGMP and Nonstop Active Routing

Nonstop active routing (NSR) configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. These NSR configurations include passive support with IGMP in connection with PIM. The master

Copyright © 2016, Juniper Networks, Inc.

Routing Engine uses IGMP to determine its PIM multicast state, and this IGMP-derived information is replicated on the backup Routing Engine. IGMP on the new master Routing Engine (after failover) relearns the state information quickly through IGMP operation. In the interim, the new master Routing Engine retains the IGMP-derived PIM state as received by the replication process from the old master Routing Engine. This state information times out unless refreshed by IGMP on the new master Routing Engine. No additional IGMP configuration is required.

Related Documentation

- Understanding Nonstop Active Routing for PIM on page 4227
- Examples: Configuring MLD on page 4063

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 177

Examples: Configuring MLD

- Understanding MLD on page 4063
- Configuring MLD on page 4066
- Enabling MLD on page 4067
- Modifying the MLD Version on page 4068
- Modifying the MLD Host-Query Message Interval on page 4068
- Modifying the MLD Query Response Interval on page 4069
- Modifying the MLD Last-Member Query Interval on page 4070
- Specifying Immediate-Leave Host Removal for MLD on page 4071
- Filtering Unwanted MLD Reports at the MLD Interface Level on page 4072
- Example: Modifying the MLD Robustness Variable on page 4073
- Limiting the Maximum MLD Message Rate on page 4075
- Enabling MLD Static Group Membership on page 4075
- Example: Recording MLD Join and Leave Events on page 4082
- Configuring the Number of MLD Multicast Group Joins on Logical Interfaces on page 4084
- Disabling MLD on page 4085

Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each router maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the router does not need to know the address of each listener—just the address of each host. The router provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.

MLD is an integral part of IPv6 and must be enabled on all IPv6 routers and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

Copyright © 2016, Juniper Networks, Inc.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast router can be either a querier or a nonquerier. A querier router, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier router that it has interested listeners, the querier router forwards the membership information to the rendezvous point (RP) router by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP router. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routers do not transmit MLD queries on a subnet but can do so if the querier router fails.

All MLD-configured routers start as querier routers on each attached subnet (see Figure 154). The querier router on the right is the receiver's DR.

Figure 154: Routers Start Up on a Subnet

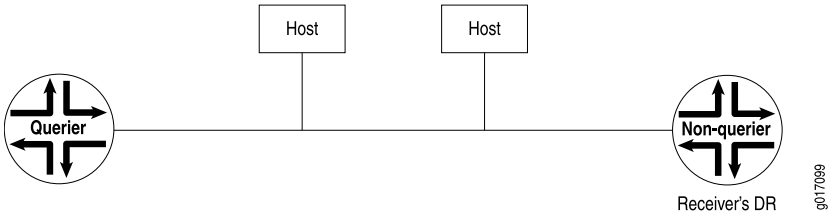

To elect the querier router, the routers exchange query messages containing their IPv6 source addresses. If a router hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In Figure 155, the router on the left has a source address numerically lower than the one on the right and therefore becomes the querier router.

NOTE: In the practical application of MLD, several routers on a subnet are nonqueriers. If the elected querier router fails, query messages are exchanged among the remaining routers. The router with the lowest IPv6 source address becomes the new querier router. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 155: Querier Router Is Determined

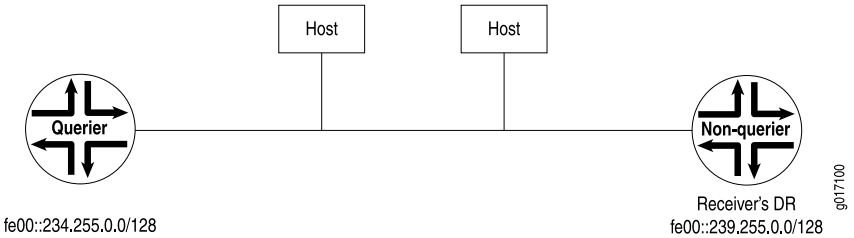

Copyright © 2016, Juniper Networks, Inc.

The querier router sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see Figure 156). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 156: General Query Message Is Issued

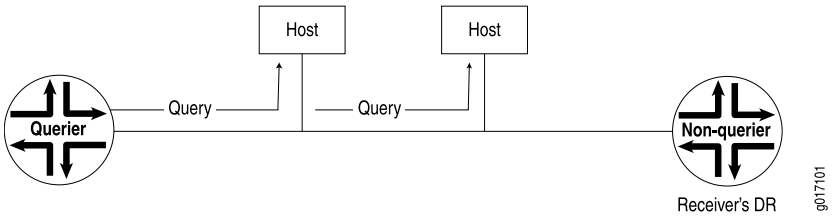

If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the router (see Figure 157). If the reported address is not yet in the router's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

Figure 157: Reports Are Received by the Querier Router

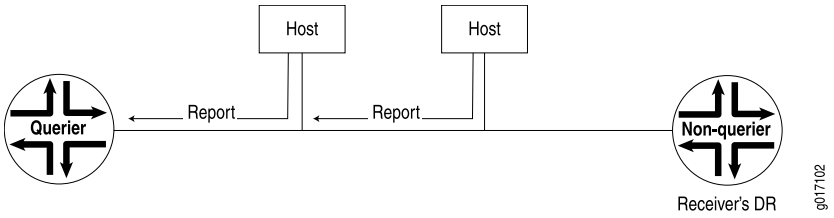

If the host has no interested multicast listeners, it sends a done message to the querier router. On receipt, the querier router issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the router does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see Figure 158).

Figure 158: Host Has No Interested Receivers and Sends a Done Message to Router

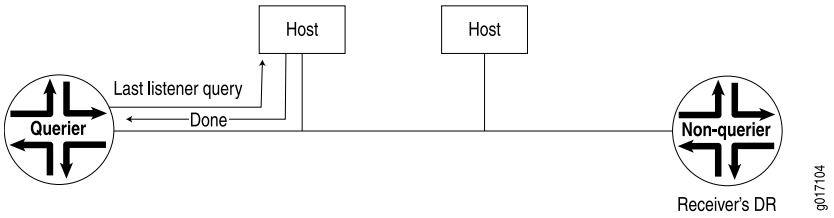

If a done message is not received by the querier router, the querier router continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier router assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see Figure 159).

Copyright © 2016, Juniper Networks, Inc.

Figure 159: Host Address Timer Expires and Address Is Removed from Multicast Address List

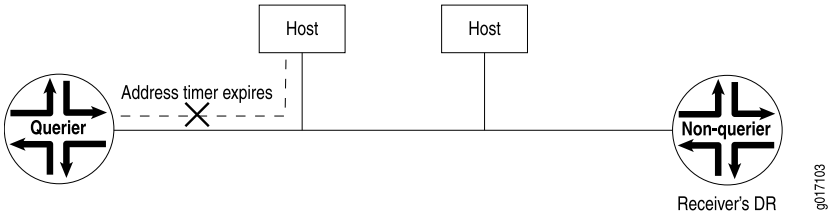

Configuring MLD

To configure the Multicast Listener Discovery (MLD) Protocol, include the **mld** statement:

```
mld {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
    oif-map [ map-names ];
    passive;
    ssm-map ssm-map-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
  maximum-transmit-rate packets-per-second;
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

Copyright © 2016, Juniper Networks, Inc.

By default, MLD is enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or the Distance Vector Multicast Routing Protocol (DVMRP).

Enabling MLD

The Multicast Listener Discovery (MLD) Protocol manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use MLD to learn which groups have members on each of their attached physical networks. MLD must be enabled for the router to receive IPv6 multicast packets. MLD is only needed for IPv6 networks, because multicast is handled differently in IPv4 networks. MLD is enabled on all IPv6 interfaces on which you configure PIM and on all IPv6 broadcast interfaces when you configure DVMRP.

MLD specifies different behaviors for multicast listeners and for routers. When a router is also a listener, the router responds to its own messages. If a router has more than one interface to the same link, it needs to perform the router behavior over only one of those interfaces. Listeners, on the other hand, must perform the listener behavior on all interfaces connected to potential receivers of multicast traffic.

If MLD is not running on an interface—either because PIM and DVMRP are not configured on the interface or because MLD is explicitly disabled on the interface—you can explicitly enable MLD.

To explicitly enable MLD:

1. If PIM and DVMRP are not running on the interface, explicitly enable MLD by including the interface name.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/0/0.0
   ```

2. Check to see if MLD is disabled on any interfaces. In the following example, MLD is disabled on a Gigabit Ethernet interface.

   ```
   [edit protocols mld]
   user@host# show
   ```

   ```
   interface fe-0/0/0.0;
   interface ge-0/0/0.0 {
       disable;
   }
   ```

3. Enable MLD on the interface by deleting the **disable** statement.

   ```
   [edit protocols mld]
   delete interface ge-0/0/0.0 disable
   ```

4. Verify the configuration.

   ```
   [edit protocols mld]
   user@host# show
   ```

   ```
   interface fe-0/0/0.0;
   interface ge-0/0/0.0;
   ```

5. Verify the operation of MLD by checking the output of the **show mld interface** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding MLD on page 4063
- Disabling MLD on page 4085
- **show mld interface on page 4699** in the CLI Explorer

Modifying the MLD Version

By default, the router supports MLD version 1 (MLDv1). To enable the router to use MLD version 2 (MLDv2) for source-specific multicast (SSM) only, include the **version 2** statement.

If you configure the MLD version setting at the individual interface hierarchy level, it overrides configuring the IGMP version using the **interface all** statement.

If a source address is specified in a multicast group that is statically configured, the version must be set to MLDv2.

To change an MLD interface to version 2:

1. Configure the MLD interface.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/0/0.0 version 2
   ```

2. Verify the configuration by checking the **version** field in the output of the **show mld interface** command. The **show mld statistics** command has version-specific output fields, such as the counters in the **MLD Message type** field.

Related Documentation

- Understanding MLD on page 4063
- Source-Specific Multicast Groups Overview on page 4186
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4187
- Example: Configuring an SSM-Only Domain on page 4190
- Example: Configuring PIM SSM on a Network on page 4191
- Example: Configuring SSM Mapping on page 4192

Modifying the MLD Host-Query Message Interval

The objective of MLD is to keep routers up to date with IPv6 group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The MLD querier router periodically sends general host-query messages on each attached network to solicit membership information. These messages solicit group membership information and are sent to the **link-scope all-nodes** address **FF02::1**. A general host-query message has a maximum response time that you can set by configuring the query response interval.

Copyright © 2016, Juniper Networks, Inc.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of MLD messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

   ```
   [edit protocols mld]
   user@host# set query-interval 200
   ```

   The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the **MLD Query Interval** field in the output of the **show mld interface** command.

3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

Related Documentation

- Understanding MLD on page 4063
- Modifying the MLD Query Response Interval on page 4069
- Example: Modifying the MLD Robustness Variable on page 4073
- **show mld interface on page 4699** in the CLI Explorer
- **show mld statistics on page 4702** in the CLI Explorer

Modifying the MLD Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

Copyright © 2016, Juniper Networks, Inc.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

   ```
   [edit protocols mld]
   user@host# set query-response-interval 0.5
   ```

2. Verify the configuration by checking the **MLD Query Response Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

Related Documentation

- Understanding MLD on page 4063
- Modifying the MLD Host-Query Message Interval on page 4068
- Example: Modifying the MLD Robustness Variable on page 4073
- **show mld interface on page 4699** in the CLI Explorer
- **show mld statistics on page 4702** in the CLI Explorer

Modifying the MLD Last-Member Query Interval

The last-member query interval (also called the last-listener query interval) is the maximum amount of time between group-specific query messages, including those sent in response to done messages sent on the **link-scope-all-routers** address FF02::2. You can lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group (done) message from a host, the routing device sends multiple group-specific queries to the group. The querier sends a specific number of these queries, and it sends them at a specific interval. The number of queries sent is called the last-listener query count. The interval at which the queries are sent is called the last-listener query interval. Both settings are configurable, thus allowing you to adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-listener query count x (times) the last-listener query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-listener query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

Copyright © 2016, Juniper Networks, Inc.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

   ```
   [edit protocols mld]
   user@host# set query-last-member-interval 0.1
   ```

2. Verify the configuration by checking the **MLD Last Member Query Interval** field in the output of the **show igmp interfaces** command.

NOTE: You can configure the last-member query count by configuring the robustness variable. The two are always equal.

Related Documentation

- Understanding MLD on page 4063
- Modifying the MLD Query Response Interval on page 4069
- Example: Modifying the MLD Robustness Variable on page 4073
- **show mld interface on page 4699** in the CLI Explorer

Specifying Immediate-Leave Host Removal for MLD

The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave:

1. Configure immediate leave on the MLD interface.

   ```
   [edit protocols mld]
   user@host# set interface ge-0/0/0.1 immediate-leave
   ```

2. Verify the configuration by checking the **Immediate Leave** field in the output of the **show mld interface** command.

Related Documentation

- Understanding MLD on page 4063
- show mld interface on page 4699 in the CLI Explorer

Filtering Unwanted MLD Reports at the MLD Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted MLD reports at the interface level.

When the **group-policy** statement is enabled on a router, after the router receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only MLD group addresses (for MLDv1) by using the policy's **route-filter** statement to match the group address. You define the policy to match MLD (source, group) addresses (for MLDv2) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted MLD reports:

1. Configure an MLDv1 policy.

   ```
   [edit policy-statement reject_policy_v1]
   user@host# set from route-filter fec0:1:1:4::/64 exact
   user@host# set then reject
   ```

2. Configure an MLDv2 policy.

   ```
   [edit policy-statement reject_policy_v2]
   user@host# set from route-filter fec0:1:1:4::/64 exact
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger
user@host# set then reject
```

3. Apply the policies to the MLD interfaces where you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running MLDv1 and **ge-0/1/1.0** is running MLDv2.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v1
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v2
```

4. Verify the operation of the filter by checking the **Rejected Report** field in the output of the **show mld statistics** command.

Related Documentation

- Understanding MLD on page 4063
- **show mld statistics on page 4702** in the CLI Explorer

Example: Modifying the MLD Robustness Variable

This example shows how to configure and verify the MLD robustness variable in a multicast domain.

- Requirements on page 4073
- Overview on page 4073
- Configuration on page 4074
- Verification on page 4074

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable PIM. See "PIM Overview" on page 4089.

Overview

The MLD robustness variable can be fine-tuned to allow for expected packet loss on a subnet. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

The value of the robustness variable is used in calculating the following MLD message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is

Copyright © 2016, Juniper Networks, Inc.

calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).

- Other querier present interval—Amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The default number is the value of the robustness variable.

By default, the robustness variable is set to 2. The number can be from 2 through 10. You might want to increase this value if you expect a subnet to lose packets.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols mld robust-count 5
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To change the value of the robustness variable:

1. Configure the robust count.

   ```
   [edit protocols mld]
   user@host# set robust-count 5
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit protocols mld]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, check the **MLD Robustness Count** field in the output of the **show mld interfaces** command.

Related Documentation

- Understanding MLD on page 4063
- Modifying the MLD Query Response Interval on page 4069
- Modifying the MLD Last-Member Query Interval on page 4070
- **show mld interface on page 4699** in the CLI Explorer

Copyright © 2016, Juniper Networks, Inc.

Limiting the Maximum MLD Message Rate

You can change the limit for the maximum number of MLD packets transmitted in 1 second by the router.

Increasing the maximum number of MLD packets transmitted per second might be useful on a router with a large number of interfaces participating in MLD.

To change the limit for the maximum number of MLD packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

Enabling MLD Static Group Membership

You can create MLD static group membership to test multicast forwarding without a receiver host. When you enable MLD static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

Class-of-service (CoS) adjustment is not supported with MLD static group membership.

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created.

In this example, you create static group ff02::1:ff05:1a8d.

1. Configure the static groups to be created by including the **static** statement and **group** statement and specifying which IPv6 multicast address of the group to be created.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d
   ```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

   ```
   user@host> show configuration protocol mld

   interface fe-0/1/2.0 {
     static {
       group ff02::1:ff05:1a8d;
     }
   }
   ```

3. After you have committed the configuration and after the source is sending traffic, use the **show mld group** command to verify that static group ff02::1:ff05:1a8d has been created.

   ```
   user@host> show mld group
   Interface: fe-0/1/2
   Group: ff02::1:ff05:1a8d
   Group mode: Include
   Source: fe80::2e0:81ff:fe05:1a8d
   Last reported by: Local
   Timeout: 0 Type: Static
   ```

Copyright © 2016, Juniper Networks, Inc.

NOTE: You must specify a unique address for each group.

Automatically create static groups

When you create MLD static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. Configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d group-count 3
   ```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

   ```
   user@host> show configuration protocol mld

   interface fe-0/1/2.0 {
     static {
       group ff02::1:ff05:1a8d {
         group-count 3;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff02::1:ff05:1a8d, ff02::1:ff05:1a8e, and ff02::1:ff05:1a8f have been created.

   ```
   user@host> show mld group

   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8e
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8f
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

Automatically increment group addresses

Copyright © 2016, Juniper Networks, Inc.

When you configure static groups on an interface on which you want to receive multicast traffic and you specify the number of static groups to be automatically created, you can also configure the group address to be automatically incremented by some number of addresses.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. Configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in a format similar to an IPv6 address.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d group-count 3
     group-increment ::2
   ```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

   ```
   user@host> show configuration protocol mld

   interface fe-0/1/2.0 {
     static {
       group ff02::1:ff05:1a8d {
         group-increment ::2;
         group-count 3;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff02::1:ff05:1a8d, ff02::1:ff05:1a8f, and ff02::1:ff05:1a91 have been created.

   ```
   user@host> show mld group

   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8f
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a91
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

Specify multicast source address (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify the multicast source address to be accepted.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you create group ff02::1:ff05:1a8d and accept IPv6 address fe80::2e0:81ff:fe05:1a8d as the only source.

1. Configure the source address by including the **source** statement and specifying the IPv6 address of the source host.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d source
     fe80::2e0:81ff:fe05:1a8d
   ```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

   ```
   user@host> show configuration protocol mld

   interface fe-0/1/2.0 {
     static {
       group ff02::1:ff05:1a8d {
         source fe80::2e0:81ff:fe05:1a8d;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff02::1:ff05:1a8d has been created and that source fe80::2e0:81ff:fe05:1a8d has been accepted.

   ```
   user@host> show mld group

   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

Automatically specify multicast sources

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify a number of multicast sources to be automatically accepted.

In this example, you create static group ff02::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f as the source addresses.

1. Configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

   ```
   [edit protocols mld]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d source
  fe80::2e0:81ff:fe05:1a8d source-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff02::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff02::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8e
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8f
         Last reported by: Local
         Timeout: 0 Type: Static
```

Automatically increment source addresses

When you configure static groups on an interface on which you want to receive multicast traffic, and specify a number of multicast sources to be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted.

In this example, you create static group ff02::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 as the sources.

1. Configure the number of multicast source addresses to be accepted by including the **source-increment** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d source
  fe80::2e0:81ff:fe05:1a8d source-count 3 source-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff02::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
        source-increment ::2;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff02::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8f, and fe80::2e0:81ff:fe05:1a91 have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8d
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a8f
         Last reported by: Local
         Timeout: 0 Type: Static
Interface: fe-0/1/2
     Group: ff02::1:ff05:1a8d
         Source: fe80::2e0:81ff:fe05:1a91
         Last reported by: Local
         Timeout: 0 Type: Static

Interface: fe-0/1/2
Group: ff02::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
Group: ff02::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8f
Last reported by: Local
Timeout: 0 Type: Static
Group: ff02::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a91
Last reported by: Local
Timeout: 0 Type: Static
```


Copyright © 2016, Juniper Networks, Inc.

Exclude multicast source addresses (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the configured source address. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the configured source address.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you exclude address fe80::2e0:81ff:fe05:1a8d as a source for group ff02::1:ff05:1a8d.

1. Configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv6 source address to be excluded.

   ```
   [edit protocols mld]
   user@host# set interface fe-0/1/2 static group ff02::1:ff05:1a8d exclude source
     fe80::2e0:81ff:fe05:1a8d
   ```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

   ```
   user@host> show configuration protocol mld

   interface fe-0/1/2.0 {
     static {
       group ff02::1:ff05:1a8d {
         exclude;
         source fe80::2e0:81ff:fe05:1a8d;
       }
     }
   }
   ```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group detail** command to verify that static group ff02::1:ff05:1a8d has been created and that the static group is operating in exclude mode.

   ```
   user@host> show mld group detail
   Interface: fe-0/1/2
        Group: ff02::1:ff05:1a8d
            Group mode: Exclude
            Source: fe80::2e0:81ff:fe05:1a8d
            Last reported by: Local
            Timeout: 0 Type: Static
   ```

Similar configuration is available for IPv4 multicast traffic using the IGMP protocol.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

Example: Recording MLD Join and Leave Events

This example shows how to determine whether MLD tuning is needed in a network by configuring the routing device to record MLD join and leave events.

- Requirements on page 4082
- Overview on page 4082
- Configuration on page 4082
- Verification on page 4084

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable PIM. See "PIM Overview" on page 4089.

Overview

Table 382 describes the recordable MLD join and leave events.

Table 382: MLD Event Messages

ERRMSG Tag	Definition
RPD_MLD_JOIN	Records MLD join events.
RPD_MLD_LEAVE	Records MLD leave events.
RPD_MLD_ACCOUNTING_ON	Records when MLD accounting is enabled on an MLD interface.
RPD_MLD_ACCOUNTING_OFF	Records when MLD accounting is disabled on an MLD interface.
RPD_MLD_MEMBERSHIP_TIMEOUT	Records MLD membership timeout events.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set protocols mld interface fe-0/1/0.2 accounting
set system syslog file mld-events any info
set system syslog file mld-events match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
  .*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
set system syslog file mld-events archive size 100000
set system syslog file mld-events archive files 3
set system syslog file mld-events archive transfer-interval 1440
set system syslog file mld-events archive archive-sites "ftp://user@host1//var/tmp"
  password "anonymous"
set system syslog file mld-events archive archive-sites "ftp://user@host2//var/tmp"
  password "test"
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure recording of MLD join and leave events:

1. Enable accounting globally or on an MLD interface. This example shows the interface configuration.

   ```
   [edit protocols mld]
   user@host# set interface fe–0/1/0.2 accounting
   ```

2. Configure the events to be recorded, and filter the events to a system log file with a descriptive filename, such as **mld-events**.

   ```
   [edit system syslog file mld-events]
   user@host# set any info
   [edit system syslog file mld-events]
   user@host# set match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
     .*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
   ```

3. Periodically archive the log file.

   This example rotates the file every 24 hours (1440 minutes) when it reaches 100 KB and keeps three files.

   ```
   [edit system syslog file mld-events]
   user@host# set archive size 100000
   [edit system syslog file mld-events]
   user@host# set archive files 3
   [edit system syslog file mld-events]
   user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
     "anonymous"
   [edit system syslog file mld-events]
   user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
   [edit system syslog file mld-events]
   user@host# set archive transfer-interval 1440
   [edit system syslog file mld-events]
   user@host# set archive start-time 2011–01–07:12:30
   ```

4. If you are done configuring the device, commit the configuration.

   ```
   [edit system syslog file mld-events]]
   user@host# commit
   ```

Copyright © 2016, Juniper Networks, Inc.

Verification

You can view the system log file by running the **file show** command.

```
user@host> file show mld-events
```

You can monitor the system log file as entries are added to the file by running the ***monitor start*** and ***monitor stop*** commands.

```
user@host> monitor start mld-events
```

```
*** mld-events ***
Apr 16 13:08:23  host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run
 monitor start mld-events '
monitor
```

Related Documentation

- Understanding MLD on page 4063

Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of MLD multicast group joins for logical interfaces. When this statement is enabled on a router running MLD version 2, the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for MLD multicast groups, keep the following in mind:

- Each any-source group (*,G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in MLDv2 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on MLD logical interfaces by using dynamic profiles.

To limit multicast group joins on an MLD logical interface:

1. Access the logical interface at the MLD protocol hierarchy level.

   ```
   [edit]
   user@host# edit protocols mld interface interface-name
   ```


Copyright © 2016, Juniper Networks, Inc.

2. Specify the group limit for the interface.

```
[edit protocols mld interface interface-name]
user@host# set group-limit limit
```

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Disabling MLD

To disable MLD on an interface, include the **disable** statement:

```
interface interface-name {
  disable;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols** mld**]**
- **[edit logical-systems** *logical-system-name* **protocols** mld**]**

Related Documentation

- Enabling MLD on page 4067
- Configuring IGMP on page 4037

Copyright © 2016, Juniper Networks, Inc.

PART 61

Configuring Protocol Independent Multicast

- Understanding PIM on page 4089
- Configuring PIM Basics on page 4093
- Routing Content to Densely Clustered Receivers with PIM Dense Mode on page 4107
- Routing Content to Larger, Sparser Groups with PIM Sparse Mode on page 4113
- Receiving Content Directly from the Source with SSM on page 4183
- Minimizing Routing State Information with Bidirectional PIM on page 4199
- Rapidly Detecting Communication Failures with PIM and the BFD Protocol on page 4219
- Configuring PIM Options on page 4227

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 178

Understanding PIM

- PIM Overview on page 4089

PIM Overview

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same router and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routers connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize

Copyright © 2016, Juniper Networks, Inc.

PIMv1 packets and automatically switch the router interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the router processes the PIM message, a router can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called "protocol independent" because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [*,G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routers must join and leave multicast groups explicitly. Upstream routers do not forward multicast traffic to a downstream router unless the downstream router has sent an explicit request (by means of a join message) to the rendezvous point (RP) router to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

  Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routers build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (*,G) state.

- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routers sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a router receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the router has no interested receivers for the data, and the outgoing interface list becomes empty, the router sends a PIM prune message upstream.

Copyright © 2016, Juniper Networks, Inc.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

NOTE: On SRX Series devices, PIM does not support upstream and downstream interfaces across different virtual routers in flow mode.

Basic PIM Network Components

PIM dense mode requires only a multicast source and series of multicast-enabled routers running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routers called *rendezvous points* (*RPs*) in the network core. These routers are where upstream join messages from interested receivers meet downstream traffic from the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routers find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router* (*BSR*) running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Supported IP Multicast Protocol Standards on page 4032

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 179

Configuring PIM Basics

- Configuring Basic PIM Settings on page 4093
- Configuring a Designated Router for PIM on page 4103

Configuring Basic PIM Settings

- PIM Configuration Statements on page 4093
- Changing the PIM Version on page 4096
- Modifying the PIM Hello Interval on page 4096
- Preserving Multicast Performance by Disabling Response to the ping Utility on page 4097
- PIM on Aggregated Interfaces on page 4098
- Configuring PIM Trace Options on page 4098
- Disabling PIM on page 4100

PIM Configuration Statements

To configure Protocol Independent Multicast (PIM), include the **pim** statement:

```
pim {
  disable;
  default-vpn-source {
    interface-name interface-name;
  }
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  graceful-restart {
    disable;
    no-bidirectional-mode;
    restart-duration seconds;
  }
  idle-standby-path-switchover-delay seconds;
  import [ policy-names ];
  interface interface-name {
    bidirectional {
      df-election {
```

Copyright © 2016, Juniper Networks, Inc.

```
        backoff-period milliseconds;
        offer-period milliseconds;
        robustness-count number;
      }
    }
    import;
    hello-interval seconds;
    mode bidirectional-sparse | bidirectional-sparse-dense | (dense | sparse |
      sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance {
    automatic;
  }
  join-prune-timeout;
  nonstop-routing {
    disable;
  }
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group {
    inet group-name;
    inet6 group-name;
  }
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
      address address {
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
        }
      }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
```


Copyright © 2016, Juniper Networks, Inc.

```
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
      }
      local {
        family (inet | inet6) {
          address address;
          anycast-pim {
            rp-set {
              address address <forward-msdp-sa>;
            }
            local-address address;
          }
          disable;
          group-ranges {
            destination-ip-prefix</prefix-length>;
          }
          hold-time seconds;
          override;
          priority number;
        }
      }
      rp-register-policy [ policy-names ];
      static {
        address address {
          override;
          version version;
          group-ranges {
            destination-ip-prefix</prefix-length>;
          }
          spt-threshold {
            infinity [ policy-names ];
          }
          traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
          }
        }
      }
    }
  }
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

By default, PIM is disabled.

Copyright © 2016, Juniper Networks, Inc.

NOTE: You cannot configure PIM within a nonforwarding instance. If you try to do so, the router displays a commit check error and does not complete the configuration commit process.

Changing the PIM Version

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
   user@host# set hello-interval 255
   ```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

   ```
   user@host> show pim neighbors detail
   Instance: PIM.master
   Interface: fe-3/0/2.0
   Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 255 seconds
   Hello Option DR Priority: 1
   ```


Copyright © 2016, Juniper Networks, Inc.

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

   ```
   [edit system]
   user@host# set no-multicast-echo
   ```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

   ```
   user@host> show system statistics icmp
   ```

   ```
   icmp:
   0 drops due to rate limit
   0 calls to icmp_error
   0 errors not generated because old message was icmp
   Output histogram:
   echo reply: 21
   0 messages with bad code fields
   0 messages less than the minimum length
   0 messages with bad checksum
   0 messages with bad source address
   0 messages with bad length
   100 echo drops with broadcast or multicast destination address
   0 timestamp drops with broadcast or multicast destination address
   Input histogram:
   echo: 21
   21 message responses generated
   ```

Copyright © 2016, Juniper Networks, Inc.

PIM on Aggregated Interfaces

You can configure several Protocol Independent Multicast (PIM) features on an interface regardless of its PIM mode (bidirectional, sparse, dense, or sparse-dense mode).

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is "expanded" into its constituent interfaces in the next-hop database.

Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
assert	Trace assert messages, which are used to resolve which of the parallel routing devices connected to a multiaccess LAN is responsible for forwarding packets to the LAN.
autorp	Trace bootstrap, RP, and auto-RP messages.
bidirectional-df-election	Trace bidirectional PIM designated-forwarder (DF) election events.
bootstrap	Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap routing device and are forwarded, hop by hop, to all routing devices in that domain.
general	Trace general events.
graft	Trace graft and graft acknowledgment messages.
hello	Trace hello packets, which are sent so that neighboring routing devices can discover one another.
join	Trace join messages, which are sent to join a branch onto the multicast distribution tree.
mdt	Trace messages related to multicast data tunnels.
normal	Trace normal events.
nsr-synchronization	Trace nonstop routing synchronization events
packets	Trace all PIM packets.

Copyright © 2016, Juniper Networks, Inc.

Flag	Description
policy	Trace poison-route-reverse packets.
prune	Trace prune messages, which are sent to prune a branch off the multicast distribution tree.
register	Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.
route	Trace routing information.
rp	Trace candidate RP advertisements.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the PIM trace file.

   ```
   [edit protocols pim traceoptions]
   user@host# set file pim-trace
   ```

3. (Optional) Configure the maximum number of trace files.

   ```
   [edit protocols pim traceoptions]
   user@host# set file files 5
   ```

4. (Optional) Configure the maximum size of each trace file.

   ```
   [edit protocols pim traceoptions]
   user@host# set file size 1m
   ```

5. (Optional) Enable unrestricted file access.

   ```
   [edit protocols pim traceoptions]
   user@host# set file world-readable
   ```

6. Configure tracing flags.

Copyright © 2016, Juniper Networks, Inc.

Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match “Rx V1 Require V2”
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

- Disabling the PIM Protocol on page 4101
- Disabling PIM on an Interface on page 4101
- Disabling PIM for a Family on page 4102
- Disabling PIM for a Rendezvous Point on page 4102

Copyright © 2016, Juniper Networks, Inc.

Disabling the PIM Protocol

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
    disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim disable
   ```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Disabling PIM on an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
    interface interface-name {
        disable;
    }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim interface fe-0/1/0 disable
   ```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Copyright © 2016, Juniper Networks, Inc.

Disabling PIM for a Family

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
   family inet {
      disable;
   }
   family inet6 {
      disable;
   }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

   ```
   user@host# set protocols pim family inet disable
   user@host# set protocols pim family inet6 disable
   ```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```

Disabling PIM for a Rendezvous Point

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
   rp {
      local {
         family inet {
            disable;
         }
         family inet6 {
            disable;
         }
      }
   }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.

   ```
   user@host# set protocols pim rp local family inet disable
   user@host# set protocols pim rp local family inet6 disable
   ```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

   ```
   user@host# run show protocols pim
   ```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Dense Mode on page 4107
- Configuring a Designated Router for PIM on page 4103
- Configuring PIM Filtering on page 4159
- Example: Configuring Nonstop Active Routing for PIM on page 4227
- Examples: Configuring PIM RPT and SPT Cutover on page 4166

Configuring a Designated Router for PIM

- Configuring Interface Priority for PIM Designated Router Selection on page 4103
- Configuring PIM Designated Router Election on Point-to-Point Links on page 4104

Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) routing device learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has an equal probability (priority 1) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
   user@host# set priority 5
   ```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

   ```
   user@host> show pim neighbors detail
   ```

   ```
   Instance: PIM.master
   Interface: ge-0/0/0.0
   Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
   Hello Option Holdtime: 65535 seconds
   Hello Option DR Priority: 5
   Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
   Join Suppression supported
   Rx Join: Group Source Timeout
   225.1.1.1 192.168.195.78 0
   225.1.1.1 0

   Interface: lo0.0
   Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
   ```

Copyright © 2016, Juniper Networks, Inc.

```
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

Configuring PIM Designated Router Election on Point-to-Point Links

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim]
   user@host# set dr-election-on-p2p
   ```

2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.

3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routing devices on the point-to-point link. Then recheck the state.

   CAUTION: Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

   ```
   [edit]
   user@host# run restart routing
   ```

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Dense Mode on page 4107
- Configuring PIM Filtering on page 4159
- Example: Configuring Nonstop Active Routing for PIM on page 4227
- Examples: Configuring PIM RPT and SPT Cutover on page 4166

Copyright © 2016, Juniper Networks, Inc.

- Examples: Configuring PIM Sparse Mode on page 4113

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 180

Routing Content to Densely Clustered Receivers with PIM Dense Mode

- Configuring PIM Dense Mode on page 4107
- Configuring PIM Sparse-Dense Mode on page 4110

Configuring PIM Dense Mode

- Understanding PIM Dense Mode on page 4107
- Configuring PIM Dense Mode Properties on page 4109

Understanding PIM Dense Mode

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.

PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.

Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS (M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.

Unlike sparse mode, in which data is forwarded only to routing devices sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A routing device receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see Figure 160).

Copyright © 2016, Juniper Networks, Inc.

Figure 160: Multicast Traffic Flooded from the Source Using PIM Dense Mode

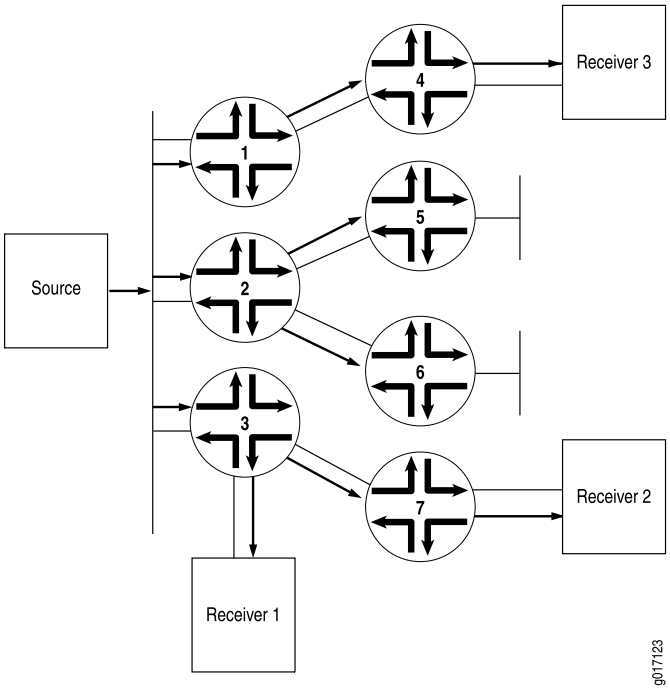

Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the OIL becomes empty, the routing device sends a prune message upstream to stop delivery of multicast traffic (see Figure 161).

Copyright © 2016, Juniper Networks, Inc.

Figure 161: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic

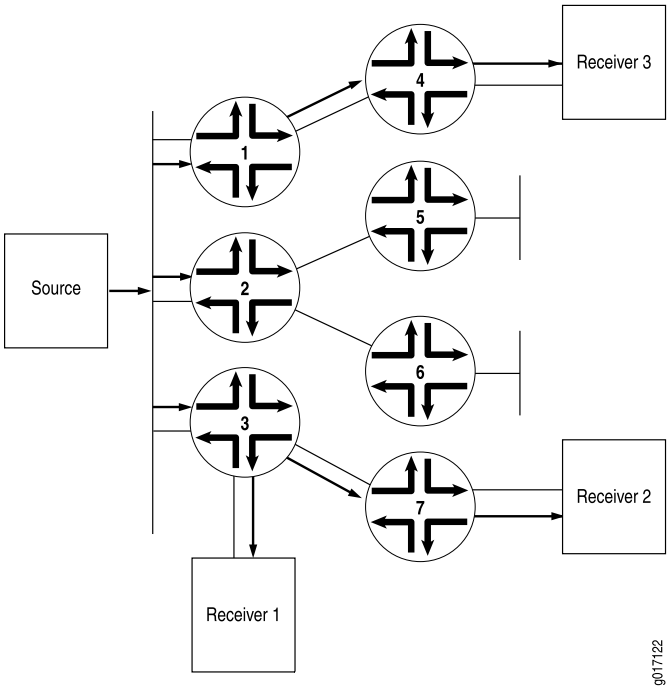

Configuring PIM Dense Mode Properties

In PIM dense mode (PIM-DM), the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is flooded with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). LANs are appropriate networks for dense-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM dense mode globally or for a routing instance. This example shows how to configure the routing instance and how to specify that PIM dense mode use **inet.2** as its RPF routing table instead of **inet.0**.

To configure the routing device properties for PIM dense mode:

1. (Optional) Create an IPv4 routing table group so that interface routes are installed into two routing tables, **inet.0** and **inet.2**.

   ```
   [edit routing-options rib-groups]
   user@host# set pim-rg export-rib inet.0
   user@host# set pim-rg import-rib [ inet.0 inet.2 ]
   ```

2. (Optional) Associate the routing table group with a PIM routing instance.

   ```
   [edit routing-instances PIM.dense protocols pim]
   user@host# set rib-group inet pim-rg
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Configure the PIM interface. If you do not specify any interfaces, PIM is enabled on all routing device interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set interface fe-0/0/1.0 mode dense
```

NOTE: You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

4. Monitor the operation of PIM dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

Related Documentation

- Configuring PIM Sparse-Dense Mode on page 4110
- Configuring Basic PIM Settings on page 4093

Configuring PIM Sparse-Dense Mode

- Understanding PIM Sparse-Dense Mode on page 4110
- Mixing PIM Sparse and Dense Modes on page 4110
- Configuring PIM Sparse-Dense Mode Properties on page 4111

Understanding PIM Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules.

For information about PIM sparse-mode and PIM dense-mode rules, see "Understanding PIM Sparse Mode" on page 4113 and "Understanding PIM Dense Mode" on page 4107.

Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same routing device, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.

NOTE: PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

Copyright © 2016, Juniper Networks, Inc.

A multicast routing device employing sparse-dense mode is a good example of mixing PIM modes on the same network or routing device or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

Configuring PIM Sparse-Dense Mode Properties

Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM sparse-dense mode globally or for a routing instance. This example shows how to configure PIM sparse-dense mode globally on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode.

To configure the routing device properties for PIM sparse-dense mode:

1. Configure the dense-mode groups.

   ```
   [protocols pim]
   user@host# set dense-groups 224.0.1.39
   user@host# set dense-groups 224.0.1.40
   ```

2. Configure all interfaces on the routing device to use sparse-dense mode. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the disable statement for that interface.

   ```
   [edit protocols pim]
   user@host# set interface all mode sparse-dense
   user@host# set interface fxp0.0 disable
   ```

3. Monitor the operation of PIM sparse-dense mode by running the show pim interfaces, show pim join, show pim neighbors, and show pim statistics commands.

Related Documentation

- Configuring PIM Dense Mode on page 4107
- Configuring Basic PIM Settings on page 4093

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 181

Routing Content to Larger, Sparser Groups with PIM Sparse Mode

- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Static RP on page 4135
- Example: Configuring Anycast RP on page 4139
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Auto-RP on page 4151
- Configuring Embedded RP on page 4156
- Configuring PIM Filtering on page 4159
- Examples: Configuring PIM RPT and SPT Cutover on page 4166

Examples: Configuring PIM Sparse Mode

- Understanding PIM Sparse Mode on page 4113
- Designated Router on page 4116
- Tunnel Services PICs and Multicast on page 4116
- Enabling PIM Sparse Mode on page 4118
- Configuring PIM Join Load Balancing on page 4119
- Modifying the Join State Timeout on page 4122
- Example: Enabling Join Suppression on page 4122
- Example: Configuring PIM Sparse Mode over an IPsec VPN on page 4127
- Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces on page 4131

Understanding PIM Sparse Mode

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each

Copyright © 2016, Juniper Networks, Inc.

router's RPF interface until it reaches the RP. The RP router receives the (*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.

NOTE: State—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and * represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).

NOTE: If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

Copyright © 2016, Juniper Networks, Inc.

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (*,G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

  The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).
- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).

PIM sparse mode has standard features for all of these issues.

Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in Figure 162, the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

Figure 162: Rendezvous Point As Part of the RPT and SPT

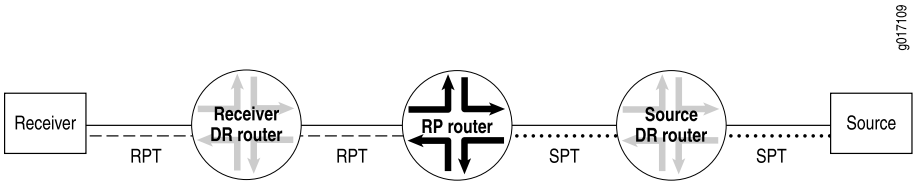

The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

Designated Router

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.

NOTE: In PIM dense mode (PIM-DM), a DR is elected by the same process that PIM-SM uses. However, the only time that a DR has any effect in PIM-DM is when IGMPv1 is used on the interface. (IGMPv2 is the default.) In this case, the DR also functions as the IGMP Query Router because IGMPv1 does not have a Query Router election mechanism.

Tunnel Services PICs and Multicast

On Juniper Networks routers, data packets are encapsulated and de-encapsulated into tunnels by means of hardware and not the software running on the router processor. The hardware used to create tunnel interfaces on M Series and T Series routers is a Tunnel Services PIC. If Juniper Networks M Series Multiservice Edge Routers and Juniper Networks T Series Core Routers are configured as rendezvous points or IP version 4 (IPv4) PIM

Copyright © 2016, Juniper Networks, Inc.

sparse-mode DRs connected to a source, a Tunnel Services PIC is required. Juniper Networks MX Series Ethernet Services Routers do not require Tunnel Services PICs. However, on MX Series routers, you must enable tunnel services with the **tunnel-services** statement on one or more online FPC and PIC combinations at the **[edit chassis fpc *number* pic *number*]** hierarchy level.

CAUTION: For redundancy, we strongly recommend that each routing device has multiple Tunnel Services PICs. In the case of MX Series routers, the recommendation is to configure multiple tunnel-services statements.

We also recommend that the Tunnel PICs be installed (or configured) on different FPCs. If you have only one Tunnel PIC or if you have multiple Tunnel PICs installed on a single FPC and then that FPC is removed, the multicast session will not come up. Having redundant Tunnel PICs on separate FPCs can help ensure that at least one Tunnel PIC is available and that multicast will continue working.

On MX Series routers, the redundant configuration looks like the following example:

```
[edit chassis]
user@mx-host# set fpc 1 pic 0 tunnel-services bandwidth 1g
user@mx-host# set fpc 2 pic 0 tunnel-services bandwidth 1g
```

In PIM sparse mode, the source DR takes the initial multicast packets and encapsulates them in PIM register messages. The source DR then unicasts the packets to the PIM sparse-mode RP router, where the PIM register message is de-encapsulated.

When a router is configured as a PIM sparse-mode RP router (by specifying an address using the **address** statement at the **[edit protocols pim rp local]** hierarchy level) and a Tunnel PIC is present on the router, a PIM register de-encapsulation interface, or **pd** interface, is automatically created. The **pd** interface receives PIM register messages and de-encapsulates them by means of the hardware.

If PIM sparse mode is enabled and a Tunnel Services PIC is present on the router, a PIM register encapsulation interface (**pe** interface) is automatically created for each RP address. The **pe** interface is used to encapsulate source data packets and send the packets to RP addresses on the PIM DR and the PIM RP. The **pe** interface receives PIM register messages and encapsulates the packets by means of the hardware.

Do not confuse the configurable **pe** and **pd** hardware interfaces with the nonconfigurable **pime** and **pimd** software interfaces. Both pairs encapsulate and de-encapsulate multicast packets, and are created automatically. However, the **pe** and **pd** interfaces appear only if a Tunnel Services PIC is present. The **pime** and **pimd** interfaces are not useful in situations requiring the **pe** and **pd** interfaces.

If the source DR is the RP, then there is no need for PIM register messages and consequently no need for a Tunnel Services PIC.

Copyright © 2016, Juniper Networks, Inc.

When PIM sparse mode is used with IP version 6 (IPv6), a Tunnel PIC is required on the RP, but not on the IPv6 PIM DR. The lack of a Tunnel PIC requirement on the IPv6 DR applies only to IPv6 PIM sparse mode and is not to be confused with IPv4 PIM sparse-mode requirements.

Table 383 shows the complete matrix of IPv4 and IPv6 PIM Tunnel PIC requirements.

Table 383: Tunnel PIC Requirements for IPv4 and IPv6 Multicast

IP Version	Tunnel PIC on RP	Tunnel PIC on DR
IPv4	Yes	Yes
IPv6	Yes	No

Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults. The following example explicitly configures PIMv2 on the interfaces.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

   ```
   [edit protocols pim]
   user@host# set rp local family inet address 192.168.3.253
   ```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

   ```
   [edit protocols pim]
   user@host# set interface all mode sparse
   user@host# set interface all version 2
   user@host# set interface fxp0.0 disable
   ```


Copyright © 2016, Juniper Networks, Inc.

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 198.58.3.253 version 2
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.

   - show pim interfaces
   - show pim join
   - show pim neighbors
   - show pim rps

Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Copyright © 2016, Juniper Networks, Inc.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the **show pim join extensive** or **show pim source** commands.

```
user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
    Source: *
    RP: 10.255.245.6
    Flags: sparse,rptree,wildcard
    Upstream interface: t1-0/2/3.0
    Upstream neighbor: 192.168.38.57
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: t1-0/2/1.0
            192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
    Source: *
    RP: 10.255.245.6
    Flags: sparse,rptree,wildcard
    Upstream interface: so-0/3/0.0
    Upstream neighbor: 192.168.38.47
    Upstream state: Join to RP
    Downstream neighbors:
        Interface: t1-0/2/3.0
            192.168.38.16 State: JOIN Flags; SRW Timeout: 164
```

   Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM join load balancing.

```
[edit protocols pim rp]
user@host# set static address 10.10.10.1
user@host# set interface all mode sparse version 2
user@host# set join-load-balance
```

   The static address is the address of the RP.

3. Monitor the operation.

   If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the **show pim interfaces** command.

```
user@host> show pim interfaces
```


Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.master

Name             Stat Mode        IP V State NbrCnt JoinCnt   DR address
lo0.0            Up   Sparse       4 2 DR         0       0   10.255.168.58
pe-1/2/0.32769   Up   Sparse       4 2 P2P        0       0
so-0/3/0.0       Up   Sparse       4 2 P2P        1       1
t1-0/2/1.0       Up   Sparse       4 2 P2P        1       0
t1-0/2/3.0       Up   Sparse       4 2 P2P        1       1
lo0.0            Up   Sparse       6 2 DR         0       0   fe80::2a0:a5ff:4b7
```

Note that the two equal-cost paths shown by the **show pim interfaces** command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the **show pim neighbors detail** output:

```
user@host> show pim neighbors detail
Interface: so-0/3/0.0

    Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1689116164
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 102 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 792890329
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: t1-0/2/3.0

    Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 678582286
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 97 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1854475503
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The **clear pim join-distribution** command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

Copyright © 2016, Juniper Networks, Inc.

Modifying the Join State Timeout

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the **join-prune-timeout** statement.

To modify the timeout, include the **join-prune-timeout** statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 240 seconds.

Example: Enabling Join Suppression

This example describes how to enable PIM join suppression.

- Requirements on page 4122
- Overview on page 4122
- Configuration on page 4125
- Verification on page 4126

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 4118.

Overview

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.

This example includes the following statements:

- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple

Copyright © 2016, Juniper Networks, Inc.

downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.

- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

  When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

Figure 163 shows the topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 163: Join Suppression

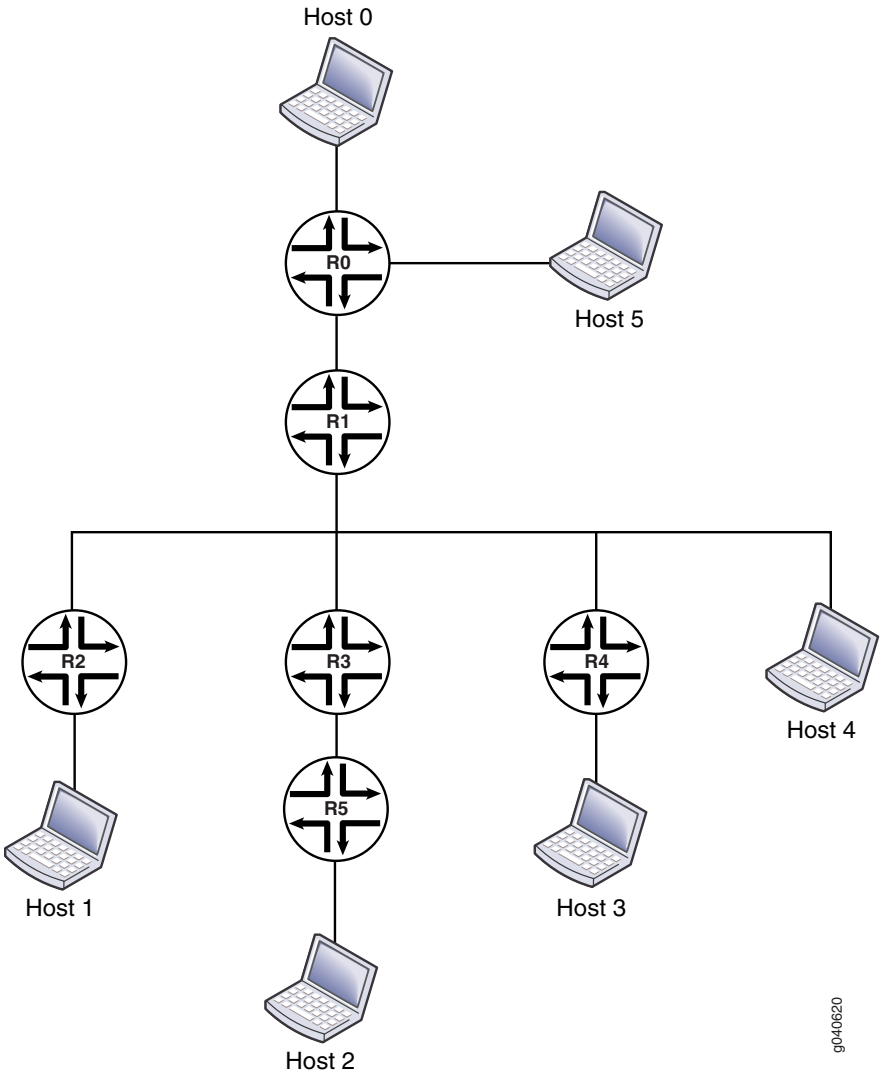

The items in Figure 163 represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```

Copyright © 2016, Juniper Networks, Inc.

5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
[edit protocols pim]
user@host# set traceoptions flag register detail
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols pim]
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
    traceoptions {
        file pim.log size 5m world-readable;
        flag join detail;
        flag prune detail;
        flag normal detail;
        flag register detail;
    }
    rp {
        static {
            address 10.255.112.160;
        }
    }
    interface all {
        mode sparse;
        version 2;
    }
    interface fxp0.0 {
        disable;
    }
    reset-tracking-bit;
    propagation-delay 500;
    override-interval 4000;
}
```

Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- show pim join **extensive**
- show multicast route **extensive**

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring PIM Sparse Mode over an IPsec VPN

IPsec VPNs create secure point-to-point connections between sites over the Internet. The Junos OS implementation of IPsec VPNs supports multicast and unicast traffic. The following example shows how to configure PIM sparse mode for the multicast solution and how to configure IPsec to secure your traffic.

The configuration shown in this example works on the following platforms:

- M Series and T Series routers with one of the following PICs:
  - Adaptive Services (AS) PIC
  - Multiservices (MS) PIC
- JCS1200 platform with a Multiservices PIC (MS-500)

The tunnel endpoints do not need to be the same platform type. For example, the device on one end of the tunnel can be a JCS1200 router, while the device on the other end can be a standalone T Series router. The two routers that are the tunnel endpoints can be in the same autonomous system or in different autonomous systems.

In the configuration shown in this example, OSPF is configured between the tunnel endpoints. In Figure 164, the tunnel endpoints are R0 and R1. The network that contains the multicast source is connected to R0. The network that contains the multicast receivers is connected to R1. R1 serves as the statically configured rendezvous point (RP).

Figure 164: PIM Sparse Mode over an IPsec VPN

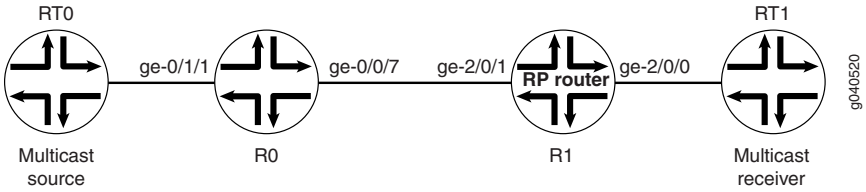

To configure PIM sparse mode with IPsec:

1. On R0, configure the incoming Gigabit Ethernet interface.

   ```
   [edit interfaces]
   user@host# set ge-0/1/1 description "incoming interface"
   user@host# set ge-0/1/1 unit 0 family inet address 10.20.0.1/30
   ```

2. On R0, configure the outgoing Gigabit Ethernet interface.

   ```
   [edit interfaces]
   user@host# set ge-0/0/7 description "outgoing interface"
   user@host# set ge-0/0/7 unit 0 family inet address 10.10.1.1/30
   ```

3. On R0, configure unit 0 on the **sp-** interface. The Junos OS uses unit 0 for service logging and other communication from the services PIC.

   ```
   [edit interfaces]
   user@host# set sp-0/2/0 unit 0 family inet
   ```

Copyright © 2016, Juniper Networks, Inc.

4. On R0, configure the logical interfaces that participate in the IPsec services. In this example, unit 1 is the inward-facing interface. Unit 1001 is the interface that faces the remote IPsec site.

   ```
   [edit interfaces]
   user@host# set sp-0/2/0 unit 1 family inet
   user@host# set sp-0/2/0 unit 1 service-domain inside
   user@host# set sp-0/2/0 unit 1001 family inet
   user@host# set sp-0/2/0 unit 1001 service-domain outside
   ```

5. On R0, direct OSPF traffic into the IPsec tunnel.

   ```
   [edit protocols ospf]
   user@host# set area 0.0.0.0 interface sp-0/2/0.1
   user@host# set parea 0.0.0.0 interface ge-0/1/1.0 passive
   user@host# set area 0.0.0.0 interface lo0.0
   ```

6. On R0, configure PIM sparse mode. This example uses static RP configuration. Because R0 is a non-RP router, configure the address of the RP router, which is the routable address assigned to the loopback interface on R1.

   ```
   [edit protocols pim]
   user@host# set rp static address 10.255.0.156
   user@host# set interfaces sp-0/2/0.1
   user@host# set interfaces ge-0/1/1.0
   user@host# set interfaces lo0.0
   ```

7. On R0, create a rule for a bidirectional dynamic IKE security association (SA) that references the IKE policy and the IPsec policy.

   ```
   [edit services ipsec-vpn rule ipsec_rule]
   user@host# set term ipsec_dynamic then remote-gateway 10.10.1.2
   user@host# set term ipsec_dynamic then dynamic ike-policy ike_policy
   user@host# set term ipsec_dynamic then dynamic ipsec-policy ipsec_policy
   user@host# set match-direction input
   ```

8. On R0, configure the IPsec proposal. This example uses the Authentication Header (AH) Protocol.

   ```
   [edit services ipsec-vpn ipsec proposal ipsec_prop]
   user@host# set protocol ah
   user@host# set authentication-algorithm hmac-md5-96
   ```

9. On R0, define the IPsec policy.

   ```
   [edit services ipsec-vpn ipsec policy ipsec_policy]
   user@host# set perfect-forward-secrecy keys group1
   user@host# set proposal ipsec_prop
   ```

10. On R0, configure IKE authentication and encryption details.

    ```
    [edit services ipsec-vpn ike proposal ike_prop]
    user@host# set authentication-method pre-shared-keys
    user@host# set dh-group group1
    user@host# set authentication-algorithm md5
    user@host# set authentication-algorithm 3des-cbc
    ```

11. On R0, define the IKE policy.

    ```
    [edit services ipsec-vpn ike policy ike_policy]
    user@host# set proposals ike_prop
    ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set pre-shared-key ascii-text "$ABC123"
```

12. On R0, create a service set that defines IPsec-specific information. The first command associates the IKE SA rule with IPsec. The second command defines the address of the local end of the IPsec security tunnel. The last two commands configure the logical interfaces that participate in the IPsec services. Unit 1 is for the IPsec inward-facing traffic. Unit 1001 is for the IPsec outward-facing traffic.

```
[edit services service-set ipsec_svc]
user@host# set ipsec-vpn-rules ipsec_rule
user@host# set ipsec-vpn-options local-gateway 10.10.1.1
user@host# set next-hop-service inside-service-interface sp-0/2/0.1
user@host# set next-hop-service outside-service-interface sp-0/2/0.1001
```

13. On R1, configure the incoming Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-2/0/1 description "incoming interface"
user@host# set ge-2/0/1 unit 0 family inet address 10.10.1.2/30
```

14. On R1, configure the outgoing Gigabit Ethernet interface.

```
[edit interfaces]
user@host# set ge-2/0/0 description "outgoing interface"
user@host# set ge-2/0/0 unit 0 family inet address 10.20.0.5/30
```

15. On R1, configure the loopback interface.

```
[edit interfaces]
user@host# set lo0.0 family inet address 10.255.0.156
```

16. On R1, configure unit 0 on the **sp-** interface. The Junos OS uses unit 0 for service logging and other communication from the services PIC.

```
[edit interfacesinterfaces]
user@host# set sp-2/1/0 unit 0 family inet
```

17. On R1, configure the logical interfaces that participate in the IPsec services. In this example, unit 1 is the inward-facing interface. Unit 1001 is the interface that faces the remote IPsec site.

```
[edit interfaces]
user@host# set sp-2/1/0 unit 1 family inet
user@host# set sp-2/1/0 unit 1 service-domain inside
user@host# set sp-2/1/0 unit 1001 family inet
user@host# set sp-2/1/0 unit 1001 service-domain outside
```

18. On R1, direct OSPF traffic into the IPsec tunnel.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface sp-2/1/0.1
user@host# set area 0.0.0.0 interface ge-2/0/0.0 passive
user@host# set area 0.0.0.0 interface lo0.0
```

19. On R1, configure PIM sparse mode. R1 is an RP router. When you configure the local RP address, use the shared address, which is the address of R1's loopback interface.

```
[edit protocols pim]
user@host# set rp local address 10.255.0.156
user@host# set interface sp-2/1/0.1
user@host# set interface ge-2/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface lo0.0 family inet
```

20. On R1, create a rule for a bidirectional dynamic Internet Key Exchange (IKE) security association (SA) that references the IKE policy and the IPsec policy.

```
[edit services ipsec-vpn rule ipsec_rule]
user@host# set term ipsec_dynamic from source-address 192.168.195.34/32
user@host# set term ipsec_dynamic then remote-gateway 10.10.1.1
user@host# set term ipsec_dynamic then dynamic ike-policy ike_policy
user@host# set term ipsec_dynamic then dynamic ipsec-policy ipsec_policy
user@host# set match-direction input
```

21. On R1, define the IPsec proposal for the dynamic SA.

```
[edit services ipsec-vpn ipsec proposal ipsec_prop]
user@host# set protocol ah
user@host# set authentication-algorithm hmac-md5-96
```

22. On R1, define the IPsec policy.

```
[edit services ipsec-vpn ipsec policy ipsec_policy]
user@host# set perfect-forward-secrecy keys group1
user@host# set proposal ipsec_prop
```

23. On R1, configure IKE authentication and encryption details.

```
[edit services ipsec-vpn ike proposal ike_prop]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group1
user@host# set authentication-algorithm md5
user@host# set authentication-algorithm 3des-cbc
```

24. On R0, define the IKE policy.

```
[edit services ipsec-vpn ike policy ike_policy]
user@host# set proposal ike_prop
user@host# set pre-shared-key ascii-text "$ABC123"
```

25. On R1, create a service set that defines IPsec-specific information. The first command associates the IKE SA rule with IPsec. The second command defines the address of the local end of the IPsec security tunnel. The last two commands configure the logical interfaces that participate in the IPsec services. Unit 1 is for the IPsec inward-facing traffic. Unit 1001 is for the IPsec outward-facing traffic.

```
[edit services service-set ipsec_svc]
user@host# set ipsec-vpn-rules ipsec_rule
user@host# set ipsec-vpn-options local-gateway 10.10.1.2
user@host# set next-hop-service inside-service-interface sp-2/1/0.1
user@host# set next-hop-service outside-service-interface sp-2/1/0.1001
```

To verify the configuration, run the following commands:

Check which RPs the various routers have learned about.

```
user@host> show pim rps extensive inet
```

Check that the IPsec SA negotiation is successful.

```
user@host> show services ipsec-vpn ipsec security-associations
```

Check that the IKE SA negotiation is successful.

Copyright © 2016, Juniper Networks, Inc.

```
user@host> show services ipsec-vpn ike security-associations
```

Check that traffic is traveling over the IPsec tunnel.

```
user@host> show services ipsec-vpn ipsec statistics
```

Example: Configuring Multicast for Virtual Routers with IPv6 Interfaces

A virtual router is a type of simplified routing instance that has a single routing table. This example shows how to configure PIM in a virtual router.

- Requirements on page 4131
- Overview on page 4131
- Configuration on page 4132
- Verification on page 4134

Requirements

Before you begin, configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.

Overview

You can configure PIM for the **virtual-router** instance type as well as for the **vrf** instance type. The **virtual-router** instance type is similar to the **vrf** instance type used with Layer 3 VPNs, except that it is used for non-VPN-related applications.

The **virtual-router** instance type has no VPN routing and forwarding (VRF) import, VRF export, VRF target, or route distinguisher requirements. The **virtual-router** instance type is used for non-Layer 3 VPN situations.

When PIM is configured under the **virtual-router** instance type, the VPN configuration is not based on RFC 2547, *BGP/MPLS VPNs*, so PIM operation does not comply with the Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*. In the **virtual-router** instance type, PIM operates in a routing instance by itself, forming adjacencies with PIM neighbors over the routing instance interfaces as the other routing protocols do with neighbors in the routing instance.

This example includes the following general steps:

1. On R1, configure a virtual router instance with three interfaces (**ge-0/0/0.0**, **ge-0/1/0.0**, and **ge-0/1/1.0**).
2. Configure PIM and the RP.
3. Configure an MLD static group containing interfaces **ge-0/1/0.0** and **ge-0/1/1.0**.

After you configure this example, you should be able to send multicast traffic from R2 through **ge-0/0/0** on R1 to the static group and verify that the traffic egresses from **ge-0/1/0.0** and **ge-0/1/1.0**.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Do not include the **group-address** statement for the **virtual-router** instance type.

Figure 165 shows the topology for this example.

Figure 165: Virtual Router Instance with Three Interfaces

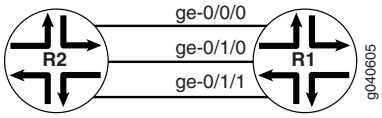

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet6 address 2001:4:4:4::1/64
set interfaces ge-0/1/0 unit 0 family inet6 address 2001:24:24:24::1/64
set interfaces ge-0/1/1 unit 0 family inet6 address 2001:7:7:7::1/64
set protocols mld interface ge-0/1/0.0 static group ff0e::10
set protocols mld interface ge-0/1/1.0 static group ff0e::10
set routing-instances mvrf1 instance-type virtual-router
set routing-instances mvrf1 interface ge-0/0/0.0
set routing-instances mvrf1 interface ge-0/1/0.0
set routing-instances mvrf1 interface ge-0/1/1.0
set routing-instances mvrf1 protocols pim rp local family inet6 address 2001:1:1:1::1
set routing-instances mvrf1 protocols pim interface ge-0/0/0.0
set routing-instances mvrf1 protocols pim interface ge-0/1/0.0
set routing-instances mvrf1 protocols pim interface ge-0/1/1.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multicast for virtual routers:

1. Configure the interfaces.

```
[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family inet6 address 2001:4:4:4::1/64
[edit interfaces]
user@host# set ge-0/1/0 unit 0 family inet6 address 2001:24:24:24::1/64
[edit interfaces]
user@host# set ge-0/1/1 unit 0 family inet6 address 2001:7:7:7::1/64
[edit interfaces]
user@host# exit
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure the routing instance type.

```
[edit]
user@host# edit routing-instances
[edit routing-instances]
user@host# set mvrf1 instance-type virtual-router
```

3. Configure the interfaces in the routing instance.

```
[edit routing-instances]
user@host# set mvrf1 interface ge-0/0/0
[edit routing-instances]
user@host# set mvrf1 interface ge-0/1/0
[edit routing-instances]
user@host# set mvrf1 interface ge-0/1/1
```

4. Configure PIM and the RP in the routing instance.

```
[edit routing-instances]
user@host# set mvrf1 protocols pim rp local family inet6 address 2001:1:1:1::1
```

5. Configure PIM on the interfaces.

```
[edit routing-instances]
user@host# set mvrf1 protocols pim interface ge-0/0/0
[edit routing-instances]
user@host# set mvrf1 protocols pim interface ge-0/1/0
[edit routing-instances]
user@host# set mvrf1 protocols pim interface ge-0/1/1
[edit routing-instances]
user@host# exit
```

6. Configure the MLD group.

```
[edit]
user@host# edit protocols mld
[edit protocols mld]
user@host# set interface ge-0/1/0.0 static group ff0e::10
[edit protocols mld]
user@host# set interface ge-0/1/1.0 static group ff0e::10
```

7. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show interfaces**, **show routing-instances**, and **show protocols** commands.

```
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet6 {
            address 2001:4:4:4::1/64;
        }
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
  ge-0/1/0 {
    unit 0 {
      family inet6 {
        address 2001:24:24:24::1/64;
      }
    }
  }
  ge-0/1/1 {
    unit 0 {
      family inet6 {
        address 2001:7:7:7::1/64;
      }
    }
  }
```

```
user@host# show routing-instances
mvrf1 {
  instance-type virtual-router;
  interface ge-0/0/0.0;
  interface ge-0/1/0.0;
  interface ge-0/1/1.0;
  protocols {
    pim {
      rp {
        local {
          family inet6 {
            address 2001:1:1:1::1;
          }
        }
      }
      interface ge-0/0/0.0;
      interface ge-0/1/0.0;
      interface ge-0/1/1.0;
    }
  }
}
```

```
user@host# show protocols
mld {
  interface ge-0/1/0.0 {
    static {
      group ff0e::10;
    }
  }
  interface ge-0/1/1.0 {
    static {
      group ff0e::10;
    }
  }
}
```

Verification

To verify the configuration, run the following commands:

Copyright © 2016, Juniper Networks, Inc.

- show mld group
- show mld interface
- show mld statistics
- show multicast interface
- show multicast route
- show multicast rpf
- show pim interfaces
- show pim join
- show pim neighbors
- *show route forwarding-table*
- *show route instance*
- show route table

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Dense Mode on page 4107
- Configuring a Designated Router for PIM on page 4103
- Configuring PIM Filtering on page 4159
- Configuring PIM Sparse-Dense Mode on page 4110
- Configuring Basic PIM Settings on page 4093

Configuring Static RP

- Understanding Static RP on page 4135
- Configuring Local PIM RPs on page 4136
- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4137

Understanding Static RP

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

Copyright © 2016, Juniper Networks, Inc.

Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

   ```
   [routing-instances VPN-A protocols pim]
   user@host# set rp local
   ```

2. Configure the IP protocol family and IP address.

   IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface *interface-name*]** hierarchy level and **family inet6** at the **[edit protocols pim interface *interface-name*]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

   By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
   user@host# set family inet address 10.1.2.254
   ```

3. (IPv4 only) Configure the routing device's RP priority.

   NOTE: The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set priority 5
   ```

4. Configure the groups for which the routing device is the RP.

Copyright © 2016, Juniper Networks, Inc.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set group-ranges fec0::/10
user@host# set group-ranges 10.1.2.0/24
```

5. (IPv4 only) Modify the local RP hold time.

   If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set hold-time 200
   ```

6. (Optional) Override dynamic RP for the specified group address range.

   If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

   If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

   ```
   [edit routing-instances VPN-A protocols pim rp local]
   user@host# set override
   ```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

Configuring the Static PIM RP Address on the Non-RP Routing Device

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

Copyright © 2016, Juniper Networks, Inc.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

   ```
   [routing-instances VPN-A protocols pim rp]
   user@host# set static address 2001:db8:85a3::8a2e:370:7334
   ```

   NOTE: Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

   For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

   ```
   [edit routing-instances VPN-A protocols pim rp]
   user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
   ```

   NOTE: The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode (**[edit pim rp static address *address*]**). PIM version 2 is the default for interface mode (**[edit pim interface *interface-name*]**). Explicitly configured versions override the defaults.

3. (Optional) Set the group address range.

   By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

   ```
   [edit routing-instances VPN-A protocols pim rp]
   user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
   ```

   The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

   If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

   If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-instances VPN-A protocols pim rp static address
  2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring a Designated Router for PIM on page 4103
- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Basic PIM Settings on page 4093

Example: Configuring Anycast RP

- Understanding RP Mapping with Anycast RP on page 4139
- Example: Configuring Multiple RPs in a Domain with Anycast RP on page 4140
- Example: Configuring PIM Anycast With or Without MSDP on page 4142
- Configuring a PIM Anycast RP Router Using Only PIM on page 4146

Understanding RP Mapping with Anycast RP

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in Internet draft draft-ietf-mboned-anycast-rp-08.txt, *Anycast RP Mechanism Using PIM and MSDP*. To access Internet RFCs and drafts, go to the IETF website at http://www.ietf.org .

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Multiple RPs in a Domain with Anycast RP

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

- Requirements on page 4140
- Overview on page 4140
- Configuration on page 4140
- Verification on page 4142

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 4118.

Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

RP Routers

```
set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary
set interfaces lo0 unit 0 family inet address 10.1.1.2/32
set protocols msdp local-address 192.168.132.1
set protocols msdp peer 192.168.12.1
set protocols pim rp local address 10.1.1.2
set routing-options router-id 192.168.132.1
```

Non-RP Routers

```
set protocols pim rp static address 10.1.1.2
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

   ```
   [edit interfaces]
   user@host# set lo0 unit 0 family inet address 10.1.1.2/32
   ```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

   ```
   [edit interfaces]
   user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary

   [edit routing-options]
   user@host# set router-id 192.168.132.1
   ```

3. On each RP router in the domain, configure the local RP address, using the shared address.

   ```
   [edit protocols pim]
   user@host# set rp local address 10.1.1.2
   ```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

   ```
   [edit protocols msdp]
   user@host# set local-address 192.168.132.1
   user@host# set peer 192.168.12.1
   ```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

   ```
   [edit protocols pim]
   user@host# set rp static address 10.1.1.2
   ```

6. If you are done configuring the devices, commit the configuration.

   ```
   user@host# commit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
lo0 {
   unit 0 {
      family inet {
         address 192.168.132.1/32 {
            primary;
         }
         address 10.1.1.2/32;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
}
```

On the RP routers:

```
user@host# show protocols
msdp {
    local-address 192.168.132.1;
    peer 192.168.12.1;
}
pim {
    rp {
        local {
            address 10.1.1.2;
        }
    }
}
```

On the non-RP routers:

```
user@host# show protocols
pim {
    rp {
        static {
            address 10.1.1.2;
        }
    }
}

user@host# show routing-options
router-id 192.168.132.1;
```

Verification

To verify the configuration, run the show pim rps **extensive inet** command.

Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

Copyright © 2016, Juniper Networks, Inc.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

   ```
   interfaces {
     lo0 {
       description "PIM RP";
       unit 0 {
         family inet {
           address 198.58.3.254/32;
           primary;
           address 198.58.3.253/32;
         }
       }
     }
   }
   ```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

   ```
   protocols {
     pim {
       rp {
         local {
           family inet;
           address 198.58.3.253;
         }
         interface all {
           mode sparse;
           version 2;
         }
         interface fxp0.0 {
           disable;
         }
       }
   ```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address address 198.58.3.254;
    }
  }
}
```

NOTE: If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}
```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

   For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by Including the **disable** statement for that interface.

   Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP

Copyright © 2016, Juniper Networks, Inc.

address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
```

Copyright © 2016, Juniper Networks, Inc.

```
            mode sparse;
            version 2;
          }
          interface fxp0.0 {
            disable;
          }
        }
      }
```

Configuring a PIM Anycast RP Router Using Only PIM

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}
```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
```


Copyright © 2016, Juniper Networks, Inc.

```
                address 198.58.3.240;
                address 198.58.3.241 forward-msdp-sa;
              }
              local-address 198.58.3.254; #If not configured, use lo0 primary
            }
          }
        }
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring a Designated Router for PIM on page 4103
- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Basic PIM Settings on page 4093

Configuring PIM Bootstrap Router

- Understanding the PIM Bootstrap Router on page 4147
- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 4150
- Example: Configuring PIM BSR Filters on page 4151

Understanding the PIM Bootstrap Router

To determine which routing device is the rendezvous point (RP), all routing devices within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routing devices that all share the same RP router. The domain bootstrap routing device initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap routing device when necessary.

Configuring PIM Bootstrap Properties for IPv4

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap

Copyright © 2016, Juniper Networks, Inc.

router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

NOTE: For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routers: one section for both IPv4 and IPv6, and this section, which is for IPv4 only. The method described in "Configuring PIM Bootstrap Properties for IPv4 or IPv6" on page 4149 is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is "duplicate IPv4 bootstrap configuration."

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

   By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. A priority of 0 disables the function for IPv4 and does not cause the routing device to send bootstrap router packets with a 0 in the priority field. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.

   ```
   [edit protocols pim rp]
   user@host# set bootstrap-priority 3
   ```

2. (Optional) Create import and export policies to control the flow of IPv4 bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **bootstrap-import** statement prevents messages from being imported into the RP. The **bootstrap-export** statement prevents messages from being exported from the RP.

   ```
   [edit protocols pim rp]
   user@host# set bootstrap-import pim-bootstrap-import
   user@host# set bootstrap-export pim-bootstrap-export
   ```

3. Configure the policies.

   ```
   [edit policy-options policy-statement pim-bootstrap-import]
   user@host# set from interface se-0/0/0
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set then reject

[edit policy-options policy-statement pim-bootstrap-export]
user@host# set from interface se-0/0/0
user@host# set then reject
```

4. Monitor the operation of PIM bootstrap routing devices by running the **show pim bootstrap** command.

Configuring PIM Bootstrap Properties for IPv4 or IPv6

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

NOTE: For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routers: one section for IPv4 only, and this section, which is for both IPv4 and IPv6. The method described in this section is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is “duplicate IPv4 bootstrap configuration.”

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

   By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.

Copyright © 2016, Juniper Networks, Inc.

NOTE: In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the configuration shown here, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the **bootstrap** statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
```

3. Configure the policies.

```
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
```

4. Monitor the operation of PIM bootstrap routing devices by running the **show pim bootstrap** command.

Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```
protocols {
  pim {
    rp {
      bootstrap {
        family inet {
          priority 1;
          import pim-import;
          export pim-export;
        }
        family inet6 {
          priority 1;
```


Copyright © 2016, Juniper Networks, Inc.

```
            import pim-import;
            export pim-export;
          }
        }
      }
    }
  }
  policy-options {
    policy-statement pim-import {
      from interface so-0/1/0;
      then reject;
    }
    policy-statement pim-export {
      to interface so-0/1/0;
      then reject;
    }
  }
```

Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
  pim {
    rp {
      bootstrap-import no-bsr;
      bootstrap-export no-bsr;
    }
  }
}
policy-options {
  policy-statement no-bsr {
    then reject;
  }
}
```

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring a Designated Router for PIM on page 4103
- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Basic PIM Settings on page 4093

Configuring PIM Auto-RP

- Understanding PIM Auto-RP on page 4151
- Configuring PIM Auto-RP on page 4152

Understanding PIM Auto-RP

You can configure a more dynamic way of assigning rendezvous points (RPs) in a multicast network by means of auto-RP. When you configure auto-RP for a router, the router learns

Copyright © 2016, Juniper Networks, Inc.

the address of the RP in the network automatically and has the added advantage of operating in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function that typically uses dense mode PIM to advertise control traffic, it provides an important failover advantage that simple static RP assignment does not. You can configure multiple routers as RP candidates. If the elected RP fails, one of the other preconfigured routers takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Configuring PIM Auto-RP

For correct operation, every multicast within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The auto-RP mechanism is one way that a multicast can learn the set of group-to-RP mappings. Auto-RP automatically distributes mapping information to routing devices. It simplifies use of multiple RPs for different multicast group ranges, thus allowing multiple RPs to act as backups for each other. Auto-RP relies on a to act as the RP mapping agent. Potential RPs announce themselves to the mapping agent, and the mapping agent resolves any conflicts.

The mapping agent sends the multicast group-RP mapping information to the other s using PIM dense mode. The specific groups used are 224.0.1.39 and .40. The first (.39) is used to advertise, the second (.40) is used for discovery. Because PIM dense mode is necessary to enable auto-RP to work, which in turns enables PIM sparse mode to work, you must configure PIM sparse-dense mode in the PIM domains that use auto-RP.

Although auto-RP is a nonstandard (non-RFC-based) function requiring dense mode PIM to advertise control traffic, it provides an important failover advantage that static RP assignment does not. That is, you can configure multiple routing devices as RP candidates. If the elected RP fails, one of the other preconfigured routing devices takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Auto-RP operates in PIM version 1 and version 2.

In most cases, how the routing device handles auto-RP discovery, announce, or mapping messages depends on whether the routing device is an RP (configured as local RP) or not. Table 384 shows how the routing device behaves depending on the local RP configuration.

Table 384: Local RP and Auto-RP Message Types

Auto-RP Message Type	Local RP?	Routing Device Behavior
discovery	No	Listen for auto-RP mapping messages.
discovery	Yes	Listen for auto-RP mapping messages.
announce	No	Listen for auto-RP mapping messages.
announce	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages.

Copyright © 2016, Juniper Networks, Inc.

Table 384: Local RP and Auto-RP Message Types *(continued)*

Auto-RP Message Type	Local RP?	Routing Device Behavior
mapping	No	Listen for auto-RP mapping messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.
mapping	Yes	Listen for auto-RP mapping messages. Send auto-RP announce messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.

NOTE: If the routing device receives auto-RP announcements split across multiple messages, the routing device loses the information in the previous part of the message as soon as the next part of the message is received.

You can configure auto-RP properties globally or for a routing instance. This example shows the global configuration.

To configure auto-RP properties:

1. Configure PIM in sparse-dense mode on all routing devices in the PIM domain.

   ```
   [edit protocols pim]
   user@host# edit
   user@host# set interface all mode sparse-dense
   ```

   This configuration allows the routing device to operate in sparse mode for most groups and dense mode for others. The default is to operate in sparse mode unless the routing device is specifically informed of a dense mode group.

2. Configure a routable loopback interface address on all routing devices in the PIM domain.

   The routing device joins the auto-RP groups on the configured interfaces and on the loopback interface **lo0.0**. For auto-RP to work correctly, configure a routable IP address on the loopback interface. The router ID is used as the address for auto-RP updates. You cannot use the loopback address 127.0.0.1. Also, you must enable PIM sparse-dense mode on the **lo0.0** interface if you do not specify **interface all**.

   ```
   [edit interfaces lo0.0 unit 0 family inet]
   user@host# set address 192.168.0.3 preferred
   ```

3. Configure the two multicast dense groups on all the routing devices.

   Auto-RP requires multicast flooding to announce potential RP candidates and to discover the elected RPs in the network. Multicast flooding occurs through a PIM dense mode model, where group 224.0.1.39 is used for **announce** messages and group 224.0.1.40 is used for **discovery** messages.

   ```
   [edit protocols pim]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set dense-groups 224.0.1.39/32
user@host# set dense-groups 224.0.1.40/32
```

TIP: Step 3 is required. When auto-RP is enabled, the auto-RP announce group (224.0.1.39) and auto-RP-discovery group (224.0.1.40) must be configured explicitly as dense groups. When the auto-RP discovery group is not configured as a dense group, auto-RP is not enabled. When the auto-RP announce group is not configured as a dense group, auto-RP is enabled in the discovery mode only, and mapping and announce modes are disabled.

4. Configure the auto-RP **announce** option.

   At least one routing device in the PIM domain must announce auto-RP messages and at least one must map them, or you can configure a routing device to perform both functions.

   When a routing device sends announce messages in the network, it is advertising itself as a candidate RP. A routing device configured with this option must also be configured as an RP, or announce messages are not sent.

   ```
   [edit protocols pim rp]
   user@host# set local address 192.168.0.1
   user@host# set auto-rp announce
   ```

   NOTE: You cannot include the **auto-rp announce** option at the **[edit logical-systems** *logical-system-name* **routing-instances** *routing-instance-name* **protocols pim]** hierarchy level.

5. Configure the auto-RP mapping agent.

   The mapping agent sends discovery messages to the network, informing all routing devices in a multicast group of which RP to use. If the mapping agent is also an RP, the **mapping** option also allows the routing device to send auto-RP announcements (mapping on an RP allows the routing device to perform both the announcement and mapping functions).

   ```
   [edit protocols pim rp]
   user@host# set auto-rp mapping
   ```

   If the mapping agent is also an RP, configure the mapping agent as a local RP.

   ```
   [edit protocols pim rp]
   user@host# set local address 192.168.0.2
   ```

6. Configure mapping agent election.

   If you configure the **mapping** option on more than one routing device in the PIM domain, configure mapping agent election on each potential mapping agent.

   Auto-RP specifications state that mapping agents do not send mapping messages if they receive messages from a mapping agent with a higher IP address. However,

Copyright © 2016, Juniper Networks, Inc.

some vendors' mapping agents continue to announce mappings, even in the presence of higher-addressed mapping agents. In other words, some mapping agents will always send mapping messages.

The default auto-RP operation is to perform mapping agent election. To explicitly configure mapping agent election, you can include the **mapping-agent-election** statement. When this option is configured, the mapping agent will stop sending mapping messages if it receives messages from a mapping agent with a higher IP address.

```
[edit protocols pim rp]
user@host# set auto-rp mapping mapping-agent-election
```

Mapping message suppression is disabled with the **no-mapping-agent-election** statement. When this option is configured, the mapping agent will always send mapping messages even in the presence of higher-addressed mapping agents.

To disable mapping agent election for compatibility with other vendors' equipment, include the **no-mapping-agent-election** statement.

```
[edit protocols pim rp]
user@host# set auto-rp mapping no-mapping-agent-election
```

7. Configure the remaining routing devices in the PIM domain to discover the RP.

   Discovery enables the routing devices to receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.

   ```
   [edit protocols pim rp]
   user@host# set auto-rp discovery
   ```

8. Monitor the operation of PIM auto-RP routers by running the following commands:

   - show pim interfaces
   - show pim rps

9. Issue the **show pim rps extensive** command to see information about how an RP is learned, what groups it handles, and the number of groups actively using the RP.

   ```
   user@host> show pim rps extensive
   RP: 192.168.5.1
   Learned from 192.168.5.1 via: auto-rp
   Time Active: 00:34:29
   Holdtime: 150 with 108 remaining
   Device Index: 6
   Subunit: 32769
   Interface: pd-0/0/0.32769
   Group Ranges:
           224.0.0.0/4
   Active groups using RP:
           224.2.2.100
           total 1 groups active
   Register State for RP:
   Group     Source FirstHop       RP Address      StateRP address Type Holdtime
    Timeout
   ```

   In the example, the RP at 192.168.5.1 was learned through auto-RP. The RP is able to support all groups in the 224.0.0.0/4 range (all possible groups). The local router has sent PIM control traffic for the 224.2.2.100 group to the RP.

Copyright © 2016, Juniper Networks, Inc.

Additionally, the presence of a Tunnel Physical Interface Card (PIC) in an RP router creates a de-encapsulation interface, which allows the RP to receive multicast traffic from the source. This interface is indicated by **pd-0/0/0.32769**.

Related Documentation

- Configuring PIM Bootstrap Router on page 4147
- Configuring a Designated Router for PIM on page 4103
- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Basic PIM Settings on page 4093

Configuring Embedded RP

- Understanding Embedded RP for IPv6 Multicast on page 4156
- Configuring PIM Embedded RP for IPv6 on page 4158

Understanding Embedded RP for IPv6 Multicast

Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. This feature embeds the RP address in an IPv6 multicast address.

All IPv6 multicast addresses begin with 8 1-bits (1111 1111) followed by a 4-bit flag field normally set to 0011. The flag field is set to 0111 when embedded RP is used. Then the low-order bits of the normally reserved field in the IPv6 multicast address carry the 4-bit RP interface identifier (RIID).

When the IPv6 address of the RP is embedded in a unicast-prefix-based any-source multicast (ASM) address, all of the following conditions must be true:

- The address must be an IPv6 multicast address and have 0111 in the flags field (that is, the address is part of the prefix FF70::/12).
- The 8-bit prefix length (plen) field must not be all 0. An all 0 plen field implies that SSM is in use.
- The 8-bit prefix length field value must not be greater than 64, which is the length of the network prefix field in unicast-prefix-based ASM addresses.

The routing platform derives the value of the interdomain RP by copying the prefix length field number of bits from the 64-bit network prefix field in the received IPv6 multicast address to an empty 128-bit IPv6 address structure and copying the last bits from the 4-bit RIID. For example, if the prefix length field bits have the value 32, then the routing platform copies the first 32 bits of the IPv6 multicast address network prefix field to an all-0 IPv6 address and appends the last four bits determined by the RIID. See Figure 166 for an illustration of this process.

Copyright © 2016, Juniper Networks, Inc.

Figure 166: Extracting the Embedded RP IPv6 Address

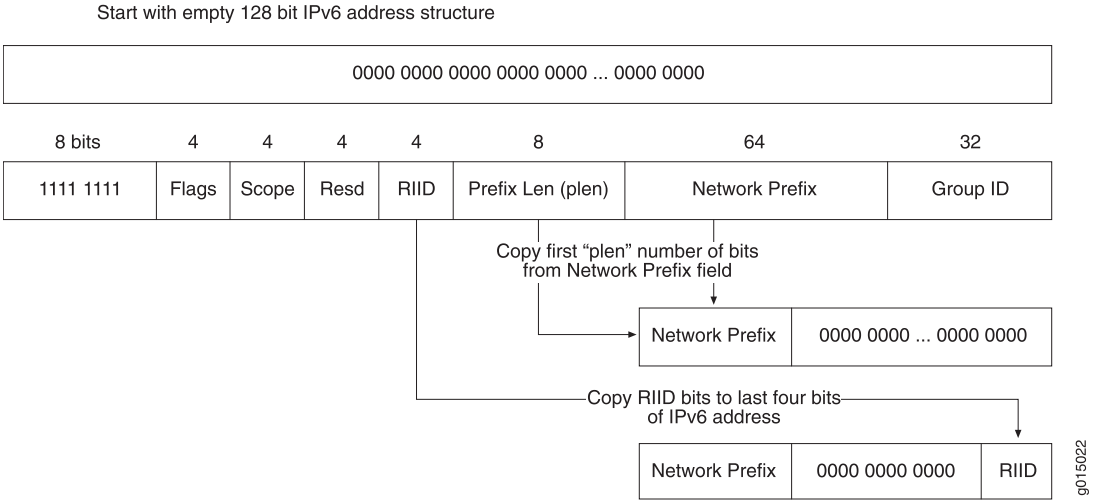

For example, the administrator of IPv6 network 2001:DB8::/32 sets up an RP for the 2001:DB8:BEEF:FEED::/96 subnet. In that case, the received embedded RP IPv6 ASM address has the form:

FF70:y40:2001:DB8:BEEF:FEED::/96

and the derived RP IPv6 address has the form:

2001:DB8:BEEF:FEED::y

where **y** is the RIID (**y** cannot be 0).

When configured, the routing platform checks for embedded RP information in every PIM join request received for IPv6. The use of embedded RP does not change the processing of IPv6 multicast and RPs in any way, except that the embedded RP address is used if available and selected for use. There is no need to specify the IPv6 address family for embedded RP configuration because the information can be used only if IPv6 multicast is properly configured on the routing platform.

The following receive events trigger extraction of an IPv6 embedded RP address on the routing platform:

- Multicast Listener Discovery (MLD) report for an embedded RP multicast group address
- PIM join message with an embedded RP multicast group address
- Static embedded RP multicast group address associated with an interface
- Packets sent to an embedded RP multicast group address received on the DR

The embedded RP node discovered through these events is added if it does not already exist on the routing platform. The routing platform chooses the embedded RP as the RP for a multicast group before choosing an RP learned through BSRs or a statically configured RP. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.

Copyright © 2016, Juniper Networks, Inc.

Configuring PIM Embedded RP for IPv6

You configure embedded RP to allow multidomain IPv6 multicast networks to find RPs in other routing domains. Embedded RP embeds an RP address inside PIM join messages and other types of messages sent between routing domains. Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. Thus, embedded RP enables you can deploy IPv6 with any-source multicast (ASM).

Embedded RP is disabled by default.

When you configure embedded RP for IPv6, embedded RPs are preferred to RPs discovered by IPv6 any other way. You configure embedded RP independent of any other IPv6 multicast properties. This feature is applied only when IPv6 multicast is properly configured.

You can configure embedded RP globally or for a routing instance. This example shows the routing instance configuration.

To configure embedded RP for IPv6 PIM sparse mode:

1. Define which multicast addresses or prefixes can embed RP address information. If messages within a group range contain embedded RP information and the group range is not configured, the embedded RP in that group range is ignored. Any valid unicast-prefix-based ASM address can be used as a group range. The default group range is FF70::/12 to FFF0::/12. Messages with embedded RP information that do not match any configured group ranges are treated as normal multicast addresses.

   ```
   [edit routing-instances vpn-A protocols pim rp embedded-rp]
   user@host# set group-ranges fec0::/10
   ```

   If the derived RP address is not a valid IPv6 unicast address, it is treated as any other multicast group address and is not used for RP information. Verification fails if the extracted RP address is a local interface, unless the routing device is configured as an RP and the extracted RP address matches the configured RP address. Then the local RP determines whether it is configured to act as an RP for the embedded RP multicast address.

2. Limit the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.

   ```
   [edit routing-instances vpn-A protocols pim rp]
   user@host# set maximum-rps 50
   ```

3. Monitor the operation by running the **show pim rps** and **show pim statistics** commands.

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring a Designated Router for PIM on page 4103

Copyright © 2016, Juniper Networks, Inc.

- Examples: Configuring PIM Sparse Mode on page 4113
- Configuring Basic PIM Settings on page 4093

Configuring PIM Filtering

- Understanding Multicast Message Filters on page 4159
- Filtering MAC Addresses on page 4160
- Filtering RP and DR Register Messages on page 4160
- Filtering MSDP SA Messages on page 4161
- Configuring Interface-Level PIM Neighbor Policies on page 4161
- Filtering Outgoing PIM Join Messages on page 4162
- Filtering Incoming PIM Join Messages on page 4163
- Configuring Register Message Filters on a PIM RP and DR on page 4164

Understanding Multicast Message Filters

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.

NOTE: If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.

Copyright © 2016, Juniper Networks, Inc.

NOTE: If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

Copyright © 2016, Juniper Networks, Inc.

- Group multicast address
- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.

NOTE: When you apply firewall filters, firewall action modifiers, such as log, sample, and count, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor

Copyright © 2016, Juniper Networks, Inc.

hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

   ```
   [edit policy-options]
   user@host# set prefix-list nbrGroup 1 20.20.20.1/32
   user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
   user@host# set policy-statement nbr-policy then reject
   ```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

   ```
   [edit routing-instances PIM.master protocols pim]
   user@host# set neighbor-policy nbr-policy
   ```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

   ```
   user@host# set policy-options policy-statement block-groups term t1 from route-filter
     224.0.1.2/32 exact
   user@host# set policy-options policy-statement block-groups term t1 from route-filter
     225.1.1.1/32 exact
   user@host# set policy-options policy-statement block-groups term t1 then reject
   user@host# set policy-options policy-statement block-groups term last then accept
   ```


Copyright © 2016, Juniper Networks, Inc.

2. Verify the policy configuration by running the **show policy-options** command.

```
user@host# show policy-options
policy-statement block-groups {
  term t1 {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 225.1.1.1/32 exact;
      then reject;
    }
    term last {
      then accept;
    }
  }
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```
user@host>  show pim statistics | grep filtered
RP Filtered Source                  0

Rx Joins/Prunes filtered            0

Tx Joins/Prunes filtered            254
```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See Table 385 for a list of match conditions.

Table 385: PIM Join Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address
neighbor	Neighbor address (the source address in the IP header of the join and prune message)

Copyright © 2016, Juniper Networks, Inc.

Table 385: PIM Join Filter Match Conditions *(continued)*

Match Condition	Matches On
route-filter	Multicast group address embedded in the join and prune message
source-address-filter	Multicast source address embedded in the join and prune message

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (*,G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

   ```
   [edit policy-statement pim-join-filter term bad-groups]
   user@host# set from route-filter 224.0.1.2/32 exact
   user@host# set from route-filter 239.0.0.0/8 orlonger
   user@host# set then reject

   [edit policy-statement pim-join-filter term bad-sources]
   user@host# set from source-address-filter 10.0.0.0/8 orlonger
   user@host# set from source-address-filter 127.0.0.0/8 orlonger
   user@host# set then reject

   [edit policy-statement pim-join-filter term last]
   user@host# set then accept
   ```

2. Apply one or more policies to routes being imported into the routing table from PIM.

   ```
   [edit protocols pim]
   user@host# set import pim-join-filter
   ```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is

Copyright © 2016, Juniper Networks, Inc.

provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP routing device.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

   The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

Copyright © 2016, Juniper Networks, Inc.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit

[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [ reject_224_1_1_1 | accept_224_1_1_5 ]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- **RP Filtered Source**
- **Rx Joins/Prunes filtered**
- **Tx Joins/Prunes filtered**
- **Rx Register msgs filtering drop**
- **Tx Register msgs filtering drop**

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Dense Mode on page 4107
- Configuring a Designated Router for PIM on page 4103
- Examples: Configuring PIM RPT and SPT Cutover on page 4166
- Configuring PIM Sparse-Dense Mode on page 4110
- Configuring Basic PIM Settings on page 4093

Examples: Configuring PIM RPT and SPT Cutover

- Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 4167
- Building an RPT Between the RP and Receivers on page 4168
- PIM Sparse Mode Source Registration on page 4169
- Multicast Shortest-Path Tree on page 4171
- SPT Cutover on page 4172
- SPT Cutover Control on page 4175

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring the PIM Assert Timeout on page 4175
- Example: Configuring the PIM SPT Threshold Policy on page 4177

Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.

To join the shared tree, or *rendezvous-point tree* (*RPT*) as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.
- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (*,G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (*,G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (*,G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined

Copyright © 2016, Juniper Networks, Inc.

- A router along the RPT that already has a multicast forwarding state for the group that is being joined

In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to "migrate" a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (*,G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table. For more information about the RPF table, see "Understanding Multicast Reverse Path Forwarding" on page 4315.

Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see Figure 167). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.
3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

Figure 167: Building an RPT Between the RP and the Receiver

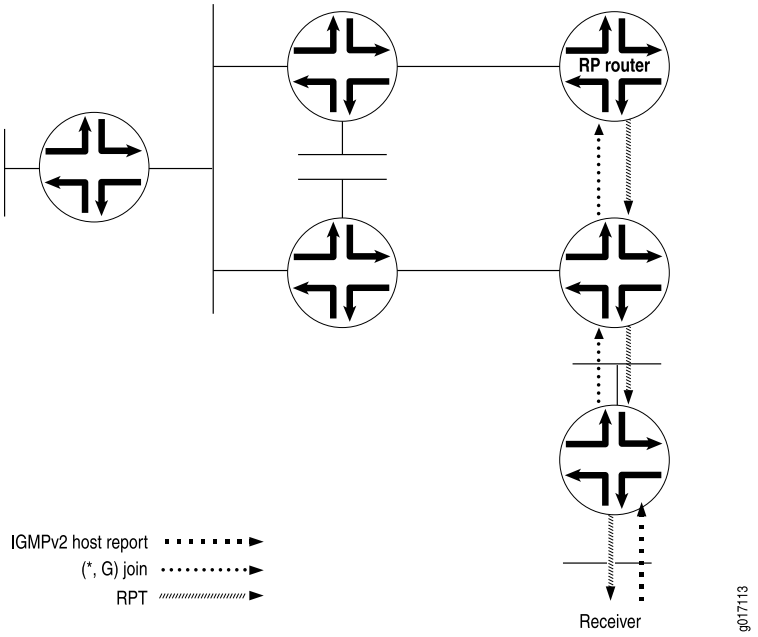

Copyright © 2016, Juniper Networks, Inc.

PIM Sparse Mode Source Registration

The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:

1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see Figure 168).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.

Figure 168: PIM Register Message and PIM Join Message Exchanged

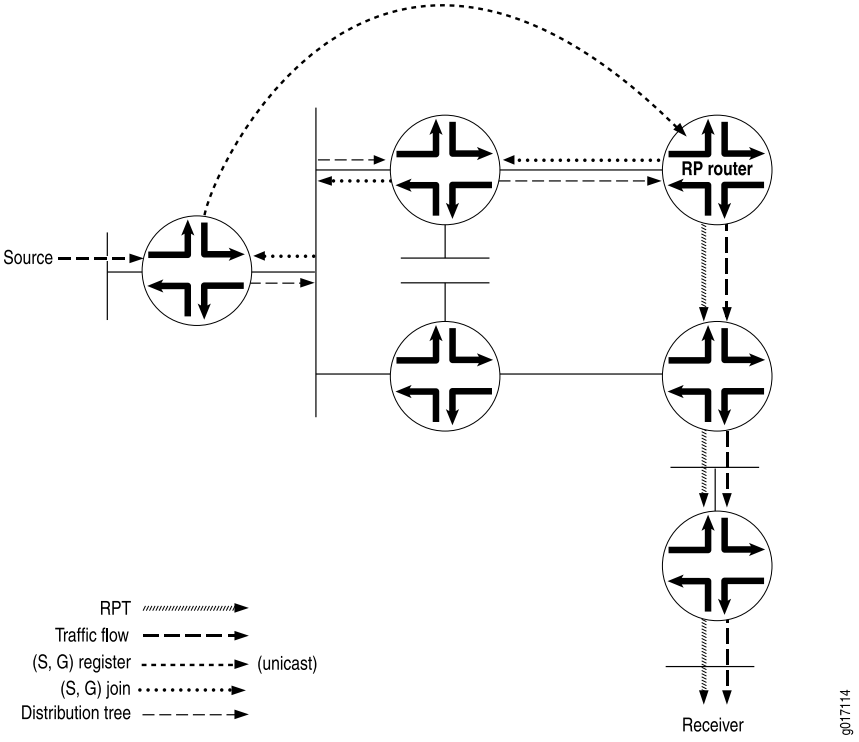

3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see Figure 169).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Copyright © 2016, Juniper Networks, Inc.

Figure 169: Traffic Sent from the Source to the RP Router

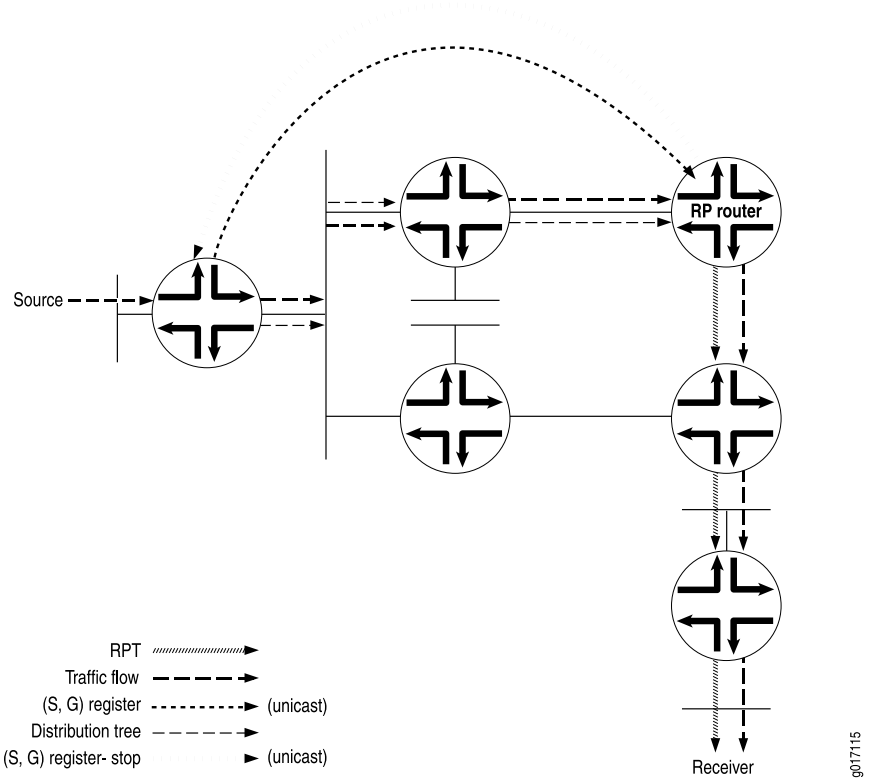

5. The RP router sends the multicast traffic down the RPT toward the receiver (see Figure 170).

Figure 170: Traffic Sent from the RP Router Toward the Receiver

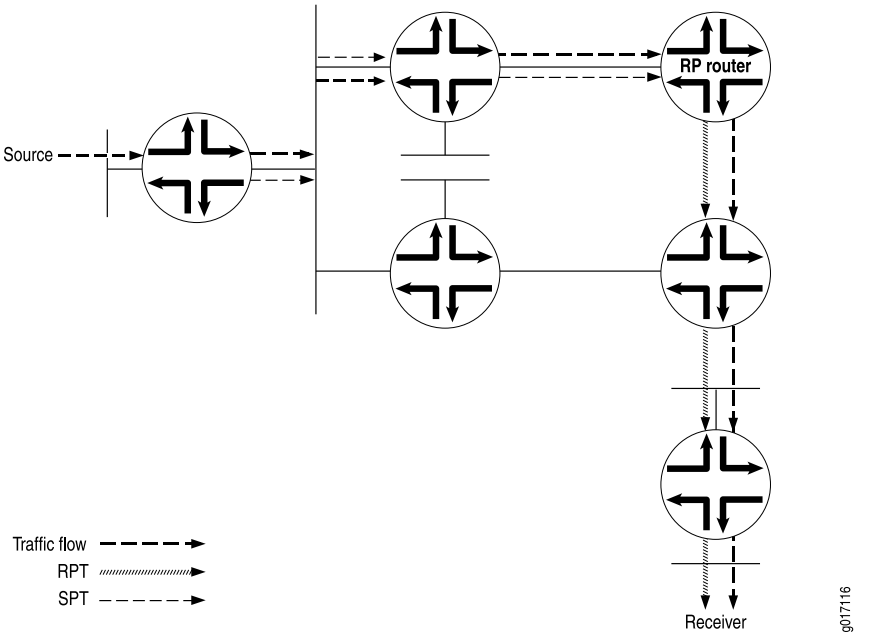

Copyright © 2016, Juniper Networks, Inc.

Multicast Shortest-Path Tree

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.

SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point. For more information about RPs, see “Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees” on page 4167.

Copyright © 2016, Juniper Networks, Inc.

SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:

1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see Figure 171).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

Figure 171: Receiver DR Sends a PIM Join Message to the Source

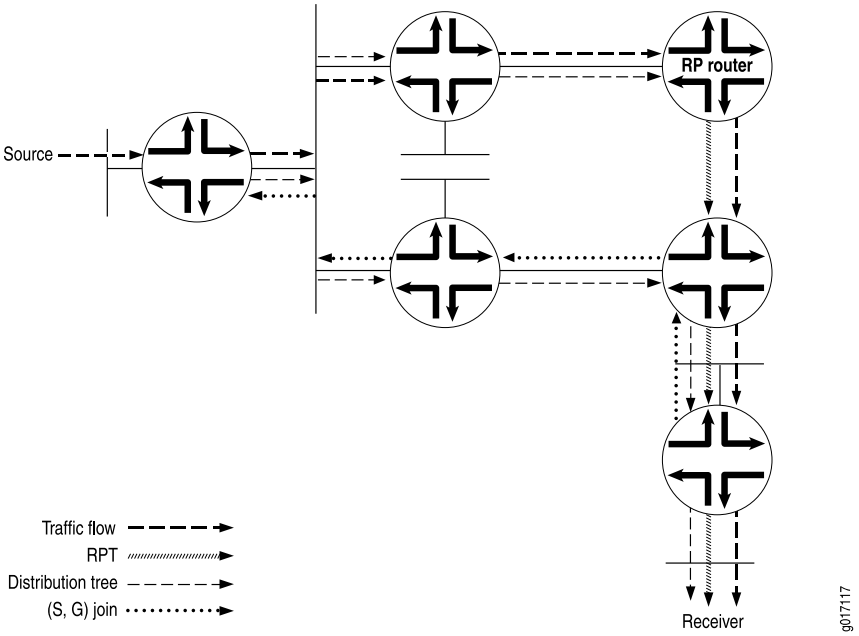

4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see Figure 172).

Copyright © 2016, Juniper Networks, Inc.

Figure 172: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router

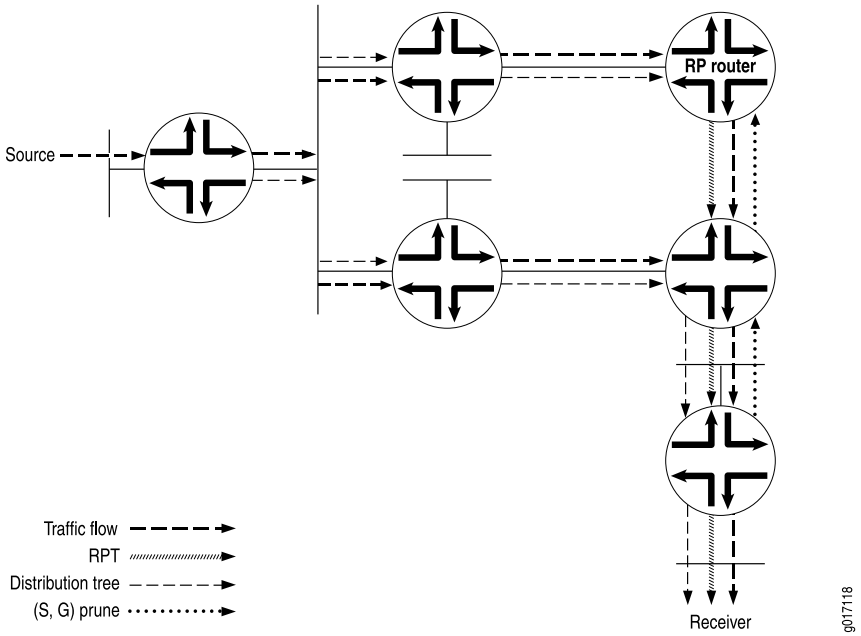

5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see Figure 173).

Figure 173: RP Router Receives PIM Prune Message

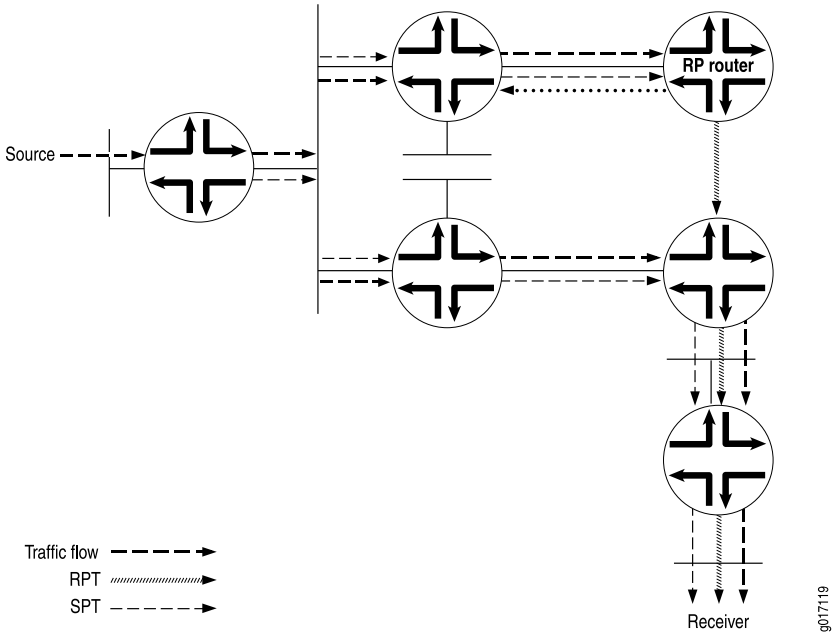

6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see Figure 174).

Copyright © 2016, Juniper Networks, Inc.

Figure 174: RP Router Sends a PIM Prune Message to the Source DR

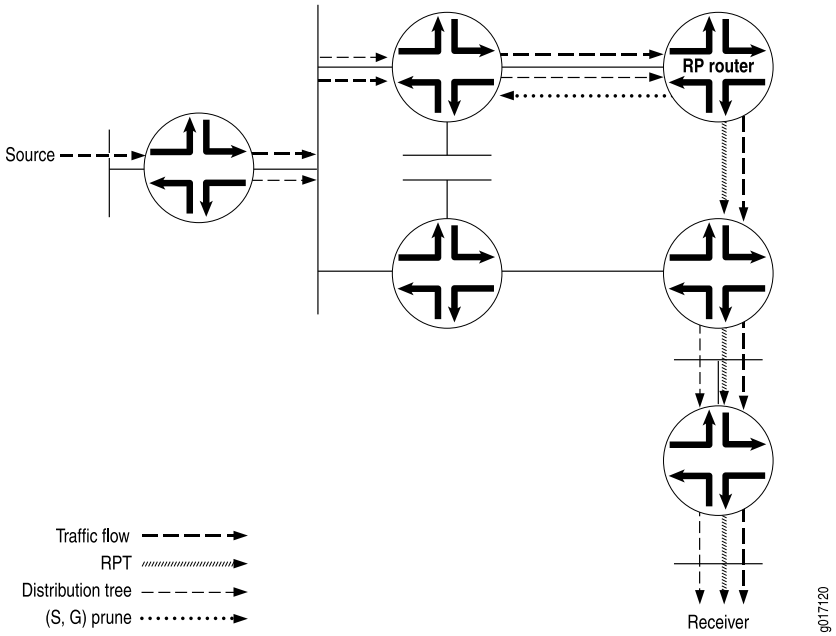

7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see Figure 175).

Figure 175: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router

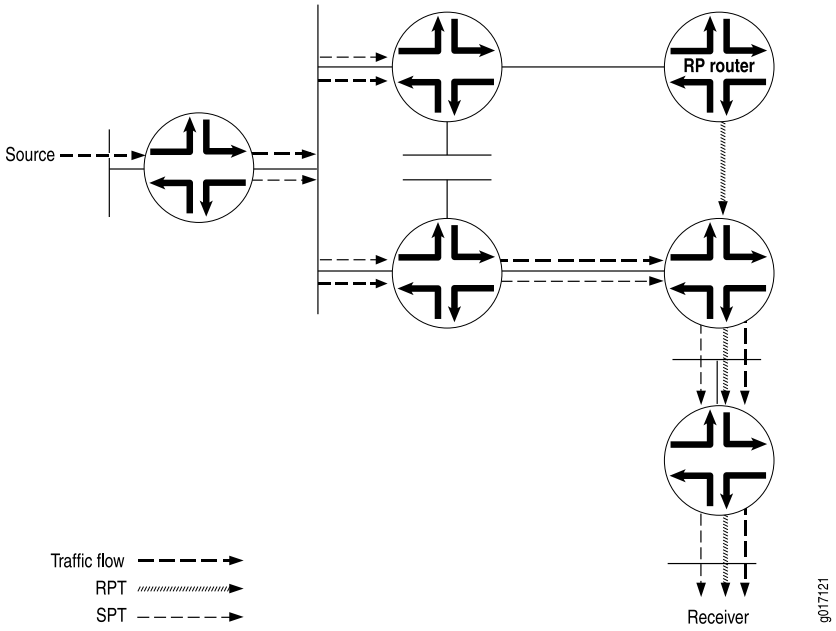

Copyright © 2016, Juniper Networks, Inc.

SPT Cutover Control

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

Example: Configuring the PIM Assert Timeout

This example shows how to configure the timeout period for a PIM assert forwarder.

- Requirements on page 4175
- Overview on page 4175
- Configuration on page 4177

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 4118.

Overview

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In Figure 176, both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the source, and a preference metric for the unicast cost. The preference metric expresses a preference

Copyright © 2016, Juniper Networks, Inc.

between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.

When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

Figure 176 shows the topology for this example.

Figure 176: PIM Assert Topology

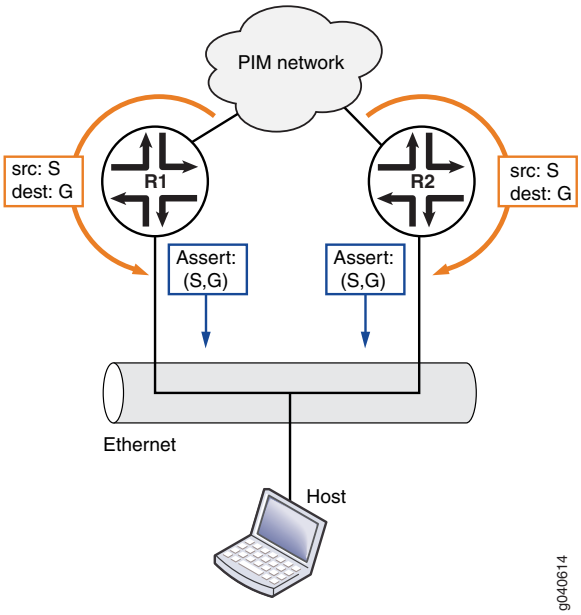

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an assert timeout:

1. Configure the timeout period, in seconds.

   ```
   [edit protocols pim]
   user@host# set assert-timeout 60
   ```

2. (Optional) Trace assert messages.

   ```
   [edit protocols pim]
   user@host# set traceoptions file PIM.log
   user@host# set traceoptions flag assert detail
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   user@host# commit
   ```

4. To verify the configuration, run the following commands:

   - show pim join
   - show pim statistics

Example: Configuring the PIM SPT Threshold Policy

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

- Requirements on page 4177
- Overview on page 4178
- Configuration on page 4179
- Verification on page 4181

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure PIM Sparse Mode on the interfaces. See "Enabling PIM Sparse Mode" on page 4118.

Copyright © 2016, Juniper Networks, Inc.

Overview

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the SPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.
- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

  This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Copyright © 2016, Juniper Networks, Inc.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter
   224.1.1.1/32 exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
   10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
  term one {
    from {
      route-filter 224.1.1.1/32 exact;
      source-address-filter 10.10.10.1/32 exact;
    }
    then accept;
  }
  term two {
    then reject;
  }
}
```

```
user@host# show protocols
pim {
  spt-threshold {
    infinity spt-infinity-policy;
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration, run the **show pim join** command.

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Dense Mode on page 4107
- Configuring a Designated Router for PIM on page 4103
- Configuring PIM Filtering on page 4159
- Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol on page 4219
- Configuring Basic PIM Settings on page 4093

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 182

Receiving Content Directly from the Source with SSM

- Example: Configuring Source-Specific Multicast on page 4183
- Example: Configuring SSM Maps for Different Groups to Different Sources on page 4195

Example: Configuring Source-Specific Multicast

- Understanding PIM Source-Specific Mode on page 4183
- PIM SSM on page 4184
- Source-Specific Multicast Groups Overview on page 4186
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4187
- Example: Configuring an SSM-Only Domain on page 4190
- Example: Configuring PIM SSM on a Network on page 4191
- Example: Configuring SSM Mapping on page 4192

Understanding PIM Source-Specific Mode

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. The RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

Copyright © 2016, Juniper Networks, Inc.

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to *subscribers* (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in Table 386.

Table 386: ASM and SSM Terminology

Term	Any-Source Multicast	Source-Specific Multicast
Address identifier	G	S,G
Address designation	group	channel
Receiver operations	join, leave	subscribe, unsubscribe
Group address range	224/4 excluding 232/8	224/4 (guaranteed only for 232/8)

Although PIM SSM describes receiver operations as *subscribe* and *unsubscribe*, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

PIM SSM

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255

Copyright © 2016, Juniper Networks, Inc.

cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

You can also configure Junos OS to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.

In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see Figure 177). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in Figure 177 that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 177: Receiver Announces Desire to Join Group G and Source S

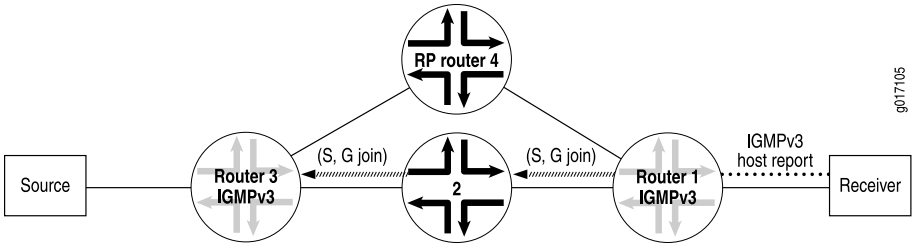

The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In Figure 178, the source tree is built across the network to Router 3, the last-hop router connected to the source.

Figure 178: Router 3 (Last-Hop Router) Joins the Source Tree

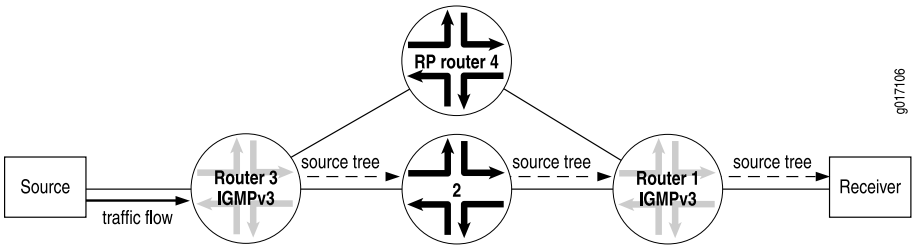

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 179).

Copyright © 2016, Juniper Networks, Inc.

Figure 179: (S,G) State Is Built Between the Source and the Receiver

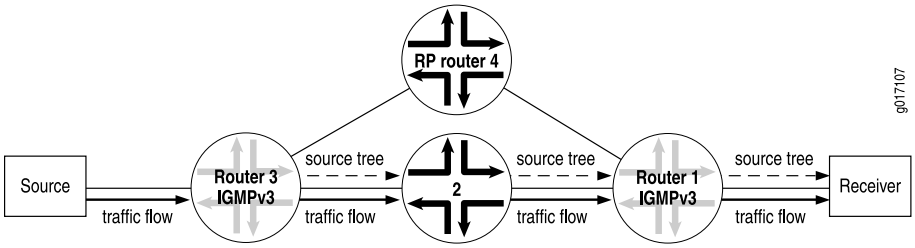

To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

Source-Specific Multicast Groups Overview

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (*,G) pairs. The (*) is a wildcard referring to any source sending to group G, and “G” refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The “S” refers to the source's unicast IP address, and the “G” refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol. IGMPv3 is available for Windows XP, Windows Vista, and most UNIX operating systems.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

- Requirements on page 4187
- Overview on page 4187
- Configuration on page 4188
- Verification on page 4190

Requirements

Before you begin, configure the router interfaces.

Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.

In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see Figure 180). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in Figure 180 that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

Figure 180: Receiver Sends Messages to Join Group G and Source S

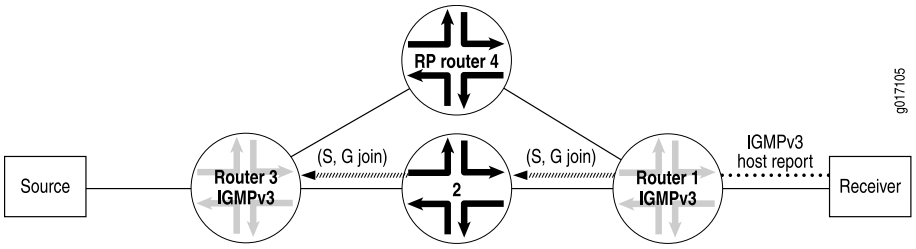

The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In Figure 181, the source tree is built across the network to Router 3, the last-hop router connected to the source.

Copyright © 2016, Juniper Networks, Inc.

Figure 181: Router 3 (Last-Hop Router) Joins the Source Tree

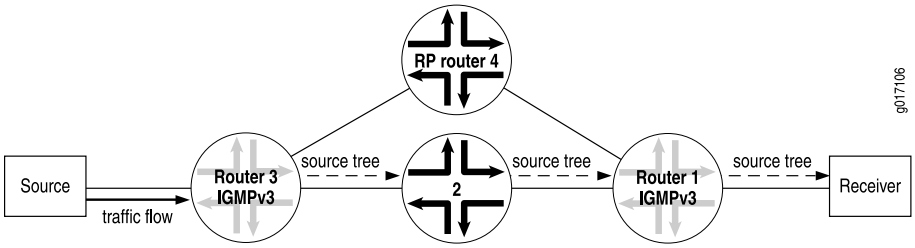

Using the source tree, multicast traffic is delivered to the subscribing host (see Figure 182).

Figure 182: (S,G) State Is Built Between the Source and the Receiver

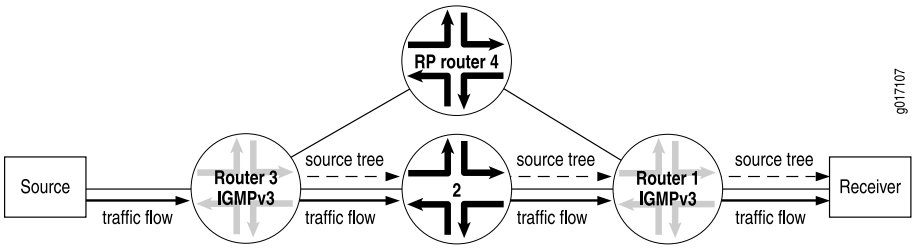

SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in Figure 183.

Figure 183: Simple RPF Topology

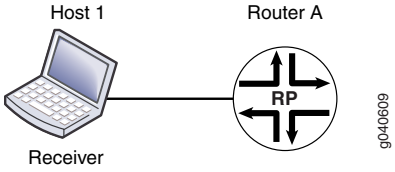

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
pim {
    rp {
        local {
            address 10.255.72.46;
            group-ranges {
                239.0.0.0/24;
            }
        }
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    interface fe-1/0/0.0 {
      mode sparse;
    }
    interface lo0.0 {
      mode sparse;
    }
  }
```

```
user@host# show routing-options
multicast {
  ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
  asm-override-ssm;
}
```

Verification

To verify the configuration, run the following commands:

- show igmp group
- show igmp statistics
- show pim join

Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
  igmp {
    interface fe-0/1/2 {
      version 3;
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in Figure 184.

Figure 184: Network on Which to Configure PIM SSM

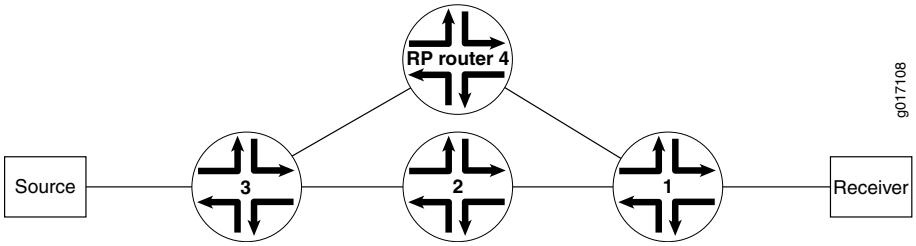

This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

   ```
   user@router1# set protocols igmp interface all version 3
   user@router1# set protocols igmp interface fxp0.0 disable
   ```

   NOTE: When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

   ```
   user@router1> show configuration protocol igmp

   [edit protocols igmp]
   interface all {
     version 3;
   }
   interface fxp0.0 {
     disable;
   }
   ```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

   ```
   user@router1> show igmp interface
   Interface      State    Querier      Timeout Version Groups
   fe-0/0/0.0     Up       198.58.3.245     213       3      0
   fe-0/0/1.0     Up       198.58.3.241     220       3      0
   fe-0/0/2.0     Up       198.58.3.237     218       3      0
   Configured Parameters:
   IGMP Query Interval (1/10 secs): 1250
   IGMP Query Response Interval (1/10 secs): 100
   IGMP Last Member Query Interval (1/10 secs): 10
   IGMP Robustness Count: 2
   Derived Parameters:
   IGMP Membership Timeout (1/10 secs): 2600
   IGMP Other Querier Present Timeout (1/10 secs): 2550
   ```

Copyright © 2016, Juniper Networks, Inc.

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```
user@router2> show pim join extensive
232.1.1.1       10.4.1.2                        sparse
     Upstream interface: fe-1/1/3.0
     Upstream State: Local Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: so-1/0/2.0
            10.10.71.1       State: Join    Flags: S    Timeout: 209
```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```
user@router1> show pim join extensive
232.1.1.1       10.4.1.2                        sparse
     Upstream interface: so-1/0/2.0
     Upstream State: Join to Source
     Keepalive timeout: 209
     Downstream Neighbors:
        Interface: fe-0/2/3.0
            10.3.1.1         State: Join    Flags: S    Timeout: Infinity
```

NOTE: IP version 6 (IPv6) multicast routers use the Multicast Listener Discovery (MLD) Protocol to manage the membership of hosts and routers in multicast groups and to learn which groups have interested listeners for each attached physical networks. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol. MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map

Copyright © 2016, Juniper Networks, Inc.

containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

   ```
   user@router1# set policy-options policy-statement ssm-policy-example term A from
     route-filter 232.1.1.1/32 exact
   user@router1# set policy-options policy-statement ssm-policy-example term A then
     accept
   user@router1# set policy-options policy-statement ssm-policy-example term B from
     route-filter ff35::1/128 exact
   user@router1# set policy-options policy-statement ssm-policy-example term B then
     accept
   ```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

   ```
   user@host> show configuration policy-options

   [edit policy-options]
   policy-statement ssm-policy-example {
     term A {
       from {
         route-filter 232.1.1.1/32 exact;
       }
       then accept;
     }
     term B {
       from {
         route-filter ff35::1/128 exact;
       }
       then accept;
     }
     then reject;
   }
   ```

   The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

   ```
   user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
     ssm-policy-example
   user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
     fec0::1 fec0::12
   user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
     ssm-policy-example
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
  10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
  192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options
```

```
[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.

5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol
```

```
[edit protocols]
igmp {
  interface fe-0/1/0.0 {
    ssm-map ssm-map-ipv4-example;
  }
}
mld {
  interface fe-/0/1/1.0 {
    ssm-map ssm-map-ipv6-example;
  }
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
     Querier: 192.168.224.28
     State:         Up Timeout:    None Version:   2 Groups:   2
     SSM Map: ssm-map-ipv4-example
```

```
user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
     Querier: fec0:0:0:0:1::12
     State:         Up Timeout:    None Version:   2 Groups:   2
     SSM Map: ssm-map-ipv6-example
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring Basic PIM Settings on page 4093

Example: Configuring SSM Maps for Different Groups to Different Sources

- Multiple SSM Maps and Groups for Interfaces on page 4195
- Example: Configuring Multiple SSM Maps Per Interface on page 4195

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

- Requirements on page 4195
- Overview on page 4195
- Configuration on page 4195
- Verification on page 4197

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, POLICY-ipv4-example1, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 1	232.1.1.1	10.10.10.4, 192.168.43.66
POLICY-ipv4-example1 term 2	232.1.1.2	10.10.10.5, 192.168.43.67

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

Configuration

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure this example, perform the following task:

Copyright © 2016, Juniper Networks, Inc.

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter
  232.1.1.1/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter
  232.1.1.2/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  10.10.10.5
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1
```

Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

   ```
   [edit policy-options policy-statement POLICY-ipv4-example1 term 1]
   user@host# set from route-filter 232.1.1.1/32 exact
   user@host# set then ssm-source 10.10.10.4
   user@host# set then ssm-source 192.168.43.66
   user@host# set then accept
   ```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

   ```
   [edit policy-options policy-statement POLICY-ipv4-example1 term 2]
   user@host# set from route-filter 232.1.1.2/32 exact
   user@host# set then ssm-source 10.10.10.5
   user@host# set then ssm-source 192.168.43.67
   user@host# set then accept
   ```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

   ```
   [edit protocols igmp interface fe-0/1/0.0]
   user@host# set ssm-map-policy POLICY-ipv4-example1
   ```

Results

After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host# show policy-options
policy-statement POLICY-ipv4-example1 {
  term 1 {
    from {
```


Copyright © 2016, Juniper Networks, Inc.

```
        route-filter 232.1.1.1/32 exact;
      }
      then {
        ssm-source [ 10.10.10.4 192.168.43.66 ];
        accept;
      }
    }
    term 2{
      from {
        route-filter 232.1.1.2/32 exact;
      }
      then {
        ssm-source [ 10.10.10.5 192.168.43.67 ];
        accept;
      }
    }
  }

  user@host# show protocols
  igmp {
    interface fe-0/1/0.0 {
      ssm-map-policy POLICY-ipv4-example1;
    }
  }
```

Verification

Confirm that the configuration is working properly.

- Displaying Information About IGMP-Enabled Interfaces on page 4197
- Displaying the PIM Groups on page 4198
- Displaying the Entries in the IP Multicast Forwarding Table on page 4198

Displaying Information About IGMP-Enabled Interfaces

Purpose Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

Action Use the show igmp interface operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface
Interface: fe-0/1/0.0
    Querier: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: POLICY-ipv4-example1;

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

Copyright © 2016, Juniper Networks, Inc.

The command output displays the name of the IGMP logical interface (fe-0/1/0.0), which is the address of the routing device that has been elected to send membership queries and group information.

Displaying the PIM Groups

Purpose Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

Action Use the **show pim join extensive 232.1.1.1** operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

Displaying the Entries in the IP Multicast Forwarding Table

Purpose Verify that the IP multicast forwarding table displays the multicast route state.

Action Use the **show multicast route extensive** operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

Related Documentation

- Example: Configuring Source-Specific Multicast on page 4183

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 183

Minimizing Routing State Information with Bidirectional PIM

- Example: Configuring Bidirectional PIM on page 4199

Example: Configuring Bidirectional PIM

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Understanding Bidirectional PIM

Bidirectional PIM (PIM-Bidir) is specified by the IETF in RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*. It provides an alternative to other PIM modes, such as PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM source-specific multicast (SSM). In bidirectional PIM, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes the amount of PIM routing state information that must be maintained, which is especially important in networks with numerous and dispersed senders and receivers. For example, one important application for bidirectional PIM is distributed inventory polling. In many-to-many applications, a multicast query from one station generates multicast responses from many stations. For each multicast group, such an application generates a large number of (S,G) routes for each station in PIM-SM, PIM-DM, or SSM. The problem is even worse in applications that use bursty sources, resulting in frequently changing multicast tables and, therefore, performance problems in routers.

Figure 185 shows the traffic flows generated to deliver traffic for one group to and from three stations in a PIM-SM network.

Copyright © 2016, Juniper Networks, Inc.

Figure 185: Example PIM Sparse-Mode Tree

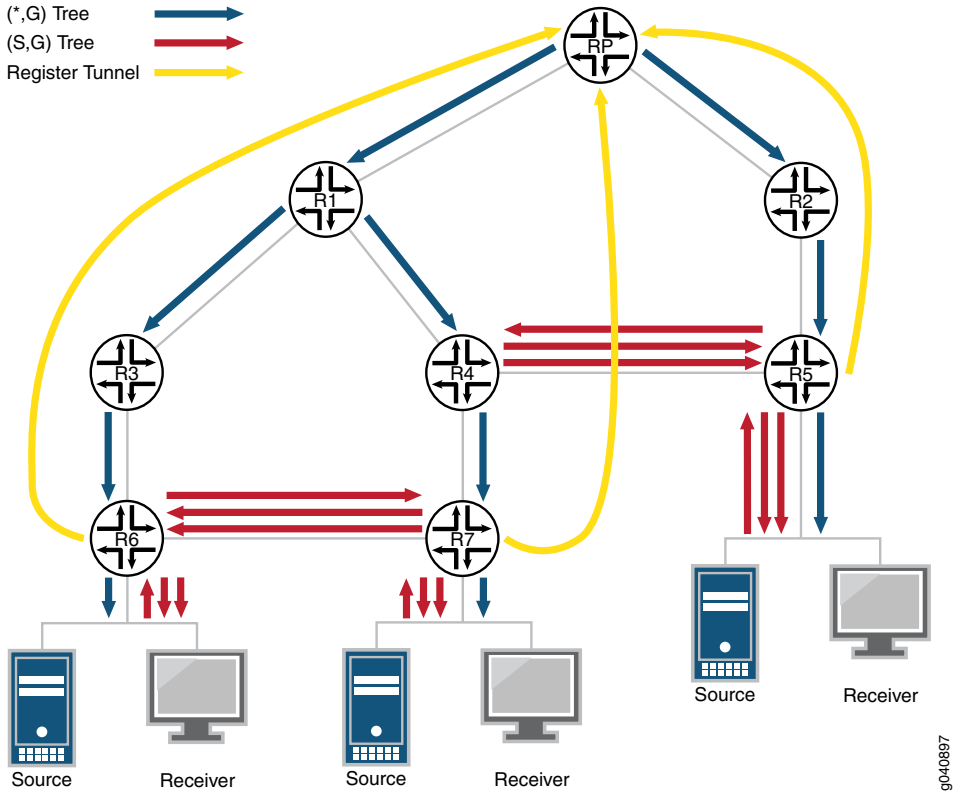

Bidirectional PIM solves this problem by building only group-specific (*,G) state. Thus, only a single (*,G) route is needed for each group to deliver traffic to and from all the sources.

Figure 186 shows the traffic flows generated to deliver traffic for one group to and from three stations in a bidirectional PIM network.

Copyright © 2016, Juniper Networks, Inc.

Figure 186: Example Bidirectional PIM Tree

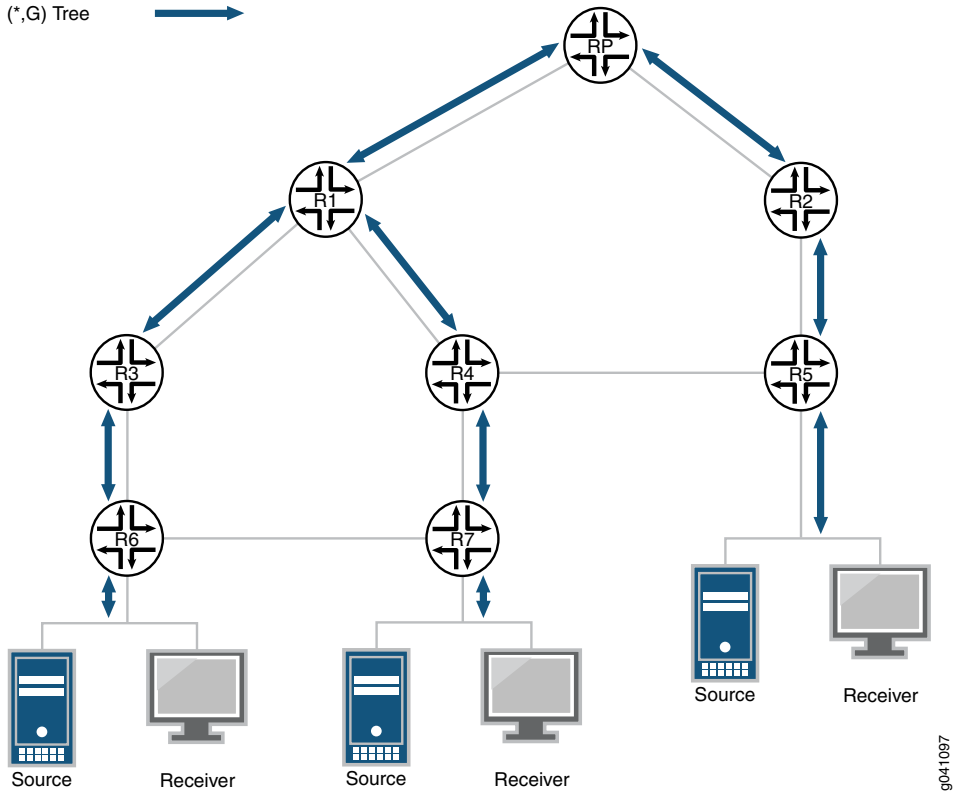

Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees (SPTs) as in PIM-SM and is therefore optimized for routing state size instead of path length. Bidirectional PIM routes are always wildcard-source (*,G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (*,G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM routes forward traffic from all sources and the RP. Thus, bidirectional PIM routers must have the ability to accept traffic on many potential incoming interfaces.

Designated Forwarder Election

To prevent forwarding loops, only one router on each link or subnet (including point-to-point links) is a designated forwarder (DF). The responsibilities of the DF are to forward downstream traffic onto the link toward the receivers and to forward upstream traffic from the link toward the RP address. Bidirectional PIM relies on a process called DF election to choose the DF router for each interface and for each RP address. Each bidirectional PIM router in a subnet advertises its interior gateway protocol (IGP) unicast route to the RP address. The router with the best IGP unicast route to the RP address wins the DF election. Each router advertises its IGP route metrics in DF Offer, Winner, Backoff, and Pass messages.

Copyright © 2016, Juniper Networks, Inc.

Junos OS implements the DF election procedures as stated in RFC 5015, except that Junos OS checks RP unicast reachability before accepting incoming DF messages. DF messages for unreachable rendezvous points are ignored.

Bidirectional PIM Modes

In the Junos OS implementation, there are two modes for bidirectional PIM: bidirectional-sparse and bidirectional-sparse-dense. The differences between bidirectional-sparse and bidirectional-sparse-dense modes are the same as the differences between sparse mode and sparse-dense mode. Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as "dense" is not mapped to an RP. Use bidirectional-sparse-dense mode when you have a mix of bidirectional groups, sparse groups, and dense groups in your network. One typical scenario for this is the use of auto-RP, which uses dense-mode flooding to bootstrap itself for sparse mode or bidirectional mode. In general, the dense groups could be for any flows that the network design requires to be flooded.

Each group-to-RP mapping is controlled by the RP **group-ranges** statement and the **ssm-groups** statement.

The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- **bidirectional-sparse**—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- **bidirectional-sparse-dense**—Use if multicast groups, except those that are specified in the **dense-groups** statement, are operating in bidirectional, sparse, or SSM mode.

Bidirectional Rendezvous Points

You can configure group-range-to-RP mappings network-wide statically, or only on routers connected to the RP addresses and advertise them dynamically. Unlike rendezvous points for PIM-SM, which must de-encapsulate PIM Register messages and perform other specific protocol actions, bidirectional PIM rendezvous points implement no specific functionality. RP addresses are simply locations in the network to rendezvous toward. In fact, RP addresses need not be loopback interface addresses or even be addresses configured on any router, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable router and advertised to the network.

Thus, for bidirectional PIM, there is no meaningful distinction between static and local RP addresses. Therefore, bidirectional PIM rendezvous points are configured at the **[edit protocols pim rp bidirectional]** hierarchy level, not under **static** or **local**.

The settings at the **[edit protocol pim rp bidirectional]** hierarchy level function like the settings at the **[edit protocols pim rp local]** hierarchy level, except that they create bidirectional PIM RP state instead of PIM-SM RP state.

Where only a single local RP can be configured, multiple bidirectional rendezvous points can be configured having group ranges that are the same, different, or overlapping. It is also permissible for a group range or RP address to be configured as bidirectional and either static or local for sparse-mode.

Copyright © 2016, Juniper Networks, Inc.

If a bidirectional PIM RP is configured without a group range, the default group range is 224/4 for IPv4. For IPv6, the default is ff00::/8. You can configure a bidirectional PIM RP group range to cover an SSM group range, but in that case the SSM or DM group range takes precedence over the bidirectional PIM RP configuration for those groups. In other words, because SSM always takes precedence, it is not permitted to have a bidirectional group range equal to or more specific than an SSM or DM group range.

PIM Bootstrap and Auto-RP Support

Group ranges for the specified RP address are flagged by PIM as bidirectional PIM group-to-RP mappings and, if configured, are advertised using PIM bootstrap or auto-RP. Dynamic advertisement of bidirectional PIM-flagged group-to-RP mappings using PIM bootstrap, and auto-RP is controlled as normal using the **bootstrap** and **auto-rp** statements.

Bidirectional PIM RP addresses configured at the **[edit protocols pim rp bidirectional address]** hierarchy level are advertised by auto-RP or PIM bootstrap if the following prerequisites are met:

- The routing instance must be configured to advertise candidate rendezvous points by way of auto-RP or PIM bootstrap, and an auto-RP mapping agent or bootstrap router, respectively, must be elected.
- The RP address must either be configured locally on an interface in the routing instance, or the RP address must belong to a subnet connected to an interface in the routing instance.

IGMP and MLD Support

Internet Group Management Protocol (IGMP) version 1, version 2, and version 3 are supported with bidirectional PIM. Multicast Listener Discovery (MLD) version 1 and version 2 are supported with bidirectional PIM. However, in all cases, only anysource multicast (ASM) state is supported for bidirectional PIM membership.

The following rules apply to bidirectional PIM:

- IGMP and MLD (*,G) membership reports trigger the PIM DF to originate bidirectional PIM (*,G) join messages.
- IGMP and MLD (S,G) membership reports do not trigger the PIM DF to originate bidirectional PIM (*,G) join messages.

Bidirectional PIM and Graceful Restart

Bidirectional PIM accepts packets for a bidirectional route on multiple interfaces. This means that some topologies might develop multicast routing loops if all PIM neighbors are not synchronized with regard to the identity of the designated forwarder (DF) on each link. If one router is forwarding without actively participating in DF elections, particularly after unicast routing changes, multicast routing loops might occur.

If graceful restart for PIM is enabled and bidirectional PIM is enabled, the default graceful restart behavior is to continue forwarding packets on bidirectional routes. If the gracefully

Copyright © 2016, Juniper Networks, Inc.

restarting router was serving as a DF for some interfaces to rendezvous points, the restarting router sends a DF Winner message with a metric of 0 on each of these RP interfaces. This ensures that a neighbor router does not become the DF due to unicast topology changes that might occur during the graceful restart period. Sending a DF Winner message with a metric of 0 prevents another PIM neighbor from assuming the DF role until after graceful restart completes. When graceful restart completes, the gracefully restarted router sends another DF Winner message with the actual converged unicast metric.

The no-bidirectional-mode statement at the **[edit protocols pim graceful-restart]** hierarchy level overrides the default behavior and disables forwarding for bidirectional PIM routes during graceful restart recovery, both in cases of simple routing protocol process (rpd) restart and graceful Routing Engine switchover. This configuration statement provides a very conservative alternative to the default graceful restart behavior for bidirectional PIM routes. The reason to discontinue forwarding of packets on bidirectional routes is that the continuation of forwarding might lead to short-duration multicast loops in rare double-failure circumstances.

Junos OS Enhancements to Bidirectional PIM

In addition to the functionality specified in RFC 5015, the following functions are included in the Junos OS implementation of bidirectional PIM:

- Source-only branches without PIM join state
- Support for both IPv4 and IPv6 domain and multicast addresses
- Nonstop routing (NSR) for bidirectional PIM routes
- Support for bidirectional PIM in logical systems
- Support for non-forwarding and virtual router instances

Limitations of Bidirectional PIM

The Junos OS implementation of bidirectional PIM does not support the following functionality:

- SNMP for bidirectional PIM.
- Graceful Routing Engine switchover is configurable with bidirectional PIM enabled, but bidirectional routes do not forward packets during the switchover.
- Multicast VPNs (Draft Rosen and NextGen).

The bidirectional PIM protocol does not support the following functionality:

- Embedded RP
- Anycast RP

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Bidirectional PIM

This example shows how to configure bidirectional PIM, as specified in RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*.

- Requirements on page 4205
- Overview on page 4205
- Configuration on page 4207
- Verification on page 4211

Requirements

This example uses the following hardware and software components:

- Eight Juniper Networks routers that can be M120, M320, MX Series, or T Series platforms. To support bidirectional PIM, M Series platforms must have I-chip FPCs. M7i, M10i, M40e, and other older M Series routers do not support bidirectional PIM.
- Junos OS Release 12.1 or later running on all eight routers.

Overview

Compared to PIM sparse mode, bidirectional PIM requires less PIM router state information. Because less state information is required, bidirectional PIM scales well and is useful in deployments with many dispersed sources and receivers.

In this example, two rendezvous points are configured statically. One RP is configured as a phantom RP. A phantom RP is an RP address that is a valid address on a subnet, but is not assigned to a PIM router interface. The subnet must be reachable by the bidirectional PIM routers in the network. For the other (non-phantom) RP in this example, the RP address is assigned to a PIM router interface. It can be assigned to either the loopback interface or any physical interface on the router. In this example, it is assigned to a physical interface.

OSPF is used as the interior gateway protocol (IGP) in this example. The OSPF metric determines the designated forwarder (DF) election process. In bidirectional PIM, the DF establishes a loop-free shortest-path tree that is rooted at the RP. On every network segment and point-to-point link, all PIM routers participate in DF election. The procedure selects one router as the DF for every RP of bidirectional groups. This router forwards multicast packets received on that network upstream to the RP. The DF election uses the same tie-break rules used by PIM assert processes.

This example uses the default DF election parameters. Optionally, at the **[edit protocols pim interface (*interface-name* | all)** bidirectional**]** hierarchy level, you can configure the following parameters related to the DF election:

- The robustness-count is the minimum number of DF election messages that must be lost for election to fail.
- The offer period is the interval to wait between repeated DF Offer and Winner messages.

Copyright © 2016, Juniper Networks, Inc.

- The backoff period is the period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

This example uses bidirectional-sparse-dense mode on the interfaces. The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- **bidirectional-sparse**—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- **bidirectional-sparse-dense**—Use if multicast groups, except those that are specified in the **dense-groups** statement, are operating in bidirectional, sparse, or SSM mode.

Topology Diagram

Figure 187 shows the topology used in this example.

Figure 187: Bidirectional PIM with Statically Configured Rendezvous Points

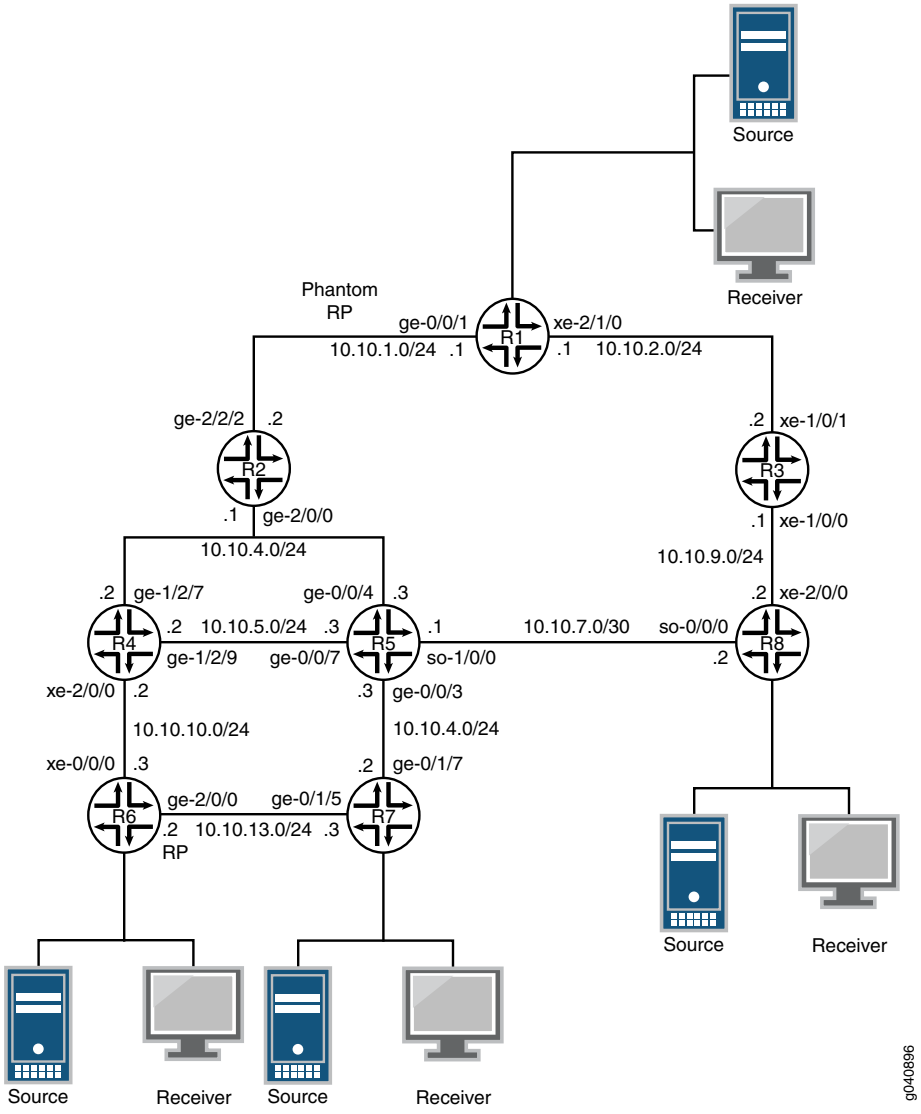

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Router R1

```
set interfaces ge-0/0/1 unit 0 family inet address 10.10.1.1/24
set interfaces xe-2/1/0 unit 0 family inet address 10.10.2.1/24
set interfaces lo0 unit 0 family inet address 10.255.11.11/32
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface xe-2/1/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim interface ge-0/0/1.0 mode bidirectional-sparse-dense
set protocols pim interface xe-2/1/0.0 mode bidirectional-sparse-dense
```

Router R2

```
set interfaces ge-2/0/0 unit 0 family inet address 10.10.4.1/24
set interfaces ge-2/2/2 unit 0 family inet address 10.10.1.2/24
set interfaces lo0 unit 0 family inet address 10.255.22.22/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-2/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim interface fxp0.0 disable
set protocols pim interface ge-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-2/2/2.0 mode bidirectional-sparse-dense
```

Router R3

```
set interfaces xe-1/0/0 unit 0 family inet address 10.10.9.1/24
set interfaces xe-1/0/1 unit 0 family inet address 10.10.2.2/24
set interfaces lo0 unit 0 family inet address 10.255.33.33/32
set protocols ospf area 0.0.0.0 interface xe-1/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface xe-1/0/0.0
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim interface xe-1/0/1.0 mode bidirectional-sparse-dense
set protocols pim interface xe-1/0/0.0 mode bidirectional-sparse-dense
```

Copyright © 2016, Juniper Networks, Inc.

Router R4
```
set interfaces ge-1/2/7 unit 0 family inet address 10.10.4.2/24
set interfaces ge-1/2/8 unit 0 family inet address 10.10.5.2/24
set interfaces xe-2/0/0 unit 0 family inet address 10.10.10.2/24
set interfaces lo0 unit 0 family inet address 10.255.44.44/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/7.0
set protocols ospf area 0.0.0.0 interface ge-1/2/8.0
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface xe-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-1/2/7.0 mode bidirectional-sparse-dense
set protocols pim interface ge-1/2/8.0 mode bidirectional-sparse-dense
```

Router R5
```
set interfaces ge-0/0/3 unit 0 family inet address 10.10.12.3/24
set interfaces ge-0/0/4 unit 0 family inet address 10.10.4.3/24
set interfaces ge-0/0/7 unit 0 family inet address 10.10.5.3/24
set interfaces so-1/0/0 unit 0 family inet address 10.10.7.1/30
set interfaces lo0 unit 0 family inet address 10.255.55.55/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/0/7.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface so-1/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface ge-0/0/7.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/0/4.0 mode bidirectional-sparse-dense
set protocols pim interface so-1/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/0/3.0 mode bidirectional-sparse-dense
```

Router R6
```
set interfaces xe-0/0/0 unit 0 family inet address 10.10.10.3/24
set interfaces ge-2/0/0 unit 0 family inet address 10.10.13.2/24
set interfaces lo0 unit 0 family inet address 10.255.66.66/32
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols ospf area 0.0.0.0 interface xe-0/0/0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface fxp0.0 disable
set protocols pim interface xe-0/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface ge-2/0/0.0 mode bidirectional-sparse-dense
```

Router R7
```
set interfaces ge-0/1/5 unit 0 family inet address 10.10.13.3/24
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/1/7 unit 0 family inet address 10.10.12.2/24
set interfaces lo0 unit 0 family inet address 10.255.77.77/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/1/5.0
set protocols ospf area 0.0.0.0 interface ge-0/1/7.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface ge-0/1/5.0 mode bidirectional-sparse-dense
set protocols pim interface ge-0/1/7.0 mode bidirectional-sparse-dense
```

Router R8

```
set interfaces so-0/0/0 unit 0 family inet address 10.10.7.2/30
set interfaces xe-2/0/0 unit 0 family inet address 10.10.9.2/24
set interfaces lo0 unit 0 family inet address 10.255.88.88/32
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface xe-2/0/0.0
set protocols ospf area 0.0.0.0 interface so-0/0/0.0
set protocols pim traceoptions file df
set protocols pim traceoptions flag bidirectional-df-election detail
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 224.1.1.0/24
set protocols pim rp bidirectional address 10.10.13.2 group-ranges 225.1.1.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 224.1.3.0/24
set protocols pim rp bidirectional address 10.10.1.3 group-ranges 225.1.3.0/24
set protocols pim interface xe-2/0/0.0 mode bidirectional-sparse-dense
set protocols pim interface so-0/0/0.0 mode bidirectional-sparse-dense
```

Router R1

Step-by-Step Procedure

To configure Router R1:

1. Configure the router interfaces.

```
[edit interfaces]
user@R1# set ge-0/0/1 unit 0 family inet address 10.10.1.1/24
user@R1# set xe-2/1/0 unit 0 family inet address 10.10.2.1/24
user@R1# set lo0 unit 0 family inet address 10.255.11.11/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface ge-0/0/1.0
user@R1# set interface xe-2/1/0.0
user@R1# set interface lo0.0
user@R1# set interface fxp0.0 disable
```

3. Configure the group-to-RP mappings.

```
[edit protocols pim rp bidirectional]
user@R1# set address 10.10.1.3 group-ranges 224.1.3.0/24
user@R1# set address 10.10.1.3 group-ranges 225.1.3.0/24
user@R1# set address 10.10.13.2 group-ranges 224.1.1.0/24
user@R1# set address 10.10.13.2 group-ranges 225.1.1.0/24
```

Copyright © 2016, Juniper Networks, Inc.

The RP represented by IP address 10.10.1.3 is a phantom RP. The 10.10.1.3 address is not assigned to any interface on any of the routers in the topology. It is, however, a reachable address. It is in the subnet between Routers R1 and R2.

The RP represented by address 10.10.13.2 is assigned to the **ge-2/0/0** interface on Router R6.

4. Enable bidirectional PIM on the interfaces.

```
[edit protocols pim]
user@R1# set interface ge-0/0/1.0 mode bidirectional-sparse-dense
user@R1# set interface xe-2/1/0.0 mode bidirectional-sparse-dense
```

5. (Optional) Configure tracing operations for the DF election process.

```
[edit protocols pim]
user@R1# set traceoptions file df
user@R1# set traceoptions flag bidirectional-df-election detail
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.10.1.1/24;
        }
    }
}
xe-2/1/0 {
    unit 0 {
        family inet {
            address 10.10.2.1/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.11.11/32;
        }
    }
}
```

```
user@R1# show protocols
ospf {
    area 0.0.0.0 {
        interface ge-0/0/1.0;
        interface xe-2/1/0.0;
        interface lo0.0;
        interface fxp0.0 {
            disable;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  pim {
    rp {
      bidirectional {
        address 10.10.1.3 { # phantom RP
          group-ranges {
            224.1.3.0/24;
            225.1.3.0/24;
          }
        }
        address 10.10.13.2 {
          group-ranges {
            224.1.1.0/24;
            225.1.1.0/24;
          }
        }
      }
    }
    interface ge-0/0/1.0 {
      mode bidirectional-sparse-dense;
    }
    interface xe-2/1/0.0 {
      mode bidirectional-sparse-dense;
    }
    traceoptions {
      file df;
      flag bidirectional-df-election detail;
    }
  }
```

If you are done configuring the router, enter **commit** from configuration mode.

Repeat the procedure for every Juniper Networks router in the bidirectional PIM network, using the appropriate interface names and addresses for each router.

Verification

Confirm that the configuration is working properly.

- Verifying Rendezvous Points on page 4212
- Verifying Messages on page 4212
- Checking the PIM Join State on page 4212
- Displaying the Designated Forwarder on page 4214
- Displaying the PIM Interfaces on page 4214
- Checking the PIM Neighbors on page 4214
- Checking the Route to the Rendezvous Points on page 4215
- Verifying Multicast Routes on page 4215
- Viewing Multicast Next Hops on page 4217

Copyright © 2016, Juniper Networks, Inc.

Verifying Rendezvous Points

Purpose Verify the group-to-RP mapping information.

Action

```
user@R1> show pim rps
Instance: PIM.master
Address family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
10.10.1.3       static      bidir       150    None      2 224.1.3.0/24
                                                           225.1.3.0/24
10.10.13.2      static      bidir       150    None      2 224.1.1.0/24
                                                           225.1.1.0/24
```

Verifying Messages

Purpose Check the number of DF election messages sent and received, and check bidirectional join and prune error statistics.

Action

```
user@R1> show pim statistics
PIM Message type        Received       Sent  Rx errors
V2 Hello                      16         34          0
...
V2 DF Election                18         38          0
...

Global Statistics

...
Rx Bidir Join/Prune on non-Bidir if           0
Rx Bidir Join/Prune on non-DF if              0
```

Checking the PIM Join State

Purpose Confirm the upstream interface, neighbor, and state information.

Copyright © 2016, Juniper Networks, Inc.

Action

```
user@R1> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-2/1/0.0    (DF Winner)

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-2/1/0.0    (DF Winner)
```

Meaning The output shows a (*,G-range) entry for each active bidirectional RP group range. These entries provide a hierarchy from which the individual (*,G) routes inherit RP-derived state (upstream information and accepting interfaces). These entries also provide the control plane basis for the (*, G-range) forwarding routes that implement the sender-only branches of the tree.

Copyright © 2016, Juniper Networks, Inc.

Displaying the Designated Forwarder

Purpose Display RP address information and confirm the DF elected.

Action

```
user@R1> show pim bidirectional df-election
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
    ge-0/0/1.0    (RPL)      DF: none
    lo0.0         (Win)      DF: 10.255.179.246
    xe-2/1/0.0    (Win)      DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
    ge-0/0/1.0    (Lose)     DF: 10.10.1.2
    lo0.0         (Win)      DF: 10.255.179.246
    xe-2/1/0.0    (Lose)     DF: 10.10.2.2
```

Displaying the PIM Interfaces

Purpose Verify that the PIM interfaces have bidirectional-sparse-dense (**SDB**) mode assigned.

Action

```
user@R1> show pim interfaces
Instance: PIM.master

Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/0/1.0          Up   SDB   4 2 NotDR,Active      1 0/0            10.10.1.2
lo0.0              Up   SDB   4 2 DR,Active        0 9901/100      10.255.179.246
xe-2/1/0.0          Up   SDB   4 2 NotDR,Active      1 0/0            10.10.2.2
```

Checking the PIM Neighbors

Purpose Check that the router detects that its neighbors are enabled for bidirectional PIM by verifying that the **B** option is displayed.

Copyright © 2016, Juniper Networks, Inc.

Action

```
user@R1> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface           IP V Mode        Option       Uptime Neighbor addr
ge-0/0/1.0           4 2             HPLGBT     00:06:46 10.10.1.2
xe-2/1/0.0           4 2             HPLGBT     00:06:46 10.10.2.2
```

Checking the Route to the Rendezvous Points

Purpose Check the interface route to the rendezvous points.

Action

```
user@R1> show route 10.10.13.2
inet.0: 56 destinations, 56 routes (55 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.13.0/24      *[OSPF/10] 00:04:35, metric 4
                    > to 10.10.1.2 via ge-0/0/1.0

user@R1> show route 10.10.1.3
inet.0: 56 destinations, 56 routes (55 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.1.0/24       *[Direct/0] 00:06:25
                    > via ge-0/0/1.0
```

Verifying Multicast Routes

Purpose Verify the multicast traffic route for each group.

For bidirectional PIM, the **show multicast route extensive** command shows the (*, G*/prefix*) forwarding routes and the list of interfaces that accept bidirectional PIM traffic.

Copyright © 2016, Juniper Networks, Inc.

Action

```
user@R1> show multicast route extensive
Family: INET

Group: 224.0.0.0/4
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Session description: zeroconfaddr
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097157
    Incoming interface list ID: 559
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 224.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097157
    Incoming interface list ID: 579
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097157
    Incoming interface list ID: 556
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 225.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Session description: Unknown
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097157
```


Copyright © 2016, Juniper Networks, Inc.

```
    Incoming interface list ID: 579
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 225.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Session description: Unknown
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097157
    Incoming interface list ID: 556
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
```

Meaning For information about how the incoming and outgoing interface lists are derived, see the forwarding rules in RFC 5015.

Viewing Multicast Next Hops

Purpose Verify that the correct accepting interfaces are shown in the incoming interface list.

Action

```
user@R1> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface
2097157           10         5 ge-0/0/1.0


Family: Incoming interface list
ID          Refcount KRefcount Downstream interface
579                5         2 lo0.0
                               ge-0/0/1.0
556                5         2 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
559                3         1 lo0.0
                               ge-0/0/1.0
                               xe-4/1/0.0
```

Meaning The nexthop IDs for the outgoing and incoming next hops are referenced directly in the **show multicast route extensive** command.

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 184

Rapidly Detecting Communication Failures with PIM and the BFD Protocol

- Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol on page 4219

Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol

- Understanding Bidirectional Forwarding Detection Authentication for PIM on page 4219
- Configuring BFD for PIM on page 4221
- Configuring BFD Authentication for PIM on page 4222

Understanding Bidirectional Forwarding Detection Authentication for PIM

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

NOTE: **Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over PIM. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.**

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- BFD Authentication Algorithms on page 4220
- Security Authentication Keychains on page 4220
- Strict Versus Loose Authentication on page 4221

Copyright © 2016, Juniper Networks, Inc.

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.

NOTE: **Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.**

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

Copyright © 2016, Juniper Networks, Inc.

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

   This example shows the global configuration.

   ```
   [edit protocols pim]
   user@host# edit interface fe-1/0/0.0 bfd-liveness-detection
   ```

2. Configure the minimum transmit interval.

   This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

   Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

   ```
   [edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
   user@host# set minimum-receive-interval 350
   ```

4. (Optional) Configure other BFD settings.

   As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

   ```
   [edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
   user@host# set minimum-interval 350
   ```

5. Configure the threshold for the adaptation of the BFD session detection time.

   When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

   ```
   [edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
   user@host# set detection-time threshold 800
   ```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

   ```
   [edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
   user@host# set multiplier 50
   ```

7. Configure the BFD version.

   ```
   [edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
   user@host# set version 1
   ```

8. Specify that BFD sessions should not adapt to changing network conditions.

   We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

   ```
   [edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
   user@host# set no-adaptation
   ```

9. Verify the configuration by checking the output of the **show bfd session** command.

Configuring BFD Authentication for PIM

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

Copyright © 2016, Juniper Networks, Inc.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- Configuring BFD Authentication Parameters on page 4223
- Viewing Authentication Information for BFD Sessions on page 4224

Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

   ```
   [edit protocols pim]
   user@host# set interface if3-pim bfd-liveness-detection authentication algorithm
     keyed-sha-1
   ```

   NOTE: Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

   The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

   ```
   [edit protocols pim]
   user@host# set interface if3-pim bfd-liveness-detection authentication keychain
     bfd-pim
   ```

   NOTE: The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

   - The matching keychain name as specified in Step 2.
   - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
   - The secret data used to allow access to the session.
   - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret "$ABC123/"
  start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat these steps to configure the other end of the BFD session.

Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if3-pim** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key **1** contains the secret data "**$ABC123/**" and a start time of June 1, 2009, at 9:46:02 AM PST. Key **2** contains the secret data "**$ABC123/**" and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface if3–pim {
  bfd-liveness-detection {
    authentication {
      algorithm keyed-sha-1;
      key-chain bfd-pim;
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$ABC123/";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$ABC123/";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in

Copyright © 2016, Juniper Networks, Inc.

the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

show bfd session detail

```
user@host# show bfd session detail

                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
50.0.0.2                Up        ge-0/1/5.0    0.900     0.300        3
 Client PIM, TX interval 0.300, RX interval 0.300, Authenticate
 Session up time 3d 00:34
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
```

show bfd session extensive

```
user@host# show bfd session extensive
                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
50.0.0.2                Up        ge-0/1/5.0    0.900     0.300        3
 Client PIM, TX interval 0.300, RX interval 0.300, Authenticate
        keychain bfd-pim, algo keyed-sha-1, mode strict
 Session up time 00:04:42
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.300, min slow interval 1.000
 Adaptive async TX interval 0.300, RX interval 0.300
 Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
 Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
 Local discriminator 2, remote discriminator 2
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict
```

Related Documentation

- Configuring PIM Auto-RP on page 4151
- Configuring PIM Bootstrap Router on page 4147
- Configuring PIM Dense Mode on page 4107
- Configuring a Designated Router for PIM on page 4103
- Configuring PIM Filtering on page 4159
- Example: Configuring Nonstop Active Routing for PIM on page 4227

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 185

Configuring PIM Options

- Example: Configuring Nonstop Active Routing for PIM on page 4227
- Configuring PIM-to-IGMP and PIM-to-MLD Message Translation on page 4240

Example: Configuring Nonstop Active Routing for PIM

- Understanding Nonstop Active Routing for PIM on page 4227
- Example: Configuring Nonstop Active Routing with PIM on page 4228
- Configuring PIM Sparse Mode Graceful Restart on page 4238

Understanding Nonstop Active Routing for PIM

Nonstop active routing configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. When nonstop active routing is configured on a dual Routing Engine platform, the PIM control state is replicated on both Routing Engines.

This PIM state information includes:

- Neighbor relationships
- Join and prune information
- RP-set information
- Synchronization between routes and next hops and the forwarding state between the two Routing Engines

The PIM control state is maintained on the backup Routing Engine by the replication of state information from the master to the backup Routing Engine and having the backup Routing Engine react to route installation and modification in the **[instance].inet.1** routing table on the master Routing Engine. The backup Routing Engine does not send or receive PIM protocol packets directly. In addition, the backup Routing Engine uses the dynamic interfaces created by the master Routing Engine. These dynamic interfaces include PIM encapsulation, de-encapsulation, and multicast tunnel interfaces.

NOTE: **The clear pim join, clear pim register, and clear pim statistics operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.**

Copyright © 2016, Juniper Networks, Inc.

To enable nonstop active routing for PIM (in addition to the PIM configuration on the master Routing Engine), you must include the following statements at the **[edit]** hierarchy level:

- **chassis redundancy graceful-switchover**
- **routing-options nonstop-routing**
- **system commit synchronize**

Example: Configuring Nonstop Active Routing with PIM

This example shows how to configure nonstop active routing for PIM-based multicast IPv4 and IPv6 traffic.

- Requirements on page 4228
- Overview on page 4228
- Configuration on page 4229
- Verification on page 4238

Requirements

Before you begin:

- Configure the router interfaces. See the *Network Interfaces Configuration Guide*.
- Configure an interior gateway protocol or static routing. See the *Routing Protocols Configuration Guide*.
- Configure a multicast group membership protocol (IGMP or MLD). See "Understanding IGMP" on page 4038 and "Understanding MLD" on page 4063.
- For this feature to work with IPv6, the routing device must be running Junos OS Release 10.4 or above.

Overview

Junos OS supports nonstop active routing in the following PIM scenarios:

- Dense mode
- Sparse mode
- SSM
- Static RP
- Auto-RP (for IPv4 only)
- Bootstrap router
- Embedded RP on the non-RP router (for IPv6 only)
- BFD support

Copyright © 2016, Juniper Networks, Inc.

NOTE: Multicast VPNs are not supported with nonstop active routing. Policy-based features (such as neighbor policy, join policy, BSR policy, scope policy, flow maps, and RPF check policy) are not supported with nonstop active routing.

This example uses static RP. The interfaces are configured to receive both IPv4 and IPv6 traffic. R2 provides RP services as the local RP. Note that nonstop active routing is not supported on the RP router. The configuration shown in this example is on R1.

Figure 188 shows the topology used in this example.

Figure 188: Nonstop Active Routing in PIM Domain

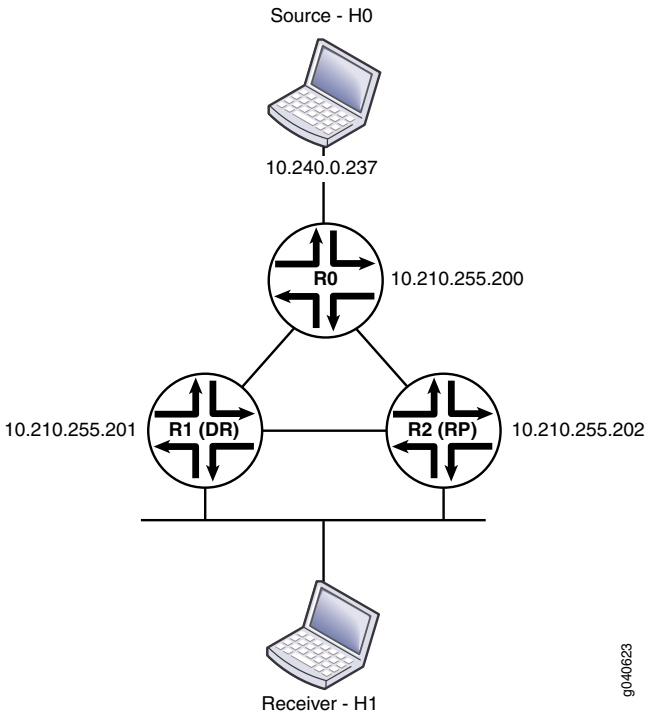

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

R1

```
set system syslog archive size 10m
set system syslog file messages any info
set system commit synchronize
set chassis redundancy graceful-switchover
set interfaces traceoptions file dcd-trace
set interfaces traceoptions file size 10m
set interfaces traceoptions file files 10
set interfaces traceoptions flag all
```

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces so-0/0/1 unit 0 description "to R0 so-0/0/1.0"
set interfaces so-0/0/1 unit 0 family inet address 10.210.1.2/30
set interfaces so-0/0/1 unit 0 family inet6 address FDCA:9E34:50CE:0001::2/126
set interfaces fe-0/1/3 unit 0 description "to R2 fe-0/1/3.0"
set interfaces fe-0/1/3 unit 0 family inet address 10.210.12.1/30
set interfaces fe-0/1/3 unit 0 family inet6 address FDCA:9E34:50CE:0012::1/126
set interfaces fe-1/1/0 unit 0 description "to H1"
set interfaces fe-1/1/0 unit 0 family inet address 10.240.0.250/30
set interfaces fe-1/1/0 unit 0 family inet6 address ::10.240.0.250/126
set interfaces lo0 unit 0 description "R1 Loopback"
set interfaces lo0 unit 0 family inet address 10.210.255.201/32 primary
set interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00
set interfaces lo0 unit 0 family inet6 address abcd::10:210:255:201/128
set protocols ospf traceoptions file r1-nsr-ospf2
set protocols ospf traceoptions file size 10m
set protocols ospf traceoptions file files 10
set protocols ospf traceoptions file world-readable
set protocols ospf traceoptions flag error
set protocols ospf traceoptions flag lsa-update detail
set protocols ospf traceoptions flag flooding detail
set protocols ospf traceoptions flag lsa-request detail
set protocols ospf traceoptions flag state detail
set protocols ospf traceoptions flag event detail
set protocols ospf traceoptions flag hello detail
set protocols ospf traceoptions flag nsr-synchronization detail
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface so-0/0/1.0 metric 100
set protocols ospf area 0.0.0.0 interface fe-0/1/3.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface fe-1/1/0.0 passive
set protocols ospf3 traceoptions file r1-nsr-ospf3
set protocols ospf3 traceoptions file size 10m
set protocols ospf3 traceoptions file world-readable
set protocols ospf3 traceoptions flag lsa-update detail
set protocols ospf3 traceoptions flag flooding detail
set protocols ospf3 traceoptions flag lsa-request detail
set protocols ospf3 traceoptions flag state detail
set protocols ospf3 traceoptions flag event detail
set protocols ospf3 traceoptions flag hello detail
set protocols ospf3 traceoptions flag nsr-synchronization detail
set protocols ospf3 area 0.0.0.0 interface fe-1/1/0.0 passive
set protocols ospf3 area 0.0.0.0 interface fe-1/1/0.0 metric 1
set protocols ospf3 area 0.0.0.0 interface lo0.0 passive
set protocols ospf3 area 0.0.0.0 interface so-0/0/1.0 metric 1
set protocols ospf3 area 0.0.0.0 interface fe-0/1/3.0 metric 1
set protocols pim traceoptions file r1-nsr-pim
set protocols pim traceoptions file size 10m
set protocols pim traceoptions file files 10
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag mdt detail
set protocols pim traceoptions flag rp detail
set protocols pim traceoptions flag register detail
set protocols pim traceoptions flag packets detail
set protocols pim traceoptions flag autorp detail
```


Copyright © 2016, Juniper Networks, Inc.

```
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag hello detail
set protocols pim traceoptions flag assert detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag state detail
set protocols pim traceoptions flag nsr-synchronization
set protocols pim rp static address 10.210.255.202
set protocols pim rp static address abcd::10:210:255:202
set protocols pim interface lo0.0
set protocols pim interface fe-0/1/3.0 mode sparse
set protocols pim interface fe-0/1/3.0 version 2
set protocols pim interface so-0/0/1.0 mode sparse
set protocols pim interface so-0/0/1.0 version 2
set protocols pim interface fe-1/1/0.0 mode sparse
set protocols pim interface fe-1/1/0.0 version 2
set policy-options policy-statement load-balance then load-balance per-packet
set routing-options nonstop-routing
set routing-options router-id 10.210.255.201
set routing-options forwarding-table export load-balance
set routing-options forwarding-table traceoptions file r1-nsr-krt
set routing-options forwarding-table traceoptions file size 10m
set routing-options forwarding-table traceoptions file world-readable
set routing-options forwarding-table traceoptions flag queue
set routing-options forwarding-table traceoptions flag route
set routing-options forwarding-table traceoptions flag routes
set routing-options forwarding-table traceoptions flag synchronous
set routing-options forwarding-table traceoptions flag state
set routing-options forwarding-table traceoptions flag asynchronous
set routing-options forwarding-table traceoptions flag consistency-checking
set routing-options traceoptions file r1-nsr-sync
set routing-options traceoptions file size 10m
set routing-options traceoptions flag nsr-synchronization
set routing-options traceoptions flag commit-synchronize
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure nonstop active routing on R1:

1. Synchronize the Routing Engines.

   ```
   [edit]
   user@host# edit system
   [edit system]
   user@host# set commit synchronize
   user@host# exit
   ```

2. Enable graceful Routing Engine switchover.

   ```
   [edit]
   user@host# set chassis redundancy graceful-switchover
   ```

3. Configure R1's interfaces.

   ```
   [edit]
   user@host# edit interfaces
   ```

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set so-0/0/1 unit 0 description "to R0 so-0/0/1.0"
user@host# set so-0/0/1 unit 0 family inet address 10.210.1.2/30
user@host# set so-0/0/1 unit 0 family inet6 address FDCA:9E34:50CE:0001::2/126
user@host# set fe-0/1/3 unit 0 description "to R2 fe-0/1/3.0"
user@host# set fe-0/1/3 unit 0 family inet address 10.210.12.1/30
user@host# set fe-0/1/3 unit 0 family inet6 address FDCA:9E34:50CE:0012::1/126
user@host# set fe-1/1/0 unit 0 description "to H1"
user@host# set fe-1/1/0 unit 0 family inet address 10.240.0.250/30
user@host# set fe-1/1/0 unit 0 family inet6 address ::10.240.0.250/126
user@host# set lo0 unit 0 description "R1 Loopback"
user@host# set lo0 unit 0 family inet address 10.210.255.201/32 primary
user@host# set lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00
user@host# set lo0 unit 0 family inet6 address abcd::10:210:255:201/128
user@host# exit
```

4. Configure OSPF for IPv4 on R1.

```
[edit]
user@host# edit protocols ospf
[edit protocols ospf]
user@host# set traffic-engineering
user@host# set area 0.0.0.0 interface so-0/0/1.0 metric 100
user@host# set area 0.0.0.0 interface fe-0/1/3.0 metric 100
user@host# set area 0.0.0.0 interface lo0.0 passive
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface fe-1/1/0.0 passive
```

5. Configure OSPF for IPv6 on R1.

```
[edit]
user@host# edit protocols ospf3
[edit protocols ospf3]
user@host# set area 0.0.0.0 interface fe-1/1/0.0 passive
user@host# set area 0.0.0.0 interface fe-1/1/0.0 metric 1
user@host# set area 0.0.0.0 interface lo0.0 passive
user@host# set area 0.0.0.0 interface so-0/0/1.0 metric 1
user@host# set area 0.0.0.0 interface fe-0/1/3.0 metric 1
```

6. Configure PIM on R1. The PIM static address points to the RP router (R2).

```
[edit]
user@host# edit
[edit protocols pim]
user@host# set protocols pim rpstatic address 10.210.255.202
user@host# set protocols pim rp static address abcd::10:210:255:202
user@host# set protocols pim interface lo0.0
user@host# set protocols pim interface fe-0/1/3.0 mode sparse
user@host# set protocols pim interface fe-0/1/3.0 version 2
user@host# set protocols pim interface so-0/0/1.0 mode sparse
user@host# set protocols pim interface so-0/0/1.0 version 2
user@host# set protocols pim interface fe-1/1/0.0 mode sparse
user@host# set protocols pim interface fe-1/1/0.0 version 2
```

7. Configure per-packet load balancing on R1.

```
[edit]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# edit policy-options policy-statement load-balance
[edit policy-options policy-statement load-balance]
user@host# set then load-balance per-packet
```

8. Apply the load-balance policy on R1.

```
[edit]
user@host# set routing-options forwarding-table export load-balance
```

9. Configure nonstop routing on R1.

```
[edit]
user@host# set routing-options nonstop-routing
user@host# set routing-options router-id 10.210.255.201
```

Step-by-Step Procedure

For troubleshooting, configure system log and tracing operations.

1. Enable system log messages.

```
[edit]
user@host# set system syslog archive size 10m
user@host# set system syslog file messages any info
```

2. Trace interface operations.

```
[edit]
user@host# set interfaces traceoptions file dcd-trace
user@host# set interfaces traceoptions file size 10m
user@host# set interfaces traceoptions file files 10
user@host# set interfaces traceoptions flag all
```

3. Trace IGP operations for IPv4.

```
[edit]
user@host# set protocols ospf traceoptions file r1-nsr-ospf2
user@host# set protocols ospf traceoptions file size 10m
user@host# set protocols ospf traceoptions file files 10
user@host# set protocols ospf traceoptions file world-readable
user@host# set protocols ospf traceoptions flag error
user@host# set protocols ospf traceoptions flag lsa-update detail
user@host# set protocols ospf traceoptions flag flooding detail
user@host# set protocols ospf traceoptions flag lsa-request detail
user@host# set protocols ospf traceoptions flag state detail
user@host# set protocols ospf traceoptions flag event detail
user@host# set protocols ospf traceoptions flag hello detail
user@host# set protocols ospf traceoptions flag nsr-synchronization detail
```

4. Trace IGP operations for IPv6.

```
[edit]
user@host# set protocols ospf3 traceoptions file r1-nsr-ospf3
user@host# set protocols ospf3 traceoptions file size 10m
user@host# set protocols ospf3 traceoptions file world-readable
user@host# set protocols ospf3 traceoptions flag lsa-update detail
user@host# set protocols ospf3 traceoptions flag flooding detail
user@host# set protocols ospf3 traceoptions flag lsa-request detail
user@host# set protocols ospf3 traceoptions flag state detail
user@host# set protocols ospf3 traceoptions flag event detail
user@host# set protocols ospf3 traceoptions flag hello detail
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set protocols ospf3 traceoptions flag nsr-synchronization detail
```

5. Trace PIM operations.

```
[edit]
user@host# set protocols pim traceoptions file r1-nsr-pim
user@host# set protocols pim traceoptions file size 10m
user@host# set protocols pim traceoptions file files 10
user@host# set protocols pim traceoptions file world-readable
user@host# set protocols pim traceoptions flag mdt detail
user@host# set protocols pim traceoptions flag rp detail
user@host# set protocols pim traceoptions flag register detail
user@host# set protocols pim traceoptions flag packets detail
user@host# set protocols pim traceoptions flag autorp detail
user@host# set protocols pim traceoptions flag join detail
user@host# set protocols pim traceoptions flag hello detail
user@host# set protocols pim traceoptions flag assert detail
user@host# set protocols pim traceoptions flag normal detail
user@host# set protocols pim traceoptions flag state detail
user@host# set protocols pim traceoptions flag nsr-synchronization
```

6. Trace all routing protocol functionality.

```
[edit]
user@host# set routing-options traceoptions file r1-nsr-sync
user@host# set routing-options traceoptions file size 10m
user@host# set routing-options traceoptions flag nsr-synchronization
user@host# set routing-options traceoptions flag commit-synchronize
```

7. Trace forwarding table operations.

```
[edit]
user@host# set routing-options forwarding-table traceoptions file r1-nsr-krt
user@host# set routing-options forwarding-table traceoptions file size 10m
user@host# set routing-options forwarding-table traceoptions file world-readable
user@host# set routing-options forwarding-table traceoptions flag queue
user@host# set routing-options forwarding-table traceoptions flag route
user@host# set routing-options forwarding-table traceoptions flag routes
user@host# set routing-options forwarding-table traceoptions flag synchronous
user@host# set routing-options forwarding-table traceoptions flag state
user@host# set routing-options forwarding-table traceoptions flag asynchronous
user@host# set routing-options forwarding-table traceoptions flag
  consistency-checking
```

8. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, and **show system** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show chassis
redundancy {
```


Copyright © 2016, Juniper Networks, Inc.

```
    graceful-switchover;
}
user@host# show interfaces
traceoptions {
    file dcd-trace size 10m files 10;
    flag all;
}
so-0/0/1 {
    unit 0 {
        description "to R0 so-0/0/1.0";
        family inet {
            address 10.210.1.2/30;
        }
        family inet6 {
            address FDCA:9E34:50CE:0001::2/126;
        }
    }
}
fe-0/1/3 {
    unit 0 {
        description "to R2 fe-0/1/3.0";
        family inet {
            address 10.210.12.1/30;
        }
        family inet6 {
            address FDCA:9E34:50CE:0012::1/126;
        }
    }
}
fe-1/1/0 {
    unit 0 {
        description "to H1";
        family inet {
            address 10.240.0.250/30;
        }
        family inet6 {
            address ::10.240.0.250/126;
        }
    }
}
lo0 {
    unit 0 {
        description "R1 Loopback";
        family inet {
            address 10.210.255.201/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.1025.5201.00;
        }
        family inet6 {
            address abcd::10:210:255:201/128;
        }
    }
```

```
}

user@host# show policy-options
policy-statement load-balance {
  then {
    load-balance per-packet;
  }
}

user@host# show protocols
ospf {
  traceoptions {
    file r1-nsr-ospf2 size 10m files 10 world-readable;
    flag error;
    flag lsa-update detail;
    flag flooding detail;
    flag lsa-request detail;
    flag state detail;
    flag event detail;
    flag hello detail;
    flag nsr-synchronization detail;
  }
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/0/1.0 {
      metric 100;
    }
    interface fe-0/1/3.0 {
      metric 100;
    }
    interface lo0.0 {
      passive;
    }
    interface fxp0.0 {
      disable;
    }
    interface fe-1/1/0.0 {
      passive;
    }
  }
}
ospf3 {
  traceoptions {
    file r1-nsr-ospf3 size 10m world-readable;
    flag lsa-update detail;
    flag flooding detail;
    flag lsa-request detail;
    flag state detail;
    flag event detail;
    flag hello detail;
    flag nsr-synchronization detail;
  }
  area 0.0.0.0 {
    interface fe-1/1/0.0 {
      passive;
      metric 1;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
      interface lo0.0 {
        passive;
      }
      interface so-0/0/1.0 {
        metric 1;
      }
      interface fe-0/1/3.0 {
        metric 1;
      }
    }
}
pim {
    traceoptions {
      file r1-nsr-pim size 10m files 10 world-readable;
      flag mdt detail;
      flag rp detail;
      flag register detail;
      flag packets detail;
      flag autorp detail;
      flag join detail;
      flag hello detail;
      flag assert detail;
      flag normal detail;
      flag state detail;
      flag nsr-synchronization;
    }
    rp {
      static {
        address 10.210.255.202;
        address abcd::10:210:255:202;
      }
    }
    interface lo0.0;
    interface fe-0/1/3.0 {
      mode sparse;
      version 2;
    }
    interface so-0/0/1.0 {
      mode sparse;
      version 2;
    }
    interface fe-1/1/0.0 {
      mode sparse;
      version 2;
    }
}
user@host# show routing-options
traceoptions {
    file r1-nsr-sync size 10m;
    flag nsr-synchronization;
    flag commit-synchronize;
}
nonstop-routing;
router-id 10.210.255.201;
forwarding-table {
```

Copyright © 2016, Juniper Networks, Inc.

```
    traceoptions {
        file r1-nsr-krt size 10m world-readable;
        flag queue;
        flag route;
        flag routes;
        flag synchronous;
        flag state;
        flag asynchronous;
        flag consistency-checking;
    }
    export load-balance;
}

user@host# show system
syslog {
    archive size 10m;
    file messages {
        any info;
    }
}
commit synchronize;
```

Verification

To verify the configuration, run the following commands:

- show pim join extensive
- show pim neighbors inet detail
- show pim neighbors inet6 detail
- show pim rps inet detail
- show pim rps inet6 detail
- show multicast route inet extensive
- show multicast route inet6 extensive
- show route table inet.1 detail
- show route table inet6.1 detail

Configuring PIM Sparse Mode Graceful Restart

You can configure PIM sparse mode to continue to forward existing multicast packet streams during a routing process failure and restart. Only PIM sparse mode can be configured this way. The routing platform does not forward multicast packets for protocols other than PIM during graceful restart, because all other multicast protocols must restart after a routing process failure. If you configure PIM sparse-dense mode, only sparse multicast groups benefit from a graceful restart.

The routing platform does not forward new streams until after the restart is complete. After restart, the routing platform refreshes the forwarding state with any updates that were received from neighbors during the restart period. For example, the routing platform

Copyright © 2016, Juniper Networks, Inc.

relearns the join and prune states of neighbors during the restart, but it does not apply the changes to the forwarding table until after the restart.

When PIM sparse mode is enabled, the routing platform generates a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the Internet draft **draft-ietf-pim-sm-v2-new-10.txt**. When a routing platform receives PIM hello messages containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a routing platform with PIM sparse mode restarts, it creates a new generation identifier and sends it to neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase is complete when the restart duration timer expires.

Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast RPF checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

You can configure graceful restart globally or for a routing instance. This example shows how to configure graceful restart globally.

To configure graceful restart for PIM sparse mode:

1. Enable graceful restart.

   ```
   [edit protocols pim]
   user@host# set graceful-restart
   ```

2. (Optional) Configure the amount of time the routing device waits (in seconds) to complete PIM sparse mode graceful restart. By default, the router allows 60 seconds. The range is from 30 through 300 seconds. After this restart time, the Routing Engine resumes normal multicast operation.

   ```
   [edit protocols pim graceful-restart]
   user@host# set restart-duration 120
   ```

3. Monitor the operation of PIM graceful restart by running the show pim neighbors command. In the command output, look for the **G** flag in the **Option** field. The **G** flag stands for generation identifier. Also run the **show task replication** command to verify the status of GRES and NSR.

Related Documentation

- Configuring Basic PIM Settings on page 4093

Copyright © 2016, Juniper Networks, Inc.

Configuring PIM-to-IGMP and PIM-to-MLD Message Translation

- Understanding PIM-to-IGMP and PIM-to-MLD Message Translation on page 4240
- Configuring PIM-to-IGMP Message Translation on page 4241
- Configuring PIM-to-MLD Message Translation on page 4242

Understanding PIM-to-IGMP and PIM-to-MLD Message Translation

Routing devices can translate Protocol Independent Multicast (PIM) join and prune messages into corresponding Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) report or leave messages. You can use this feature to forward multicast traffic across PIM domains in certain network topologies.

In some network configurations, customers are unable to run PIM between the customer edge-facing PIM domain and the core-facing PIM domain, even though PIM is running in sparse mode within each of these domains. Because PIM is not running between the domains, customers with this configuration cannot use PIM to forward multicast traffic across the domains. Instead, they might want to use IGMP to forward IPv4 multicast traffic, or MLD to forward IPv6 multicast traffic across the domains.

To enable the use of IGMP or MLD to forward multicast traffic across the PIM domains in such topologies, you can configure the rendezvous point (RP) router that resides between the edge domain and core domain to translate PIM join or prune messages received from PIM neighbors on downstream interfaces into corresponding IGMP or MLD report or leave messages. The router then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP router. As a result, this feature is sometimes referred to as *PIM-to-IGMP proxy* or *PIM-to-MLD proxy*.

To configure the RP router to translate PIM join or prune messages into IGMP report or leave messages, include the **pim-to-igmp-proxy** statement at the **[edit routing-options multicast]** hierarchy level. Similarly, to configure the RP router to translate PIM join or prune messages into MLD report or leave messages, include the **pim-to-mld-proxy** statement at the **[edit routing-options multicast]** hierarchy level. As part of the configuration, you must specify the full name of at least one, but not more than two, upstream interfaces on which to enable the PIM-to-IGMP proxy or PIM-to-MLD proxy feature.

The following guidelines apply when you configure PIM-to-IGMP or PIM-to-MLD message translation:

- Make sure that the router connecting the PIM edge domain and the PIM core domain is the static or elected RP router.
- Make sure that the RP router is using the PIM sparse mode (PIM-SM) multicast routing protocol.
- When you configure an upstream interface, use the full logical interface specification (for example, **ge-0/0/1.0**) and not just the physical interface specification (**ge-0/0/1**).
- When you configure two upstream interfaces, the RP router transmits the same IGMP or MLD report messages and multicast traffic on both upstream interfaces. As a result,

Copyright © 2016, Juniper Networks, Inc.

make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

- The router transmits IGMP or MLD report messages on one or both upstream interfaces only for the first PIM join message that it receives among all of the downstream interfaces. Similarly, the router transmits IGMP or MLD leave messages on one or both upstream interfaces only if it receives a PIM prune message for the last downstream interface.
- Upstream interfaces support both local sources and remote sources.
- Multicast traffic received from an upstream interface is accepted as if it came from a host.

Configuring PIM-to-IGMP Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages. To do so, include the **pim-to-igmp-proxy** statement at the **[edit routing-options multicast]** hierarchy level:

```
[edit routing-options multicast]
pim-to-igmp-proxy {
  upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-IGMP message translation, also referred to as *PIM-to-IGMP proxy*, is useful when you want to use IGMP to forward IPv4 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-IGMP message translation:

- Make sure that the routing device connecting the PIM edge domain and that the PIM core domain is the static or elected RP routing device.
- Make sure that the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP router transmits the same IGMP messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding IGMP report or leave messages:

1. Include the **pim-to-igmp-proxy** statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the routing device transmits IGMP report or leave messages.

   The following example configures PIM-to-IGMP message translation on a single upstream interface, **ge-0/1/0.1**.

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-options multicast]
user@host# set pim-to-igmp-proxy upstream-interface ge-0/1/0.1
```

The following example configures PIM-to-IGMP message translation on two upstream interfaces, **ge-0/1/0.1** and **ge-0/1/0.2**. You must include the logical interface names within square brackets ([ ]) when you configure a set of two upstream interfaces.

```
[edit routing-options multicast]
user@host# set pim-to-igmp-proxy upstream-interface [ge-0/1/0.1 ge-0/1/0.2]
```

2. Use the **show multicast pim-to-igmp-proxy** command to display the PIM-to-IGMP proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-igmp-proxy
Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

Configuring PIM-to-MLD Message Translation

You can configure the rendezvous point (RP) routing device to translate PIM join or prune messages into corresponding MLD report or leave messages. To do so, include the **pim-to-mld-proxy** statement at the **[edit routing-options multicast]** hierarchy level:

```
[edit routing-options multicast]
pim-to-mld-proxy {
  upstream-interface [ interface-names ];
}
```

Enabling the routing device to perform PIM-to-MLD message translation, also referred to as *PIM-to-MLD proxy*, is useful when you want to use MLD to forward IPv6 multicast traffic between a PIM sparse mode edge domain and a PIM sparse mode core domain in certain network topologies.

Before you begin configuring PIM-to-MLD message translation:

- Make sure that the routing device connecting the PIM edge domain and that the PIM core domain is the static or elected RP routing device.
- Make sure that the PIM sparse mode (PIM-SM) routing protocol is running on the RP routing device.
- If you plan to configure two upstream interfaces, make sure that reverse-path forwarding (RPF) is running in the PIM-SM core domain. Because the RP routing device transmits the same MLD messages and multicast traffic on both upstream interfaces, you need to run RPF to verify that multicast packets are received on the correct incoming interface and to avoid sending duplicate packets.

To configure the RP routing device to translate PIM join or prune messages into corresponding MLD report or leave messages:

1. Include the **pim-to-mld-proxy** statement, specifying the names of one or two logical interfaces to function as the upstream interfaces on which the router transmits MLD report or leave messages.

Copyright © 2016, Juniper Networks, Inc.

The following example configures PIM-to-MLD message translation on a single upstream interface, **ge-0/5/0.1**.

```
[edit routing-options multicast]
user@host# set pim-to-mld-proxy upstream-interface ge-0/5/0.1
```

The following example configures PIM-to-MLD message translation on two upstream interfaces, **ge-0/5/0.1** and **ge-0/5/0.2**. You must include the logical interface names within square brackets ([ ]) when you configure a set of two upstream interfaces.

```
[edit routing-options multicast]
user@host# set pim-to-mld-proxy upstream-interface [ge-0/5/0.1 ge-0/5/0.2]
```

2. Use the **show multicast pim-to-mld-proxy** command to display the PIM-to-MLD proxy state (enabled or disabled) and the name or names of the configured upstream interfaces.

```
user@host# run show multicast pim-to-mld-proxy
Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

Related Documentation

- Configuring IGMP on page 4037
- Examples: Configuring MLD on page 4063

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 62

Configuring Multicast Routing Protocols

- Improving Multicast Reliability with PGM on page 4247
- Connecting Routing Domains Using MSDP on page 4253
- Handling Session Announcements with SAP on page 4271
- Facilitating Multicast Delivery Across Unicast-Only Networks with AMT on page 4273
- Routing Content to Densely Clustered Receivers with DVMRP on page 4287

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 186

Improving Multicast Reliability with PGM

- Configuring PGM on page 4247

Configuring PGM

- Understanding Pragmatic General Multicast on page 4247
- PGM Architecture and PGM Routers on page 4248
- PGM-Enabled Source on page 4249
- PGM-Enabled Receivers on page 4249
- PGM-Enabled Routers on page 4250
- PGM Configuration Guidelines on page 4251

Understanding Pragmatic General Multicast

Multicast applications often require real-time operation. These applications cannot take advantage of Transmission Control Protocol (TCP) reliability features such as sequencing, retransmission, and flow control through windowing between sender and receiver. The User Datagram Protocol (UDP), the major transport layer alternative to TCP, is not as reliable as it needs to be for multicast traffic. Pragmatic General Multicast (PGM) is a special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and to request replacement information if the receiver application requires it. PGM is Internet Protocol number 113.

Although PGM mainly deals with the operation of multicast source and receiver, PGM-enabled routers (called PGM network elements) play a *router assistance* role in the initial delivery and potential replacement of multicast traffic. PGM routers are not mandatory in PGM, but they can provide the following benefits when placed anywhere between the source and receivers:

- Reduce the load on the multicast source by aggregating duplicate messages to the source. PGM routers are required to perform this function.
- Limit the flooding of repair data (replacement information) to only those downstream receivers that requested the repair data. PGM routers are required to perform this function.

Copyright © 2016, Juniper Networks, Inc.

- Act as designated local repairers (DLRs) by caching the repair data and resending it to receivers that request it later. DLR functions are a PGM option, and PGM routers are not required to perform this role.

PGM adds reliability to multicast traffic streams. It is not a complete multicast protocol like the Distance Vector Routing Multicast Protocol (DVMRP) or Protocol Independent Multicast (PIM). Adding PGM to a router does not enable the router to perform multicast functions. Instead, a PGM router with multicast capabilities and a preconfigured multicast protocol such as PIM can offer more reliable multicast services to PGM sources and receivers. PGM is not an alternative to multicast routing protocols, but an enhancement of the multicast capabilities already present and configured on the router.

PGM Architecture and PGM Routers

PGM is defined in RFC 3208 and forms a reliable transport layer for multicast applications. Almost any multicast application can use PGM. Applications most suitable for PGM include stock market ticker update information, news reports, weather warnings, and other information that must reach multiple listeners in its entirety and in a timely fashion.

The basic PGM architecture consists of a multicast content source, one or more receivers, and zero or more routers between the source and receivers. All end devices must be PGM-enabled, although there can be non-PGM routers between the source and receiver. If all routers are non-PGM routers, then no routers are capable of the PGM router assistance function, and all PGM functions take place directly between the source and receiver.

PGM sources send sequenced content in sessions to receivers, using multicast protocols. Other, non-PGM protocols allow receivers to learn about a particular source, its sessions, and its location. PGM receivers listen to multicast original data (ODATA), detect missing content through the sequence numbers, and send unicast negative acknowledgments (NAKs) back to the source. NAKs are answered by multicast NAK confirmations (NCFs), which suppress any NAKs from receivers on the same subnet that have not yet sent a NAK upstream. The source sends multicast repair data (RDATA) to receivers containing the missing content. PGM routers assist in this process by making sure that the negative acknowledgments follow the same path as the outbound content upstream to the source, and by suppressing duplicate negative acknowledgments and repair information.

PGM sources must maintain a “sliding window” of retransmittable information. There is no concept of group membership in PGM, so receivers never need to communicate with the source unless they request repair data with a negative acknowledgment. However, this means that the PGM source determines the window size for each receiver, in contrast to almost all other protocols, and requires a certain processing power in each receiver. The absence of positive receiver-to-source acknowledgments also means that PGM scales well and cuts down on control message traffic that can easily overwhelm a multicast network.

PGM receivers can start receiving a PGM session from a PGM source at any time and request any missing previous information that the receiving application needs. If the session is not long enough or if the transmit window is too small so that the source does not maintain a long session history, the receiver cannot get all required information.

Copyright © 2016, Juniper Networks, Inc.

PGM-Enabled Source

A PGM-enabled source of multicast content generates sequenced packets of ODATA that are multicast to receivers. Interleaved with the content packets are source path messages (SPMs), which tell PGM routers and receivers about their upstream next-hop PGM device—either another PGM router or the PGM source.

ODATA packets and SPMs are multicast from the source. A PGM router always appends its own IP address to the SPM before it is multicast on the downstream interfaces. The SPMs are sent by the source and upstream PGM routers with the router alert option set in the IP headers so that PGM routers do not have to examine every packet in the session for SPM packets.

The PGM source acknowledges a received NAK by multicasting NAK confirmations (NCFs) downstream to the next PGM device on the path to the receiver. NCFs make sure that PGM routers and receivers do not bombard sources with NAKs. Downstream PGM routers suppress all subsequent NAKs that indicate the same missing information once one NCF is received from the upstream device.

The PGM source also responds to NAKs by multicasting RDATA packets with the same sequence number as the one indicated by the NAK. RDATA packets have the router alert option set in the IP header so that PGM routers can distinguish them from ODATA packets.

PGM sources organize their packets in sessions. PGM sources are not required to retain copies of information older than the current session, although they might. Long sessions are not necessarily kept on the source in their entirety.

PGM sources identify themselves through a global source ID (GSID). This globally unique source identifier is formed from the low-order 48 bits of the Message Digest 5 (MD5) signature of the Domain Name System (DNS) name of the source.

PGM-Enabled Receivers

The PGM architecture requires one or more PGM-enabled receivers of the multicast content generated by a PGM source. PGM receivers accept all types of downstream PGM messages: ODATA, SPMs, NCFs, and RDATA.

Receivers process the ODATA packets as they arrive from the source, constantly checking the 32-bit sequence number in the ODATA PGM header for gaps in the sequence. If the receiver detects missing information, it generates a NAK for that sequence number. The NAK is unicast upstream to the PGM next hop, which is a router or the source, as determined by the last address in the received SPM.

A receiver detects that its NAK was received by the PGM next hop when it receives an NCF in response to its NAK. If several receivers on a subnet are missing the same ODATA packet, receivers getting an NCF for the packet before sending a NAK suppress the NAK. If a receiver does not get an NCF in response to a NAK, the receiving application can send a NAK again or continue, with the certainty that information is missing.

After the NCF, PGM receivers are sent an RDATA packet with the same sequence number indicated in the NAK and a copy of the missing ODATA. NCFs and RDATA can originate

Copyright © 2016, Juniper Networks, Inc.

from the source or a router acting as a designated local repairer (DLR) for a subnet. The receiver now has complete information about what is missing.

PGM receivers can request almost anything from the PGM source. However, because the source determines the window size, there is no guarantee that older information is available.

PGM-Enabled Routers

Multicast-capable routers can implement the PGM router assistance functions, although not all multicast routers must be PGM-enabled routers. Mandatory PGM router assistance functions include aggregating duplicate NAKs sent to the source to reduce the load on the multicast source, generating NCFs in response to NAKs, and flooding RDATA packets to only those downstream receivers that requested it with a NAK. Optionally, a PGM router can be a DLR, caching PGM information and cutting down on network traffic by resending RDATA packets locally.

There can be zero or more PGM-enabled network elements (routers) between the source and receiver. If there are no PGM routers between the source and receiver, then all PGM messages flow directly between the source and receiver, and no router assistance functions are possible. Both PGM and non-PGM routers can be freely mixed on a network because PGM is a transport layer protocol and is not involved with router multicast functions.

PGM routers also receive SPMs from the source or an upstream PGM router and forward them downstream, inserting the router's own downstream IP interface address into the SPM so that receivers always know their upstream PGM next hop.

When a PGM router receives unicast NAKs from a downstream PGM router or receiver, the router unicasts one NAK for each missing sequence number to the next-hop PGM device upstream toward the source. The address of the PGM next-hop device is determined by received SPMs.

The PGM router multicasts NCFs in response to received NAKs on the downstream interfaces that received the NAKs. NCFs are not multicast on interfaces that have not received NAKs.

PGM routers must multicast all ODATA and RDATA packets that they receive from upstream PGM devices. Normal multicast protocols are used to determine downstream interfaces.

If the PGM router is a DLR, it responds to received NAKs with an NCF and with its own RDATA packet. NAKs are not forwarded upstream from a DLR.

Figure 189 shows the overall PGM architecture and the role of PGM-enabled routers.

Copyright © 2016, Juniper Networks, Inc.

Figure 189: PGM Architecture and General Operation

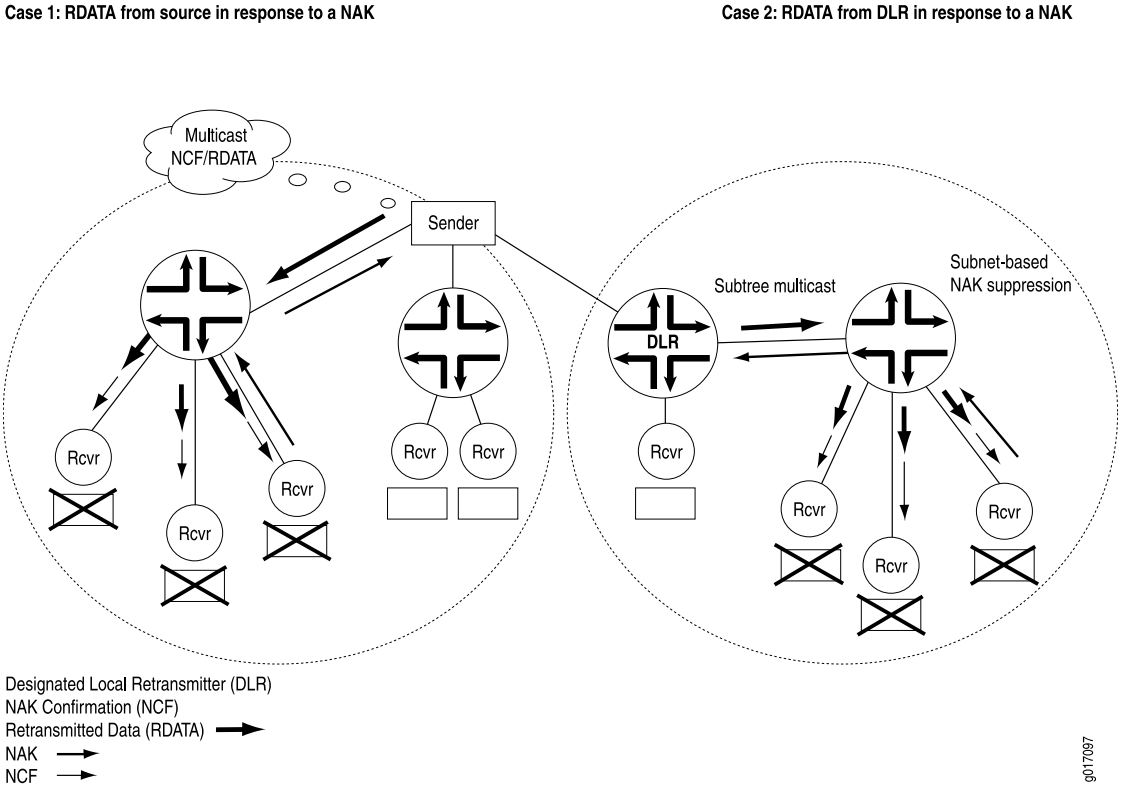

Figure 38 shows only NAKs, NCFs, and RDATA flows. RDATA can come from either the source (left) or a DLR router (right). In both cases, unicast NAKs from a receiver are forwarded upstream by the routers, and multicast NCFs are generated downstream. Subnet NAK suppression is shown, as well as RDATA from the source or DLR sent only to the portions of the network requesting it.

PGM Configuration Guidelines

Pragmatic General Multicast (PGM) allows the router to participate in defined PGM router assistance functions between PGM-enabled sources and receivers. Although PGM is a transport layer protocol and does not do IP packet routing, PGM must be explicitly configured on the router.

To enable PGM globally on the router, include the **pgm** statement:

```
pgm;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems** *logical-system-name* **protocols]**

To trace the operation of PGM, include the **traceoptions** statement:

```
traceoptions {
    flag flag <flag-modifier>;
```

```
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols** pgm**]**
- **[edit logical-systems** *logical-system-name* **protocols** pgm**]**

You can specify the following PGM-specific options in the **flag** statement:

- **all**—Trace all PGM packets.
- **init**—Trace all PGM initialization events.
- **packets**—Trace all PGM packet processing.
- **parser**—Trace all PGM parser processing.
- **route-socket**—Trace all PGM route-socket events.
- **show**—Trace all PGM **show** command servicing.
- **state**—Trace all PGM state transitions.

By default, PGM is enabled on every interface of the router, but global, explicit configuration is required. No options are available for PGM operation.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 187

Connecting Routing Domains Using MSDP

- Examples: Configuring MSDP on page 4253

Examples: Configuring MSDP

- Configuring MSDP on page 4253
- Example: Configuring MSDP in a Routing Instance on page 4254
- Configuring the Interface to Accept Traffic from a Remote Source on page 4262
- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262
- Tracing MSDP Protocol Traffic on page 4268
- Disabling MSDP on page 4269

Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  export [ policy-names ];
  group group-name {
    ... group-configuration ...
  }
  import [ policy-names ];
  local-address address;
  peer address {
    ... peer-configuration ...
  }
  rib-group group-name;
  source ip-prefix</prefix-length> {
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier > <disable>;
}
group group-name {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  mode (mesh-group | standard);
  peer address {
    ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
      just following ...
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
peer address {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  authentication-key peer-key;
  default-peer;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

By default, MSDP is disabled.

Example: Configuring MSDP in a Routing Instance

This example shows how to configure MSDP in a VRF instance.

- Requirements on page 4255
- Overview on page 4255

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 4258
- Verification on page 4261

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable PIM. See "PIM Overview" on page 4089.

Overview

You can configure MSDP in the following types of instances:

- Forwarding
- No forwarding
- Virtual router
- VPLS
- VRF

The main use of MSDP in a routing instance is to support anycast RPs in the network, which allows you to configure redundant RPs. Anycast RP addressing requires MSDP support to synchronize the active sources between RPs.

Copyright © 2016, Juniper Networks, Inc.

This example includes the following MSDP settings.

- **authentication-key**—By default, multicast routers accept and process any properly formatted MSDP messages from the configured peer address. This default behavior might violate the security policies in many organizations because MSDP messages by definition come from another routing domain beyond the control of the security practices of the multicast router's organization.

  The router can authenticate MSDP messages using the TCP message digest 5 (MD5) signature option for MSDP peering sessions. This authentication provides protection against spoofed packets being introduced into an MSDP peering session. Two organizations implementing MSDP authentication must decide on a human-readable key on both peers. This key is included in the MD5 signature computation for each MSDP segment sent between the two peers.

  You configure an MSDP authentication key on a per-peer basis, whether the MSDP peer is defined in a group or individually. If you configure different authentication keys for the same peer one in a group and one individually, the individual key is used.

  The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (“ ”).

  Adding, removing, or changing an MSDP authentication key in a peering session resets the existing MSDP session and establishes a new session between the affected MSDP peers. This immediate session termination prevents excessive retransmissions and eventual session timeouts due to mismatched keys.

- **import** and **export**—All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in, and retrieve from, the routing table.

  You can configure routing policy globally, for a group, or for an individual peer. This example shows how to configure the policy for an individual peer.

  If you configure routing policy at the group level, each peer in a group inherits the group's routing policy.

  The **import** statement applies policies to source-active messages being imported into the source-active cache from MSDP. The **export** statement applies policies to source-active messages being exported from the source-active cache into MSDP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found for the import policy, MSDP shares with the routing table only those routes that were learned from MSDP routers. If no match is found for the export policy, the default MSDP export policy is applied to entries in the source-active cache. See Table 387 for a list of match conditions.

Table 387: MSDP Source-Active Message Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address

Copyright © 2016, Juniper Networks, Inc.

Table 387: MSDP Source-Active Message Filter Match Conditions *(continued)*

Match Condition	Matches On
neighbor	Neighbor address (the source address in the IP header of the source-active message)
route-filter	Multicast group address embedded in the source-active message
source-address-filter	Multicast source address embedded in the source-active message

- **local-address**—Identifies the address of the router you are configuring as an MSDP router (the local router). When you configure MSDP, the **local-address** statement is required. The router must also be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP).
- **peer**—An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. You must configure at least one peer for MSDP to function. When you configure MSDP, the **peer** statement is required. The router must also be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP).

  You can arrange MSDP peers into groups. Each group must contain at least one peer. Arranging peers into groups is useful if you want to block sources from some peers and accept them from others, or set tracing options on one group and not others. This example shows how to configure the MSDP peers in groups. If you configure MSDP peers in a group, each peer in a group inherits all group-level options.

Figure 190 shows the topology for this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 190: MSDP in a VRF Instance Topology

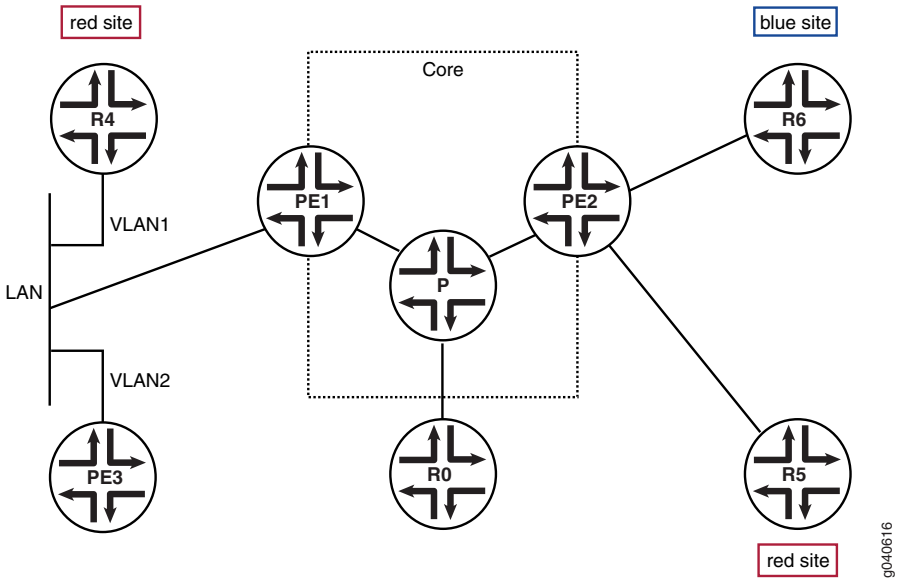

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
set policy-options policy-statement sa-filter term bad-groups from route-filter 224.0.1.2/32
    exact
set policy-options policy-statement sa-filter term bad-groups from route-filter
    224.77.0.0/16 orlonger
set policy-options policy-statement sa-filter term bad-groups then reject
set policy-options policy-statement sa-filter term bad-sources from source-address-filter
    10.0.0.0/8 orlonger
set policy-options policy-statement sa-filter term bad-sources from source-address-filter
    127.0.0.0/8 orlonger
set policy-options policy-statement sa-filter term bad-sources then reject
set policy-options policy-statement sa-filter term accept-everything-else then accept
set routing-instances VPN-100 instance-type vrf
set routing-instances VPN-100 interface ge-0/0/0.100
set routing-instances VPN-100 interface lo0.100
set routing-instances VPN-100 route-distinguisher 10.255.120.36:100
set routing-instances VPN-100 vrf-target target:100:1
set routing-instances VPN-100 protocols ospf export bgp-to-ospf
set routing-instances VPN-100 protocols ospf area 0.0.0.0 interface lo0.100
set routing-instances VPN-100 protocols ospf area 0.0.0.0 interface ge-0/0/0.100
set routing-instances VPN-100 protocols pim rp static address 11.11.47.100
set routing-instances VPN-100 protocols pim interface lo0.100 mode sparse-dense
set routing-instances VPN-100 protocols pim interface lo0.100 version 2
set routing-instances VPN-100 protocols pim interface ge-0/0/0.100 mode sparse-dense
set routing-instances VPN-100 protocols pim interface ge-0/0/0.100 version 2
```


Copyright © 2016, Juniper Networks, Inc.

```
set routing-instances VPN-100 protocols msdp export sa-filter
set routing-instances VPN-100 protocols msdp import sa-filter
set routing-instances VPN-100 protocols msdp group 100 local-address 10.10.47.100
set routing-instances VPN-100 protocols msdp group 100 peer 10.255.120.39
  authentication-key “New York”
set routing-instances VPN-100 protocols msdp group to_pe local-address 10.10.47.100
set routing-instances VPN-100 protocols msdp group to_pe peer 11.11.47.100
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an MSDP routing instance:

1. Configure the BGP export policy.

   ```
   [edit policy-options]
   user@host# set policy-statement bgp-to-ospf term 1 from protocol bgp
   user@host# set policy-statement bgp-to-ospf term 1 then accept
   ```

2. Configure a policy that filters out certain source and group addresses and accepts all other source and group addresses.

   ```
   [edit policy-options]
   user@host# set policy-statement sa-filter term bad-groups from route-filter
     224.0.1.2/32 exact
   user@host# set policy-statement sa-filter term bad-groups from route-filter
     224.0.1.2/32 exact
   user@host# set policy-statement sa-filter term bad-groups from route-filter
     224.77.0.0/16 orlonger
   user@host# set policy-statement sa-filter term bad-groups then reject
   user@host# set policy-statement sa-filter term bad-sources from
     source-address-filter 10.0.0.0/8 orlonger
   user@host# set policy-statement sa-filter term bad-sources from
     source-address-filter 127.0.0.0/8 orlonger
   user@host# set policy-statement sa-filter term bad-sources then reject
   user@host# set policy-statement sa-filter term accept-everything-else then accept
   ```

3. Configure the routing instance type and interfaces.

   ```
   [edit routing-instances]
   user@host# set VPN-100 instance-type vrf
   user@host# set VPN-100 interface ge-0/0/0.100
   user@host# set VPN-100 interface lo0.100
   ```

4. Configure the routing instance route distinguisher and VRF target.

   ```
   [edit routing-instances]
   user@host# set VPN-100 route-distinguisher 10.255.120.36:100
   user@host# set VPN-100 vrf-target target:100:1
   ```

5. Configure OSPF in the routing instance.

   ```
   [edit routing-instances]
   user@host# set VPN-100 protocols ospf export bgp-to-ospf
   user@host# set VPN-100 protocols ospf area 0.0.0.0 interface lo0.100
   user@host# set VPN-100 protocols ospf area 0.0.0.0 interface ge-0/0/0.100
   ```

6. Configure PIM in the routing instance.

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-instances]
user@host# set VPN-100 protocols pim rp static address 11.11.47.100
user@host# set VPN-100 protocols pim interface lo0.100 mode sparse-dense
user@host# set VPN-100 protocols pim interface lo0.100 version 2
user@host# set VPN-100 protocols pim interface ge-0/0/0.100 mode sparse-dense
user@host# set VPN-100 protocols pim interface ge-0/0/0.100 version 2
```

7. Configure MSDP in the routing instance.

```
[edit routing-instances]
user@host# set VPN-100 protocols msdp export sa-filter
user@host# set VPN-100 protocols msdp import sa-filter
user@host# set VPN-100 protocols msdp group 100 local-address 10.10.47.100
user@host# set VPN-100 protocols msdp group 100 peer 10.255.120.39
  authentication-key "New York"
[edit routing-instances]
user@host# set VPN-100 protocols msdp group to_pe local-address 10.10.47.100
[edit routing-instances]
user@host# set VPN-100 protocols msdp group to_pe peer 11.11.47.100
```

8. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** command and the **show routing-instances** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement bgp-to-ospf {
  term 1 {
    from protocol bgp;
    then accept;
  }
}
  policy-statement sa-filter {
    term bad-groups {
      from {
        route-filter 224.0.1.2/32 exact;
        route-filter 224.77.0.0/16 orlonger;
      }
      then reject;
    }
    term bad-sources {
      from {
        source-address-filter 10.0.0.0/8 orlonger;
        source-address-filter 127.0.0.0/8 orlonger;
      }
      then reject;
    }
    term accept-everything-else {
      then accept;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# show routing-instances
VPN-100 {
   instance-type vrf;
   interface ge-0/0/0.100; ## 'ge-0/0/0.100' is not defined
   interface lo0.100; ## 'lo0.100' is not defined
   route-distinguisher 10.255.120.36:100;
   vrf-target target:100:1;
   protocols {
      ospf {
         export bgp-to-ospf;
         area 0.0.0.0 {
            interface lo0.100;
            interface ge-0/0/0.100;
         }
      }
      pim {
         rp {
            static {
               address 11.11.47.100;
            }
         }
         interface lo0.100 {
            mode sparse-dense;
            version 2;
         }
         interface ge-0/0/0.100 {
            mode sparse-dense;
            version 2;
         }
      }
      msdp {
         export sa-filter;
         import sa-filter;
         group 100 {
            local-address 10.10.47.100;
            peer 10.255.120.39 {
               authentication-key "$ABC123"; ## SECRET-DATA
            }
         }
         group to_pe {
            local-address 10.10.47.100;
            peer 11.11.47.100;
         }
      }
   }
}
```

Verification

To verify the configuration, run the following commands:

- **show msdp instance VPN-100**
- **show msdp source-active VPN-100**

Copyright © 2016, Juniper Networks, Inc.

- **show multicast usage instance VPN-100**
- **show route table VPN-100.inet.4**

Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.

In this sample configuration, the incoming interface (**ge-1/3/0**) is on a provider edge (PE) router on the receiver side of a multicast VPN.

To accept traffic from a remote source:

1. Edit the incoming interface.

   ```
   [edit protocols pim interface ge-1/3/0.0]
   user@host# set accept-remote-source
   ```

2. If the incoming interface is not the only way to reach the remote source, ensure that the best path to reach the remote source is through the incoming interface. One way to do this is to use AS path prepending on the other possible routes.

   ```
   [edit policy-options policy-statement as-path-prepend term prepend]
   user@host# set from route-filter 192.168.0.0/16 orlonger
   user@host# set from route-filter 172.16.0.0/16 orlonger
   user@host# set then as-path-prepend "1111"
   ```

   Another way to do this might be to configure a static route on the receiver side PE router to the source.

4. After the configuration is committed, use the **show pim statistics** and **show msdp source** commands to verify that the interface is accepting traffic from the remote source.

Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

- Requirements on page 4262
- Overview on page 4263
- Configuration on page 4266
- Verification on page 4268

Requirements

Before you begin:

Copyright © 2016, Juniper Networks, Inc.

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Enable PIM sparse mode. See "PIM Overview" on page 4089.
- Configure the router as a PIM sparse-mode RP. See "Configuring Local PIM RPs" on page 4136.

Overview

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early detection (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.

NOTE: **The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.**

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Copyright © 2016, Juniper Networks, Inc.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must "pass" all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.

CAUTION: When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.

NOTE: An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the default-peer statement.

Table 388 explains how flooding is handled by peers in this example. .

Copyright © 2016, Juniper Networks, Inc.

Table 388: Source-Active Message Flooding Explanation

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	–

Figure 191 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

Figure 191: Source-Active Message Flooding

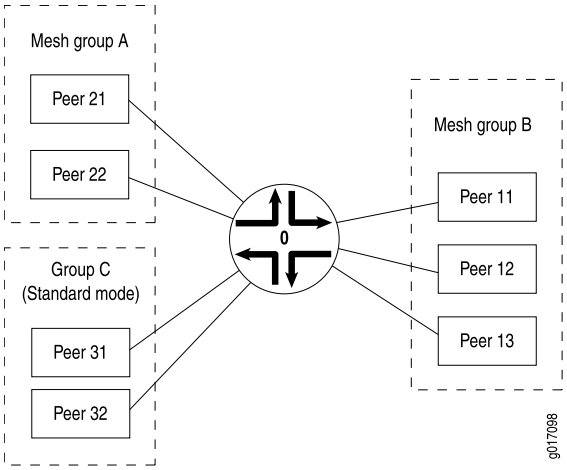

This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

  MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

  By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

Copyright © 2016, Juniper Networks, Inc.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the **MSDP-group** group are mesh group members.
- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshhold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
  7500
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

   ```
   [edit protocols msdp]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
  7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```

4. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
  data-encapsulation disable;
  active-source-limit {
    maximum 10000;
  }
  peer 10.0.0.1 {
    active-source-limit {
      maximum 5000;
      threshold 4000;
    }
  }
  source 10.1.0.0/16 {
    active-source-limit {
      maximum 500;
    }
  }
  group MSDP-group {
    mode mesh-group;
    local-address 10.1.2.3;
    peer 10.10.10.10 {
      active-source-limit {
        maximum 7500;
      }
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration, run the following commands:

- show msdp source-active
- show msdp statistics

Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
general	Trace general events.
keepalive	Trace keepalive messages.
normal	Trace normal events.
packets	Trace all MSDP packets.
policy	Trace policy processing.
route	Trace MSDP changes to the routing table.
source-active	Trace source-active packets.
source-active-request	Trace source-active request packets.
source-active-response	Trace source-active response packets.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

Copyright © 2016, Juniper Networks, Inc.

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the MSDP trace file.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file msdp-trace
   ```

3. (Optional) Configure the maximum number of trace files.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file files 5
   ```

4. (Optional) Configure the maximum size of each trace file.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file size 1m
   ```

5. (Optional) Enable unrestricted file access.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set file world-readable
   ```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

   ```
   [edit protocols msdp group groupa traceoptions]
   user@host# set flag source-active | match 230.0.0.3
   ```

7. View the trace file.

   ```
   user@host> file list /var/log
   user@host> file show /var/log/msdp-trace
   ```

Disabling MSDP

To disable MSDP on the router, include the **disable** statement:

```
disable;
```

You can disable MSDP globally for all peers, for all peers in a group, or for an individual peer.

- Globally for all MSDP peers at the following hierarchy levels:

Copyright © 2016, Juniper Networks, Inc.

- [edit protocols msdp]
- [edit logical-systems *logical-system-name* protocols msdp]
- [edit routing-instances *routing-instance-name* protocols msdp]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp]

- For all peers in a group at the following hierarchy levels:
  - [edit protocols msdp group *group-name*]
  - [edit logical-systems *logical-system-name* protocols msdp group *group-name*]
  - [edit routing-instances *routing-instance-name* protocols msdp group *group-name*]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*]
- For an individual peer at the following hierarchy levels:
  - [edit protocols msdp peer *address*]
  - [edit protocols msdp group *group-name* peer *address*]
  - [edit logical-systems *logical-system-name* protocols msdp peer *address*]
  - [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*]
  - [edit routing-instances *routing-instance-name* protocols msdp peer *address*]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*]

If you disable MSDP at the group level, each peer in the group is disabled.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 188

Handling Session Announcements with SAP

- Configuring the Session Announcement Protocol on page 4271

Configuring the Session Announcement Protocol

The SAP and SDP protocols associate multicast session names with multicast traffic addresses. Only SAP has configuration parameters that users can change. Enabling SAP allows the router to receive announcements about multimedia and other multicast sessions.

Junos OS supports the following SAP and SDP standards:

- RFC 2327, *SDP Session Description Protocol*
- RFC 2974, *Session Announcement Protocol*

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.

To enable SAP and the receipt of session announcements, include the **sap** statement:

```
sap {
  disable;
  listen address <port port>;
}
```

Copyright © 2016, Juniper Networks, Inc.

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

By default, SAP listens to the address and port 224.2.127.254:9875 for session advertisements. To add other addresses or pairs of address and port, include one or more **listen** statements.

Sessions established by SDP, SAP's higher-layer protocol, time out after 60 minutes.

To monitor the operation, use the show sap listen command.

Related Documentation

- show sap listen on page 4834

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 189

Facilitating Multicast Delivery Across Unicast-Only Networks with AMT

- Example: Configuring Automatic IP Multicast Without Explicit Tunnels on page 4273

Example: Configuring Automatic IP Multicast Without Explicit Tunnels

- Understanding AMT on page 4273
- AMT Applications on page 4274
- AMT Operation on page 4276
- Configuring the AMT Protocol on page 4277
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279
- Example: Configuring the AMT Protocol on page 4281

Understanding AMT

Automatic Multicast Tunneling (AMT) facilitates dynamic multicast connectivity between multicast-enabled networks across islands of unicast-only networks. Such connectivity enables service providers, content providers, and their customers to participate in delivering multicast traffic even if they lack end-to-end multicast connectivity.

AMT is supported on MX Series Ethernet Services Routers except the MX80 router and all Modular Port Concentrators (MPCs) that use the Junos Trio chipset. AMT supports graceful restart (GR) but does not support Graceful Routing Engine switchover (GRES).

AMT dynamically establishes unicast-encapsulated tunnels between well-known multicast-enabled relay points (AMT relays) and network points reachable only through unicast (AMT gateways). Figure 192 shows the Automatic Multicast Tunneling Connectivity.

Copyright © 2016, Juniper Networks, Inc.

Figure 192: Automatic Multicast Tunneling Connectivity

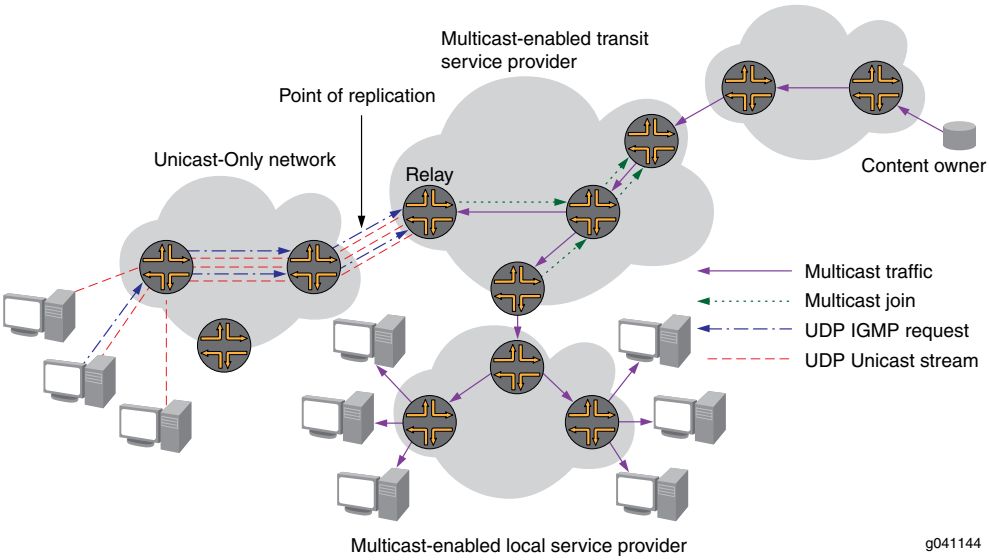

The AMT protocol provides discovery and handshaking between relays and gateways to establish tunnels dynamically without requiring explicit per-tunnel configuration.

AMT relays are typically routers with native IP multicast connectivity that aggregate a potentially large number of AMT tunnels.

The Junos OS implementation supports the following AMT relay functions:

- IPv4 multicast traffic and IPv4 encapsulation
- Well-known sources located on the multicast network
- Prevention of denial-of-service attacks by quickly discarding multicast packets that are sourced through a gateway.
- Per-route replication to the full fan-out of all AMT tunnels desired
- The ability to collect normal interface statistics on AMT tunnels

Multicast sources located behind AMT gateways are not supported."Example: Configuring the AMT Protocol" on page 4281"Example: Configuring the AMT Protocol" on page 4281

AMT supports PIM sparse mode. AMT does not support dense mode operation.

AMT Applications

Transit service providers have a challenge in the Internet because many local service providers are not multicast-enabled. The challenge is how to entice content owners to transmit video and other multicast traffic across their backbones. The cost model for the content owners might be prohibitively high if they have to pay for unicast streams for the majority of their subscribers.

Until more local providers are multicast-enabled, there is a transition strategy proposed by the Internet Engineering Task Force (IETF) and implemented in open source software. This strategy is called Automatic IP Multicast Without Explicit Tunnels (AMT). AMT

Copyright © 2016, Juniper Networks, Inc.

involves setting up relays at peering points in multicast networks that can be reached from gateways installed on hosts connected to unicast networks.

Without AMT, when a user who is connected to a unicast-only network wants to receive multicast content, the content owner can allow the user to join through unicast. However, the content owner incurs an added cost because the owner needs extra bandwidth to support the unicast subscribers.

AMT allows any host to receive multicast. On the client end is an AMT gateway that is a single host. Once the gateway has located an AMT relay, which might be a host but is more typically a router, the gateway periodically sends Internet Group Management Protocol (IGMP) messages over a dynamically created UDP tunnel to the relay. AMT relays and gateways cooperate to transmit multicast traffic sourced within the multicast network to end-user sites. AMT relays receive the traffic natively and unicast-encapsulate it to gateways. This allows anyone on the Internet to create a dynamic tunnel to download multicast data streams.

With AMT, a multicast-enabled service provider can offer multicast services to a content owner. When a customer of the unicast-only local provider wants to receive the content and subscribes using an AMT join, the multicast-enabled transit provider can then efficiently transport the content to the unicast-only local provider, which sends it on to the end user.

AMT is an excellent way for transit service providers (who can get access to the content, but do not have many end users) to provide multicast service to content owners, where it would not otherwise be economically feasible. It is also a useful transition strategy for local service providers who do not yet have multicast support on all downstream equipment.

AMT is also useful for connecting two multicast-enabled service providers that are separated by a unicast-only service provider.

Similarly, AMT can be used by local service providers whose networks are multicast-enabled to tunnel multicast traffic over legacy edge devices such as digital subscriber line access multiplexers (DSLAMs) that have limited multicast capabilities.

Technical details of the implementation of AMT are as follows:

- A three-way handshake is used to join groups from unicast receivers to prevent spoofing and denial-of-service (DoS) attacks.
- An AMT relay acting as a replication server joins the multicast group and translates multicast traffic into multiple unicast streams.
- The discovery mechanism uses anycast, enabling the discovery of the relay that is closest to the gateway in the network topology.
- An AMT gateway acting as a client is a host that joins the multicast group.
- Tunnel count limits on relays can limit bandwidth usage and avoid degradation of service.

Copyright © 2016, Juniper Networks, Inc.

AMT is described in detail in Internet draft draft-ietf-mboned-auto-multicast-10.txt, *Automatic IP Multicast Without Explicit Tunnels (AMT)*.

AMT Operation

AMT is used to create multicast tunnels dynamically between multicast-enabled networks across islands of unicast-only networks. To do this, several steps occur sequentially.

1. The AMT relay (typically a router) advertises an anycast address prefix and route into the unicast routing infrastructure.
2. The AMT gateway (a host) sends AMT relay discovery messages to the nearest AMT relay reachable across the unicast-only infrastructure. To reduce the possibility of replay attacks or dictionary attacks, the relay discovery messages contain a cryptographic nonce. A cryptographic nonce is a random number used only once.
3. The closest relay in the topology receives the AMT relay discovery message and returns the nonce from the discovery message in an AMT relay advertisement message. This enables the gateway to learn the relay's unique IP address. The AMT relay now has an address to use for all subsequent (S,G), entries it will join.
4. The AMT gateway sends an AMT request message to the AMT relay's unique IP address to begin the process of joining the (S,G).
5. The AMT relay sends an AMT membership query back to the gateway.
6. The AMT gateway receives the AMT query message and sends an AMT membership update message containing the IGMP join messages.
7. The AMT relay sends a join message toward the source to build a native multicast tree in the native multicast infrastructure.
8. As packets are received from the source, the AMT relay replicates the packets to all interfaces in the outgoing interface list, including the AMT tunnel. The multicast traffic is then encapsulated in unicast AMT multicast data messages.
9. To maintain state in the AMT relay, the AMT gateway sends periodic AMT membership updates.
10. After the tunnel is established, the AMT tunnel state is refreshed with each membership update message sent. The timeout for the refresh messages is 240 seconds.
11. When the AMT gateway leaves the group, the AMT relay can free resources associated with the tunnel.

Note the following operational details:

- The AMT relay creates an AMT pseudo interface (tunnel interface). AMT tunnel interfaces are implemented as generic UDP encapsulation (**ud**) logical interfaces. These logical interfaces have the identifier format **ud-fpc/pic/port.unit**.
- All multicast packets (data and control) are encapsulated in unicast packets. UDP encapsulation is used for all AMT control and data packets using the IANA reserved UDP port number (2268) for AMT.

Copyright © 2016, Juniper Networks, Inc.

- The AMT relay maintains a receiver list for each multicast session. The relay maintains the multicast state for each gateway that has joined a particular group or (S,G) pair.

Configuring the AMT Protocol

To configure the AMT protocol, include the **amt** statement:

```
amt {
  relay {
    family {
      inet {
        anycast-prefix ip-prefix</prefix-length>;
        local-address ip-address;
      }
    }
    secret-key-timeout minutes;
    tunnel-limit number;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**
- **[edit routing-instances *routing-instance-name* protocols]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]**

NOTE: In the following example, only the **[edit protocols]** hierarchy is identified.

The minimum configuration to enable AMT is to specify the AMT local address and the AMT anycast prefix.

1. To enable the MX Series router to create the UDP encapsulation (**ud**) logical interfaces, include the **bandwidth** statement and specify the bandwidth in gigabits per second.

   ```
   [edit chassis fpc 0 pic 1]
   user@host# set tunnel-services bandwidth 1g
   ```

2. Specify the local address by including the **local-address** statement at the **[edit protocols amt relay family inet]** hierarchy level.

   ```
   [edit protocols amt relay family inet]
   user@host# set local-address 192.168.7.1
   ```

   The local address is used as the IP source of AMT control messages and the source of AMT data tunnel encapsulation. The local address can be configured on any active

Copyright © 2016, Juniper Networks, Inc.

interface. Typically, the IP address of the router's **lo0.0** loopback interface is used for configuring the AMT local address in the default routing instance, and the IP address of the router's **lo0.*n*** loopback interface is used for configuring the AMT local address in VPN routing instances.

3. Specify the AMT anycast address by including the **anycast-prefix** statement at the **[edit protocols amt relay family inet]** hierarchy level.

   ```
   [edit protocols amt relay family inet]
   user@host# set anycast-prefix 192.168.0.0/16
   ```

   The AMT anycast prefix is advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways. Typically, the router's **lo0.0** interface loopback address is used for configuring the AMT anycast prefix in the default routing instance, and the router's **lo0.*n*** loopback address is used for configuring the AMT anycast prefix in VPN routing instances. However, the anycast address can be either the primary or secondary **lo0.0** loopback address.

   Ensure that your unicast routing protocol advertises the AMT anycast prefix in the route advertisements. If the AMT anycast prefix is advertised by BGP, ensure that the local autonomous system (AS) number for the AMT relay router is in the AS path leading to the AMT anycast prefix.

4. (Optional) Specify the AMT secret key timeout by including the **secret-key-timeout** statement at the **[edit protocols amt relay]** hierarchy level. In the following example, the secret key timeout is configured to be 120 minutes.

   ```
   [edit protocols amt relay]
   user@host# set secret-key-timeout 120
   ```

   The secret key is used to generate the AMT Message Authentication Code (MAC). Setting the secret key timeout shorter might improve security, but it consumes more CPU resources. The default is 60 minutes.

5. (Optional) Specify an AMT tunnel limit by including the **tunnel-limit** statement at the **[edit protocols amt relay]** hierarchy level. In the following example, the AMT tunnel limit is 12.

   ```
   [edit protocols amt relay]
   user@host# set tunnel-limit 12
   ```

   The tunnel limit configures the static upper limit to the number of AMT tunnels that can be established. When the limit is reached, new AMT relay discovery messages are ignored.

6. Trace AMT protocol traffic by specifying options to the **traceoptions** statement at the **[edit protocols amt]** hierarchy level. Options applied at the AMT protocol level trace only AMT traffic. In the following example, all AMT packets are logged to the file **amt-log**.

   ```
   [edit protocols amt]
   user@host# set traceoptions file amt-log
   user@host# set traceoptions flag packets
   ```


Copyright © 2016, Juniper Networks, Inc.

NOTE: For AMT operation, configure the PIM rendezvous point address as the primary loopback address of the AMT relay.

Configuring Default IGMP Parameters for AMT Interfaces

You can optionally configure default IGMP parameters for all AMT tunnel interfaces. Although, typically you do not need to change the values. To configure default IGMP attributes of all AMT relay tunnels, include the **amt** statement:

```
amt {
  relay {
    defaults {
      (accounting (Protocols IGMP AMT Interface) | no-accounting (Protocols IGMP AMT
        Interface));
      group-policy [ policy-names ];
      query-interval seconds;
      query-response-interval seconds;
      robust-count number;
      ssm-map ssm-map-name;
      version version;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols igmp]**
- **[edit logical-systems *logical-system-name* protocols igmp]**
- **[edit routing-instances *routing-instance-name* protocols igmp]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp]**

The IGMP statements included at the **[edit protocols igmp amt relay defaults]** hierarchy level have the same syntax and purpose as IGMP statements included at the **[edit protocols igmp]** or **[edit protocols igmp interface *interface-name*]** hierarchy levels. These statements are as follows:

- You can collect IGMP join and leave event statistics. To enable the collection of IGMP join and leave event statistics for all AMT interfaces, include the **accounting** statement:

  user@host# **set protocols igmp amt relay defaults accounting**

- After enabling IGMP accounting, you must configure the router to filter the recorded information to a file or display it to a terminal. You can archive the events file.
- To disable the collection of IGMP join and leave event statistics for all AMT interfaces, include the **no-accounting** statement:

  user@host# **set protocols igmp amt relay defaults no-accounting**

Copyright © 2016, Juniper Networks, Inc.

- You can filter unwanted IGMP reports at the interface level. To filter unwanted IGMP reports, define a policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. Define the policy to match IGMP (S,G) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address. In the following example, the **amt_reject** policy is created to match both the group and source addresses.

  ```
  user@host# set policy-options policy-statement amt_reject from route-filter 224.1.1.1/32
    exact
  user@host# set policy-options policy-statement amt_reject from source-address-filter
    192.168.0.0/16 orlonger
  user@host# set policy-options policy-statement amt_reject then reject
  ```

- To apply the IGMP report filtering on the interface where you prefer not to receive specific group or (S,G) reports, include the **group-policy** statement. The following example applies the **amt_reject** policy to all AMT interfaces.

  ```
  user@host# set protocols igmp amt relay defaults group-policy amt_reject
  ```

- You can change the IGMP query interval for all AMT interfaces to reduce or increase the number of host query messages sent. In AMT, host query messages are sent in response to membership request messages from the gateway. The query interval configured on the relay must be compatible with the membership request timer configured on the gateway. To modify this interval, include the **query-interval** statement. The following example sets the host query interval to 250 seconds.

  ```
  user@host# set protocols igmp amt relay defaults query-interval 250
  ```

  The IGMP querier router periodically sends general host-query messages. These messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

- You can change the IGMP query response interval. The query response interval multiplied by the robust count is the maximum amount of time that can elapse between the sending of a host query message by the querier router and the receipt of a response from a host. Varying this interval allows you to adjust the number of IGMP messages on the AMT interfaces. To modify this interval, include the **query-response-interval** statement. The following example configures the query response interval to 20 seconds.

  ```
  user@host# set protocols igmp amt relay defaults query-response-interval 20
  ```

- You can change the IGMP robust count. The robust count is used to adjust for the expected packet loss on the AMT interfaces. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork. To modify the robust count, include the **robust-count** statement. The following example configures the robust count to **3**.

  ```
  user@host# set protocols igmp amt relay defaults robust-count 3
  ```

  The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3.

  - On a shared network running IGMPv2, when the query router receives an IGMP leave message, it must send an IGMP group query message for a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

Copyright © 2016, Juniper Networks, Inc.

The interval between query messages is determined by the last member query interval. Also, the IGMPv2 query response interval is multiplied by the robust count to determine the maximum amount of time between the sending of a host query message and receipt of a response from a host.

For more information about the IGMPv2 robust count, see RFC 2236, *Internet Group Management Protocol, Version 2*.

- In IGMPv3 a change of interface state causes the system to immediately transmit a state-change report from that interface. If the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3 the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

  For more information about the IGMPv3 robust count, see RFC 3376, *Internet Group Management Protocol, Version 3*.

- You can apply a source-specific multicast (SSM) map to an AMT interface. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report, which allows hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

  SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4).

  In this example, you create a policy to match the 232.1.1.1/32 group address for translation to IGMPv3. Then you define the SSM map that associates the policy with the 192.168.43.66 source address where these group addresses are found. Finally, you apply the SSM map to all AMT interfaces.

  ```
  user@host# set policy-options policy-statement ssm-policy-example term A from
    route-filter 232.1.1.1/32 exact
  user@host# set policy-options policy-statement ssm-policy-example term A then
    accept
  user@host# set routing-options multicast ssm-map ssm-map-example policy
    ssm-policy-example
  user@host# set routing-options multicast ssm-map ssm-map-example source
    192.168.43.66
  user@host# set protocols igmp amt relay defaults ssm-map ssm-map-example
  ```

Example: Configuring the AMT Protocol

This example shows how to configure the Automatic Multicast Tunneling (AMT) Protocol to facilitate dynamic multicast connectivity between multicast-enabled networks across islands of unicast-only networks.

- Requirements on page 4282
- Overview on page 4282
- Configuration on page 4282
- Verification on page 4284

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure a multicast group membership protocol (IGMP or MLD). See "Understanding IGMP" on page 4038 and "Understanding MLD" on page 4063.

Overview

In this example, Host 0 and Host 2 are multicast receivers in a unicast cloud. Their default gateway devices are AMT gateways. R0 and R4 are configured with unicast protocols only. R1, R2, R3, and R5 are configured with PIM multicast. Host 1 is a source in a multicast cloud. R0 and R5 are configured to perform AMT relay. Host 3 and Host 4 are multicast receivers (or sources that are directly connected to receivers). This example shows R1 configured with an AMT relay local address and an anycast prefix as its own loopback address. The example also shows R0 configured with tunnel services enabled.

Figure 193 shows the topology used in this example.

Figure 193: AMT Gateway Topology

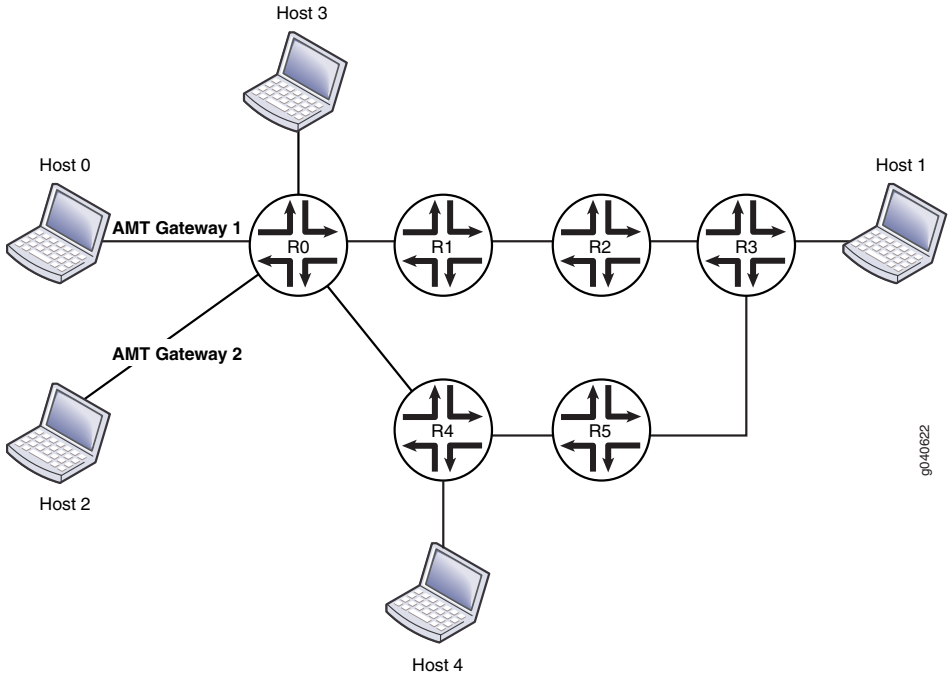

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols amt traceoptions file amt.log
set protocols amt traceoptions flag errors
set protocols amt traceoptions flag packets detail
set protocols amt traceoptions flag route detail
set protocols amt traceoptions flag state detail
set protocols amt traceoptions flag tunnels detail
set protocols amt relay family inet anycast-prefix 10.10.10.10/32
set protocols amt relay family inet local-address 10.255.112.201
set protocols amt relay tunnel-limit 10
set protocols pim interface all mode sparse-dense
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set chassis fpc 0 pic 0 tunnel-services bandwidth 1g
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the AMT protocol on R1:

1. Configure AMT tracing operations.

   ```
   [edit protocols amt traceoptions]
   user@host# set file amt.log
   user@host# set flag errors
   user@host# set flag packets detail
   user@host# set flag route detail
   user@host# set flag state detail
   user@host# set flag tunnels detail
   ```

2. Configure the AMT relay settings.

   ```
   [edit protocols amt relay]
   user@host# set relay family inet anycast-prefix 10.10.10.10/32
   user@host# set family inet local-address 10.255.112.201
   user@host# set tunnel-limit 10
   ```

3. Configure PIM on R1's interfaces.

   ```
   [edit protocols pim]
   set interface all mode sparse-dense
   set interface all version 2
   set interface fxp0.0 disable
   ```

4. Enable tunnel functionality.

   ```
   [edit chassis]
   set fpc 0 pic 0 tunnel-services bandwidth 1g
   ```

5. If you are done configuring the device, commit the configuration.

   ```
   user@host# commit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results

From configuration mode, confirm your configuration by entering the **show chassis** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show chassis
fpc 0 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
```

```
user@host# show protocols
amt {
  traceoptions {
    file amt.log;
    flag errors;
    flag packets detail;
    flag route detail;
    flag state detail;
    flag tunnels detail;
  }
  relay {
    family {
      inet {
        anycast-prefix 10.10.10.10/32;
        local-address 10.255.112.201;
      }
    }
    tunnel-limit 10;
  }
}
pim {
  interface all {
    mode sparse-dense;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
```

Verification

To verify the configuration, run the following commands:

- show amt statistics
- show amt summary
- show amt tunnel

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding AMT on page 4273

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 190

Routing Content to Densely Clustered Receivers with DVMRP

- Examples: Configuring DVMRP on page 4287

Examples: Configuring DVMRP

- Understanding DVMRP on page 4287
- Configuring DVMRP on page 4288
- Example: Configuring DVMRP on page 4288
- Example: Configuring DVMRP to Announce Unicast Routes on page 4292
- Tracing DVMRP Protocol Traffic on page 4295

Understanding DVMRP

The Distance Vector Multicast Routing Protocol (DVMRP) is a distance-vector routing protocol that provides connectionless datagram delivery to a group of hosts across an internetwork. DVMRP is a distributed protocol that dynamically generates IP multicast delivery trees by using a technique called reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces. These mechanisms allow the formation of shortest-path trees, which are used to reach all group members from each network source of multicast traffic.

DVMRP is designed to be used as an interior gateway protocol (IGP) within a multicast domain.

Because not all IP routers support native multicast routing, DVMRP includes direct support for tunneling IP multicast datagrams through routers. The IP multicast datagrams are encapsulated in unicast IP packets and addressed to the routers that do support native multicast routing. DVMRP treats tunnel interfaces and physical network interfaces the same way.

DVMRP routers dynamically discover their neighbors by sending neighbor probe messages periodically to an IP multicast group address that is reserved for all DVMRP routers.

Copyright © 2016, Juniper Networks, Inc.

Configuring DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) is the first of the multicast routing protocols and has a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G).

To configure the Distance Vector Multicast Routing Protocol (DVMRP), include the **dvmrp** statement:

```
dvmrp {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    interface interface-name {
        disable;
        hold-time seconds;
        metric metric;
        mode (forwarding | unicast-routing);
    }
    rib-group group-name;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

By default, DVMRP is disabled.

Example: Configuring DVMRP

This example shows how to use DVMRP to announce routes used for multicast routing as well as multicast data forwarding.

- Requirements on page 4288
- Overview on page 4289
- Configuration on page 4290
- Verification on page 4291

Requirements

Before you begin:

- Configure the router interfaces.

Copyright © 2016, Juniper Networks, Inc.

- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.

Overview

DVMRP is a distance vector protocol for multicast. It is similar to RIP, in that both RIP and DVMRP have issues with scalability and robustness. PIM domains are more commonly used than DVMRP domains. In some environments, you might need to configure interoperability with DVMRP.

This example includes the following DVMRP settings:

- **protocols dvmrp rib-group**—Associates the **dvmrp-rib** routing table group with the DVMRP protocol to enable multicast RPF lookup.
- **protocols dvmrp interface**—Configures the DVMRP interface. The interface of a DVMRP router can be either a physical interface to a directly attached subnetwork or a tunnel interface to another multicast-capable area of the Multicast Backbone (*MBone*). The DVMRP hold-time period is the amount of time that a neighbor is to consider the sending router (this router) to be operative (up). The default hold-time period is 35 seconds.
- **protocols dvmrp interface hold-time**—The DVMRP hold-time period is the amount of time that a neighbor is to consider the sending router (this router) to be operative (up). The default hold-time period is 35 seconds.
- **protocols dvmrp interface metric**—All interfaces can be configured with a metric specifying cost for receiving packets on a given interface. The default metric is 1.

  For each source network reported, a route metric is associated with the unicast route being reported. The metric is the sum of the interface metrics between the router originating the report and the source network. A metric of 32 marks the source network as unreachable, thus limiting the breadth of the DVMRP network and placing an upper bound on the DVMRP convergence time.
- **routing-options rib-groups**—Enables DVMRP to access route information from the unicast routing table, **inet.0**, and from a separate routing table that is reserved for DVMRP. In this example, the first routing table group named **ifrg** contains local interface routes. This ensures that local interface routes get added to both the **inet.0** table for use by unicast protocols and the **inet.2** table for multicast RPF check. The second routing table group named **dvmrp-rib** contains **inet.2** routes.

  DVMRP needs to access route information from the unicast routing table, **inet.0**, and from a separate routing table that is reserved for DVMRP. You need to create the routing table for DVMRP and to create groups of routing tables so that the routing protocol process imports and exports routes properly. We recommend that you use routing table **inet.2** for DVMRP routing information.
- **routing-options interface-routes**— After defining the **ifrg** routing table group, use the **interface-routes** statement to insert interface routes into the **ifrg** group—in other words,

Copyright © 2016, Juniper Networks, Inc.

into both **inet.0** and **inet.2**. By default, interface routes are imported into routing table **inet.0** only.

- **sap**—Enables the Session Directory Announcement Protocol (SAP) and the Session Directory Protocol (SDP). Enabling SAP allows the router to receive announcements about multimedia and other multicast sessions.

  SAP always listens to the address and port 224.2.127.254:9875 for session advertisements. To add other addresses or pairs of address and port, include one or more **listen** statements.

  Sessions learned by SDP, SAP's higher-layer protocol, time out after 60 minutes.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-options interface-routes rib-group inet ifrg
set routing-options rib-groups ifrg import-rib inet.0
set routing-options rib-groups ifrg import-rib inet.2
set routing-options rib-groups dvmrp-rib export-rib inet.2
set routing-options rib-groups dvmrp-rib import-rib inet.2
set protocols sap
set protocols dvmrp rib-group dvmrp-rib
set protocols dvmrp interface ip-0/0/0.0 metric 5
set protocols dvmrp interface ip-0/0/0.0 hold-time 40
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an MSDP routing instance:

1. Create the routing tables for DVMRP routes.

   ```
   [edit routing-options]
   user@host# set interface-routes rib-group inet ifrg
   user@host# set rib-groups ifrg import-rib [ inet.0 inet.2 ]
   user@host# set rib-groups dvmrp-rib import-rib inet.2
   user@host# set rib-groups dvmrp-rib export-rib inet.2
   ```

2. Configure SAP and SDP.

   ```
   [edit protocols]
   user@host# set sap
   ```

3. Enable DVMRP on the router and associate the **dvmrp-rib** routing table group with DVMRP to enable multicast RPF checks.

   ```
   [edit protocols]
   user@host# set dvmrp rib-group dvmrp-rib
   ```


Copyright © 2016, Juniper Networks, Inc.

4. Configure the DVMRP interface with a hold-time value and a metric. This example shows an IP-over-IP encapsulation tunnel interface.

```
[edit protocols]
user@host# set dvmrp interface ip–0/0/0.0
user@host# set dvmrp interface ip–0/0/0.0 hold-time 40
user@host# set dvmrp interface ip–0/0/0.0 metric 5
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show routing-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
interface-routes {
   rib-group inet ifrg;
}
   rib-groups {
      ifrg {
         import-rib [ inet.0 inet.2 ];
      }
      dvmrp-rib {
         export-rib inet.2;
         import-rib inet.2;
      }
   }
user@host# show protocols
sap;
dvmrp {
   rib-group dvmrp-rib;
   interface ip-0/0/0.0 {
      metric 5;
      hold-time 40;
   }
}
```

Verification

To verify the configuration, run the following commands:

- show dvmrp interfaces
- show dvmrp neighbors

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring DVMRP to Announce Unicast Routes

This example shows how to use DVMRP to announce unicast routes used solely for multicast reverse-path forwarding (RPF) to set up the multicast control plane.

- Requirements on page 4292
- Overview on page 4292
- Configuration on page 4293
- Verification on page 4295

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.

Overview

DVMRP has two modes. Forwarding mode is the default mode. In forwarding mode, DVMRP is responsible for the multicast control plane and multicast data forwarding. In the nondefault mode (which is shown in this example), DVMRP does not forward multicast data traffic. This mode is called unicast routing mode because in this mode DVMRP is only responsible for announcing unicast routes used for multicast RPF—in other words, for establishing the control plane. To forward multicast data, enable Protocol Independent Multicast (PIM) on the interface. If you have configured PIM on the interface, as shown in this example, you can configure DVMRP in unicast-routing mode only. You cannot configure PIM and DVMRP in forwarding mode at the same time.

This example includes the following settings:

- **policy-statement dvmrp-export**—Accepts static default routes.
- **protocols dvmrp export dvmrp-export**—Associates the **dvmrp-export** policy with the DVMRP protocol.

  All routing protocols use the routing table to store the routes that they learn and to determine which routes they advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table. Import and export policies are always from the point of view of the routing table. So the **dvmrp-export** policy exports static default routes from the routing table and accepts them into DVMRP.
- **protocols dvmrp interface all mode unicast-routing**—Enables all interfaces to announce unicast routes used solely for multicast RPF.
- **protocols dvmrp rib-group inet dvmrp-rg**—Associates the **dvmrp-rib** routing table group with the DVMRP protocol to enable multicast RPF checks.
- **protocols pim rib-group inet pim-rg**—Associates the **pim-rg** routing table group with the PIM protocol to enable multicast RPF checks.

Copyright © 2016, Juniper Networks, Inc.

- **routing-options rib inet.2 static route 0.0.0.0/0 discard**—Redistributes static routes to all DVMRP neighbors. The **inet.2** routing table stores unicast IPv4 routes for multicast RPF lookup. The **discard** statement silently drops packets without notice.
- **routing-options rib-groups dvmrp-rg import-rib inet.2**—Creates the routing table for DVMRP to ensure that the routing protocol process imports routes properly.
- **routing-options rib-groups dvmrp-rg export-rib inet.2**—Creates the routing table for DVMRP to ensure that the routing protocol process exports routes properly.
- **routing-options rib-groups pim-rg import-rib inet.2**—Enables access to route information from the routing table that stores unicast IPv4 routes for multicast RPF lookup. In this example, the first routing table group named **pim-rg** contains local interface routes. This ensures that local interface routes get added to the **inet.2** table.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement dvmrp-export term 10 from protocol static
set policy-options policy-statement dvmrp-export term 10 from route-filter 0.0.0.0/0
  exact
set policy-options policy-statement dvmrp-export term 10 then accept
set protocols dvmrp rib-group inet
set protocols dvmrp rib-group dvmrp-rg
set protocols dvmrp export dvmrp-export
set protocols dvmrp interface all mode unicast-routing
set protocols dvmrp interface fxp0.0 disable
set protocols pim rib-group inet pim-rg
set protocols pim interface all
set routing-options rib inet.2 static route 0.0.0.0/0 discard
set routing-options rib-groups pim-rg import-rib inet.2
set routing-options rib-groups dvmrp-rg export-rib inet.2
set routing-options rib-groups dvmrp-rg import-rib inet.2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an MSDP routing instance:

1. Configure the routing options.

```
[edit routing-options]
[edit routing -options]
user@host# set rib inet.2 static route 0.0.0.0/0 discard
user@host# set rib-groups pim-rg import-rib inet.2
user@host# set rib-groups dvmrp-rg import-rib inet.2
user@host# set rib-groups dvmrp-rg export-rib inet.2
```

2. Configure DVMRP.

Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols]
user@host# set dvmrp rib-group inet dvmrp-rg
user@host# set dvmrp export dvmrp-export
user@host# set dvmrp interface all mode unicast-routing
user@host# set dvmrp interface fxp0 disable
```

3. Configure PIM so that PIM performs multicast data forwarding.

```
[edit protocols]
user@host# set pim rib-group inet pim-rg
user@host# set pim interface all
```

4. Configure the DVMRP routing policy.

```
[edit policy-options policy-statement dvmrp-export term 10]
user@host# set from protocol static
user@host# set from route-filter 0.0.0.0/0 exact
user@host# set then accept
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** command, the **show protocols** command, and the **show routing-options** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement dvmrp-export {
    term 10 {
        from {
            protocol static;
            route-filter 0.0.0.0/0 exact;
        }
        then accept;
    }
}

user@host# show protocols
dvmrp {
    rib-group inet dvmrp-rg;
    export dvmrp-export;
    interface all {
        mode unicast-routing;
    }
    interface fxp0.0 {
        disable;
    }
}
pim {
    rib-group inet pim-rg;
    interface all;
}

user@host# show routing-options
```


Copyright © 2016, Juniper Networks, Inc.

```
rib inet.2 {
  static {
    route 0.0.0.0/0 discard;
  }
}
  rib-groups {
    pim-rg {
      import-rib inet.2;
    }
    dvmrp-rg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
```

Verification

To verify the configuration, run the following commands:

- show dvmrp interfaces
- show pim statistics

Tracing DVMRP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
general	Trace general flow.
graft	Trace graft messages.
neighbor	Trace neighbor probe packets.
normal	Trace normal events.
packets	Trace all DVMRP packets.
poison	Trace poison-route-reverse packets.
policy	Trace policy processing.
probe	Trace probe packets.
prune	Trace prune messages.

Copyright © 2016, Juniper Networks, Inc.

Flag	Description
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace task processing.
timer	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on DVMRP packets of a particular type. To configure tracing operations for DVMRP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

   ```
   [edit routing-options traceoptions]
   user@host# set file all-packets-trace
   user@host# set flag all
   ```

2. Configure the filename for the DVMRP trace file.

   ```
   [edit protocols dvmrp traceoptions]
   user@host# set file dvmrp-trace
   ```

3. (Optional) Configure the maximum number of trace files.

   ```
   [edit protocols dvmrp traceoptions]
   user@host# set file files 5
   ```

4. (Optional) Configure the maximum size of each trace file.

   ```
   [edit protocols dvmrp traceoptions]
   user@host# set file size 1m
   ```

5. (Optional) Enable unrestricted file access.

   ```
   [edit protocols dvmrp traceoptions]
   user@host# set file world-readable
   ```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular DVMRP neighbor. The following example shows how to trace neighbor probe packets that match the neighbor's IP address.

   ```
   [edit protocols dvmrp traceoptions]
   user@host# set flag neighbor | match 192.168.1.1
   ```

7. View the trace file.

   ```
   user@host> file list /var/log
   user@host> file show /var/log/dvmrp-trace
   ```

Related Documentation

- Understanding DVMRP on page 4287

Copyright © 2016, Juniper Networks, Inc.

PART 63

Configuring Multicast VPNs

- Configuring PIM Join Load Balancing on page 4299
- Configuring Next-Generation Multicast VPNs on page 4303

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 191

Configuring PIM Join Load Balancing

- PIM Join Load Balancing on Multipath MVPN Routes Overview on page 4299

PIM Join Load Balancing on Multipath MVPN Routes Overview

A multicast virtual private network (MVPN) is a technology to deploy the multicast service in an existing MPLS/BGP VPN.

The two main MVPN services are:

- Dual PIM MVPNs (also referred to as Draft-Rosen)
- Multiprotocol BGP-based MVPNs (also referred to as next-generation)

Next-generation MVPNs constitute the next evolution after the Draft-Rosen MVPN and provide a simpler solution for administrators who want to configure multicast over Layer 3 VPNs. A Draft-Rosen MVPN uses Protocol Independent Multicast (PIM) for customer multicast (C-multicast) signaling, and a next-generation MVPN uses BGP for C-multicast signaling.

Multipath routing in an MVPN is applied to make data forwarding more robust against network failures and to minimize shared backup capacities when resilience against network failures is required.

By default, PIM join messages are sent toward a source based on the reverse path forwarding (RPF) routing table check. If there is more than one equal-cost path toward the source [S, G] or rendezvous point (RP) [*, G], then one upstream interface is used to send the join messages. The upstream path can be:

- A single active external BGP (EBGP) path when both EBGP and internal BGP (IBGP) paths are present.
- A single active IBGP path when there is no EBGP path present.

With the introduction of the multipath PIM join load-balancing feature, customer PIM (C-PIM) join messages are load-balanced in the following ways:

- In the case of a Draft-Rosen MVPN, unequal EBGP and IBGP paths are utilized.
- In the case of next-generation MVPN:
  - Available IBGP paths are utilized when no EBGP path is present.

Copyright © 2016, Juniper Networks, Inc.

- Available EBGP paths are utilized when both EBGP and IBGP paths are present.

This feature is applicable to IPv4 C-PIM join messages over the Layer 3 MVPN service.

By default, a customer source (C-S) or a customer RP (C-RP) is considered remote if the active **rt_entry** is a secondary route and the primary route is present in a different routing instance. Such determination is being done without taking into consideration the (C-*,G) or (C-S,G) state for which the check is being performed. The multipath PIM join load-balancing feature determines if a source (or RP) is remote by taking into account the associated (C-*,G) or (C-S,G) state.

When the provider network does not have provider edge (PE) routers with the multipath PIM join load-balancing feature enabled, hash-based join load balancing is used. Although the decision to configure this feature does not impact PIM or overall system performance, network performance can be affected temporarily, if the feature is not enabled.

With hash-based join load balancing, adding new PE routers to the candidate upstream toward the C-S or C-RP results in C-PIM join messages being redistributed to new upstream paths. If the number of join messages is large, network performance is impacted because of join messages being sent to the new RPF neighbor and prune messages being sent to the old RPF neighbor. In next-generation MVPN, this results in BGP C-multicast data messages being withdrawn from old upstream paths and advertised on new upstream paths, impacting network performance.

Copyright © 2016, Juniper Networks, Inc.

In Figure 194, PE1 and PE2 are the upstream PE routers. Router PE1 learns route Source from EBGP and IBGP peers—the customer edge CE1 router and the PE2 router, respectively.

Figure 194: PIM Join Load Balancing

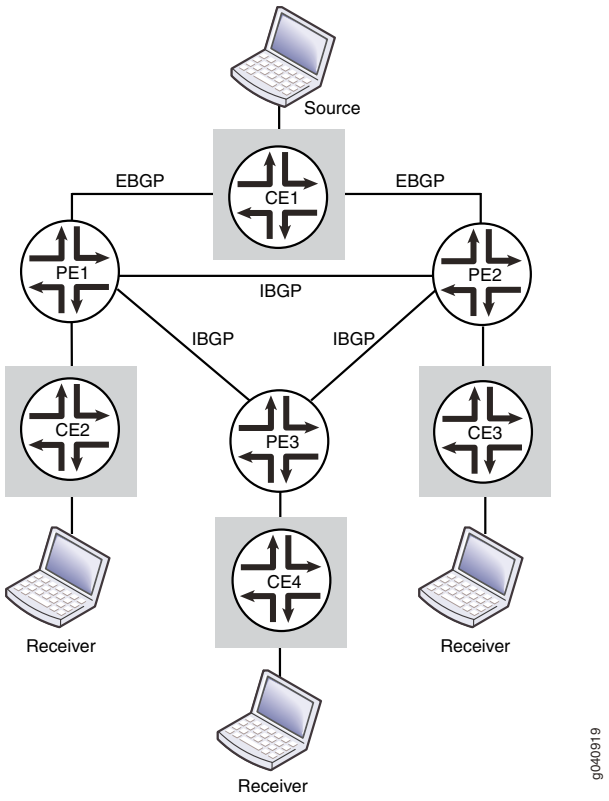

- If the PE routers run the Draft-Rosen MVPN, the PE1 router distributes C-PIM join messages between the EBGP path to the CE1 router and the IBGP path to the PE2 router. The join messages on the IBGP path are sent over a multicast tunnel interface through which the PE routers establish C-PIM adjacency with each other.

  If a PE router loses one or all EBGP paths toward the source (or RP), the C-PIM join messages that were previously using the EBGP path are moved to a multicast tunnel interface, and the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

  On discovering the first EBGP path toward the source (or RP), only new join messages get load-balanced across EBGP and IBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

- If the PE routers run the next-generation MVPN, the PE1 router sends C-PIM join messages directly to the CE1 router over the EBGP path. There is no C-PIM adjacency between the PE1 and PE2 routers. Router PE3 distributes the C-PIM join messages between the two IBGP paths to PE1 and PE2. The Bytewise-XOR hash algorithm is used to send the C-multicast data according to Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*.

Copyright © 2016, Juniper Networks, Inc.

Because the multipath PIM join load-balancing feature in a Draft-Rosen MVPN utilizes unequal EBGP and IBGP paths to the destination, loops can be created when forwarding unicast packets to the destination. To avoid or break such loops:

- Traffic arriving from a core or master instance should not be forwarded back to the core facing interfaces.
- A single multicast tunnel interface should either be selected as the upstream interface or the downstream interface.
- An upstream or downstream multicast tunnel interface should point to a non-multicast tunnel interface.

As a result of the loop avoidance mechanism, join messages arriving from an EBGP path get load-balanced across EIBGP paths as expected, whereas join messages from an IBGP path are constrained to choose the EBGP path only.

In Figure 194, if the CE2 host sends unicast data traffic to the CE1 host, the PE1 router could send the multicast flow to the PE2 router over the MPLS core due to traffic load balancing. A data forwarding loop is prevented by ensuring that PE2 does not forward traffic back on the MPLS core because of the load-balancing algorithm.

In the case of C-PIM join messages, assuming that both the CE2 host and the CE3 host are interested in receiving traffic from the source (S, G), and if both PE1 and PE2 choose each other as the RPF neighbor toward the source, then a multicast tree cannot be formed completely. This feature implements mechanisms to prevent such join loops in the multicast control plane in a Draft-Rosen MVPN scenario.

NOTE:

Disruption of multicast traffic or creation of join loops can occur, resulting in a multicast distribution tree (MDT) not being formed properly due to one of the following reasons:

- During a graceful Routing Engine switchover (GRES), the EIBGP path selection for C-PIM join messages can vary, because the upstream interface selection is performed again for the new Routing Engine based on the join messages it receives from the CE and PE neighbors. This can lead to disruption of multicast traffic depending on the number of join messages received and the load on the network at the time of the graceful restart. However, nonstop active routing (NSR) is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.
- Any PE router in the provider network is running another vendor's implementation that does not apply the same hashing algorithm implemented in this feature.
- The multipath PIM join load-balancing feature has not been configured properly.

Related Documentation

- Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN on page 4303

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 192

Configuring Next-Generation Multicast VPNs

- Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN on page 4303

Example: Configuring PIM Join Load Balancing on Next-Generation Multicast VPN

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running next-generation multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across available internal BGP (IBGP) upstream paths when there is no external BGP (EBGP) path present, and across available EBGP upstream paths when external and internal BGP (EIBGP) paths are present toward the source or rendezvous point (RP).

- Requirements on page 4303
- Overview and Topology on page 4304
- Configuration on page 4306
- Verification on page 4310

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series, MX Series, or T Series routers.
- Junos OS Release 12.1 running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
   - OSPF
   - MPLS
   - LDP

Copyright © 2016, Juniper Networks, Inc.

- PIM
- BGP

3. Configure a multicast VPN.

Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across all available IBGP paths when there are only IBGP paths present, and across all available upstream EBGP paths when EIBGP paths are present toward the source (or RP). Unlike Draft-Rosen MVPN, next-generation MVPN does not utilize unequal EIBGP paths to send C-PIM join messages. This feature is applicable to IPv4 C-PIM join messages.

By default, only one active IBGP path is used to send the C-PIM join messages for a PE router having only IBGP paths toward the source (or RP). When there are EIBGP upstream paths present, only one active EBGP path is used to send the join messages.

In a next-generation MVPN, C-PIM join messages are translated into (or encoded as) BGP customer multicast (C-multicast) MVPN routes and advertised with the BGP MCAST-VPN address family toward the sender PE routers. A PE router originates a C-multicast MVPN route in response to receiving a C-PIM join message through its PE router to customer edge (CE) router interface. The two types of C-multicast MVPN routes are:

- Shared tree join route (C-*, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a shared tree C-PIM join message through its PE-CE router interface.
- Source tree join route (C-S, C-G)
  - Originated by receiver PE routers.
  - Originated when a PE router receives a source tree C-PIM join message (C-S, C-G), or originated by the PE router that already has a shared tree join route and receives a source active autodiscovery route.

  The upstream path in a next-generation MVPN is selected using the Bytewise-XOR hash algorithm as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast, *Multicast in MPLS/BGP IP VPNs*. The hash algorithm is performed as follows:

  1. The PE routers in the candidate set are numbered from lower to higher IP address, starting from **0**.
  2. A bytewise exclusive-or of all the bytes is performed on the C-root (source) and the C-G (group) address.

Copyright © 2016, Juniper Networks, Inc.

3. The result is taken modulo *n*, where *n* is the number of PE routers in the candidate set. The result is **N**.

4. **N** represents the IP address of the upstream PE router as numbered in Step 1.

During load balancing, if a PE router with one or more upstream IBGP paths toward the source (or RP) discovers a new IBGP path toward the same source (or RP), the C-PIM join messages distributed among previously existing IBGP paths get redistributed due to the change in the candidate PE router set.

In this example, PE1, PE2, and PE3 are the PE routers that have the multipath PIM join load-balancing feature configured. Router PE1 has two EBGP paths and one IBGP upstream path, PE2 has one EBGP path and one IBGP upstream path, and PE3 has two IBGP upstream paths toward the Source. Router CE4 is the customer edge (CE) router attached to PE3. Source and Receiver are the Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The C-PIM join messages are sent using EBGP paths only. IBGP paths are not used to propagate the join messages.

   In Figure 195, the PE1 router distributes the join messages between the two EBGP paths to the CE1 router, and PE2 uses the EBGP path to CE1 to send the join messages.

2. If a PE router loses one or more EBGP paths toward the source (or RP), the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

   On discovering the first EBGP path, only new join messages get load-balanced across available EBGP paths, whereas the existing join messages on the multicast tunnel interface are not redistributed.

   If the EBGP path from the PE2 router to the CE1 router goes down, PE2 sends the join messages to PE1 using the IBGP path. When the EBGP path to CE1 is restored, only new join messages that arrive on PE2 use the restored EBGP path, whereas join messages already sent on the IBGP path are not redistributed.

On PE routers that have only IBGP paths toward the source (or RP), such as the PE3 router, PIM join load balancing is performed as follows:

1. The C-PIM join messages from CE routers get load-balanced only as BGP C-multicast data messages among IBGP paths.

   In Figure 195, assuming that the CE4 host is interested in receiving traffic from the Source, and CE4 initiates source join messages for different groups (Group 1 [C-S,C-G1] and Group 2 [C-S,C-G2]), the source join messages arrive on the PE3 router.

   Router PE3 then uses the Bytewise-XOR hash algorithm to select the upstream PE router to send the C-multicast data for each group. The algorithm first numbers the upstream PE routers from lower to higher IP address starting from **0**.

   Assuming that Router PE1 router is numbered **0** and Router PE2 is **1**, and the hash result for Group 1 and Group 2 join messages is **0** and **1**, respectively, the PE3 router

Copyright © 2016, Juniper Networks, Inc.

selects PE1 as the upstream PE router to send Group 1 join messages, and PE2 as the upstream PE router to send the Group 2 join messages to the Source.

2. The shared join messages for different groups [C-*,C-G] are also treated in a similar way to reach the destination.

Figure 195: PIM Join Load Balancing on Next-Generation MVPN

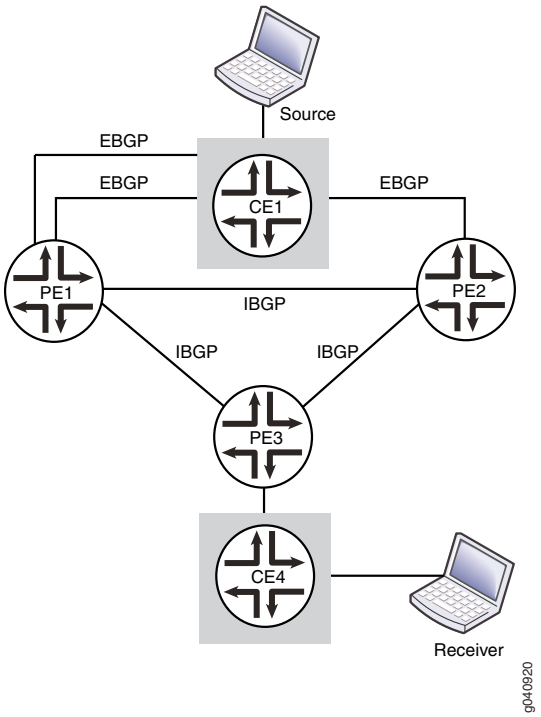

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

PE1

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-3/0/1.0
set routing-instances vpn1 interface ge-3/3/2.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 1:1
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
```


Copyright © 2016, Juniper Networks, Inc.

```
set routing-instances vpn1 protocols bgp group bgp local-address 10.40.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.40.10.2 peer-as 3
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.10.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.10.10.2 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

PE2

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-1/0/9.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 3
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

PE3

```
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/8.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 3:3
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 routing-options autonomous-system 1
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.80.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.80.10.2 peer-as 2
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bytewise-xor-hash
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*. To configure the PE1 router:

NOTE: Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing forwarding (VRF) routing instance.

```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-3/0/1.0
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set provider-tunnel rsvp-te label-switched-path-template
  default-template
user@PE1# set vrf-target target:1:1
user@PE1# set vrf-table-label
```

2. Enable protocol-independent load balancing for the VRF instance.

```
[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 10.40.10.1
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 10.40.10.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 10.10.10.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 10.10.10.2 peer-as 3
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim rp static address 10.255.10.119
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

7. Configure the mode for C-PIM join messages to use rendezvous-point trees, and switch to the shortest-path tree after the source is known.

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-mode rpt-spt
```

8. Configure the VRF instance to use the Bytewise-XOR hash algorithm.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-join-load-balance bytewise-xor-hash
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show routing-instances
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-3/0/1.0;
    interface ge-3/3/2.0;
    interface lo0.1;
    route-distinguisher 1:1;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          default-template;
        }
      }
    }
    vrf-target target:1:1;
    vrf-table-label;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
    protocols {
      bgp {
        export direct;
        group bgp {
          type external;
          local-address 10.40.10.1;
          family inet {
            unicast;
          }
          neighbor 10.40.10.2 {
            peer-as 3;
          }
        }
        group bgp1 {
          type external;
          local-address 10.10.10.1;
          family inet {
            unicast;
          }
          neighbor 10.10.10.2 {
```

Copyright © 2016, Juniper Networks, Inc.

```
            peer-as 3;
          }
        }
      }
      pim {
        rp {
          static {
            address 10.255.10.119;
          }
        }
        interface all;
        join-load-balance;
      }
      mvpn {
        mvpn-mode {
          rpt-spt;
        }
        mvpn-join-load-balance {
          bytewise-xor-hash;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages on page 4310

Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages

Purpose Verify MVPN C-multicast route information for different groups of join messages received on the PE3 router.

Action From operational mode, run the **show mvpn c-multicast** command.

```
user@PE3>
MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : vpn1
  MVPN Mode : RPT-SPT
  C-mcast IPv4 (S:G)            Ptnl                              St
  0.0.0.0/0:225.1.1.1/32        RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
  4.4.4.2/32:225.1.1.1/32       RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
  0.0.0.0/0:225.1.1.2/32        RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
  4.4.4.2/32:225.1.1.2/32       RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
```


Copyright © 2016, Juniper Networks, Inc.

Meaning The output shows how the PE3 router has load-balanced the C-multicast data for the different groups.

- For source join messages (S,G):
  - 4.4.4.2/32:225.1.1.1/32 (S,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 4.4.4.2/32:225.1.1.2/32 (S,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).
- For shared join messages (*,G):
  - 0.0.0.0/0:225.1.1.1/32 (*,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
  - 0.0.0.0/0:225.1.1.2/32 (*,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).

Related Documentation

- PIM Join Load Balancing on Multipath MVPN Routes Overview on page 4299

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 64

Configuring General Multicast Routing Options

- Preventing Routing Loops with Reverse Path Forwarding on page 4315
- Enabling Multicast Between Layer 2 and Layer 3 Devices Using Snooping on page 4329
- Configuring Multicast Routing Options on page 4353

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 193

Preventing Routing Loops with Reverse Path Forwarding

- Examples: Configuring Reverse Path Forwarding on page 4315

Examples: Configuring Reverse Path Forwarding

- Understanding Multicast Reverse Path Forwarding on page 4315
- Multicast RPF Configuration Guidelines on page 4317
- Example: Configuring a Dedicated PIM RPF Routing Table on page 4317
- Example: Configuring RPF Policies on page 4321
- Example: Configuring PIM RPF Selection on page 4323

Understanding Multicast Reverse Path Forwarding

Unicast forwarding decisions are typically based on the destination address of the packet arriving at a router. The unicast routing table is organized by destination subnet and mainly set up to forward the packet toward the destination.

In multicast, the router forwards the packet away from the source to make progress along the distribution tree and prevent routing loops. The router's multicast forwarding state runs more logically by organizing tables based on the reverse path, from the receiver back to the root of the distribution tree. This process is known as *reverse-path forwarding* (*RPF*).

The router adds a branch to a distribution tree depending on whether the request for traffic from a multicast group passes the reverse-path-forwarding check (RPF check). Every multicast packet received must pass an RPF check before it is eligible to be replicated or forwarded on any interface.

The RPF check is essential for every router's multicast implementation. When a multicast packet is received on an interface, the router interprets the source address in the multicast IP packet as the destination address for a unicast IP packet. The source multicast address is found in the unicast routing table, and the outgoing interface is determined. If the outgoing interface found in the unicast routing table is the same as the interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped because the incoming interface is not on the *shortest path* back to the source.

Copyright © 2016, Juniper Networks, Inc.

Figure 196 shows how multicast routers can use the unicast routing table to perform an RPF check and how the results obtained at each router determine where join messages are sent.

Figure 196: Multicast Routers and the RPF Check

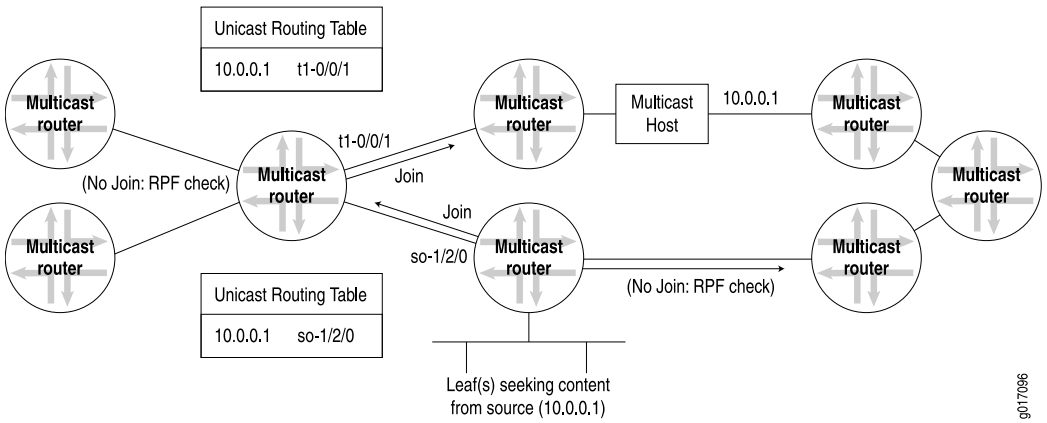

Routers can build and maintain separate tables for RPF purposes. The router must have some way to determine its RPF interface for the group, which is the interface topologically closest to the root. For greatest efficiency, the distribution tree follows the shortest-path tree topology. The RPF check helps to construct this tree.

RPF Table

The RPF table plays the key role in the multicast router. The RPF table is consulted for every RPF check, which is performed at intervals on multicast packets entering the multicast router. Distribution trees of all types rely on the RPF table to form properly, and the multicast forwarding state also depends on the RPF table.

RPF checks are performed only on unicast addresses to find the upstream interface for the multicast source or RP.

The routing table used for RPF checks can be the same routing table used to forward unicast IP packets, or it can be a separate routing table used only for multicast RPF checks. In either case, the RPF table contains only unicast routes, because the RPF check is performed on the source address of the multicast packet, not the multicast group destination address, and a multicast address is forbidden from appearing in the source address field of an IP packet header. The unicast address can be used for RPF checks because there is only one source host for a particular stream of IP multicast content for a multicast group address, although the same content could be available from multiple sources.

If the same routing table used to forward unicast packets is also used for the RPF checks, the routing table is populated and maintained by the traditional unicast routing protocols such as BGP, IS-IS, OSPF, and the Routing Information Protocol (RIP). If a dedicated multicast RPF table is used, this table must be populated by some other method. Some multicast routing protocols (such as the Distance Vector Multicast Routing Protocol [DVMRP]) essentially duplicate the operation of a unicast routing protocol and populate a dedicated RPF table. Others, such as PIM, do not duplicate routing protocol functions

Copyright © 2016, Juniper Networks, Inc.

and must rely on some other routing protocol to set up this table, which is why PIM is *protocol independent*.

Some traditional routing protocols such as BGP and IS-IS now have extensions to differentiate between different sets of routing information sent between routers for unicast and multicast. For example, there is multiprotocol BGP (MBGP) and multitopology routing in IS-IS (M-IS-IS). IS-IS routes can be added to the RPF table even when special features such as traffic engineering and "shortcuts" are turned on. Multicast Open Shortest Path First (MOSPF) also extends OSPF for multicast use, but goes further than MBGP or M-IS-IS and makes MOSPF into a complete multicast routing protocol on its own. When these routing protocols are used, routes can be tagged as multicast RPF routers and used by the receiving router differently than the unicast routing information.

Using the main unicast routing table for RPF checks provides simplicity. A dedicated routing table for RPF checks allows a network administrator to set up separate paths and routing policies for unicast and multicast traffic, allowing the multicast network to function more independently of the unicast network.

Multicast RPF Configuration Guidelines

You use multicast RPF checks to prevent multicast routing loops. Routing loops are particularly debilitating in multicast applications because packets are replicated with each pass around the routing loop.

In general, a router is to forward a multicast packet only if it arrives on the interface closest (as defined by a unicast routing protocol) to the origin of the packet, whether source host or rendezvous point (RP). In other words, if a unicast packet would be sent to the "destination" (the reverse path) on the interface that the multicast packet arrived on, the packet passes the RPF check and is processed. Multicast (or unicast) packets that fail the RPF check are not forwarded (this is the default behavior). For an overview of how a Juniper Networks router implements RPF checks with tables, see "Understanding Multicast Reverse Path Forwarding" on page 4315.

However, there are network router configurations where multicast packets that fail the RPF check need to be forwarded. For example, when point-to-multipoint label-switched paths (LSPs) are used for distributing multicast traffic to PIM "islands" downstream from the egress router, the interface on which the multicast traffic arrives is not always the RPF interface. This is because LSPs do not follow the normal next-hop rules of independent packet routing.

In cases such as these, you can configure policies on the PE router to decide which multicast groups and sources are exempt from the default RPF check.

Example: Configuring a Dedicated PIM RPF Routing Table

This example explains how to configure a dedicated Protocol Independent Multicast (PIM) reverse path forwarding (RPF) routing table.

- Requirements on page 4318
- Overview on page 4318
- Configuration on page 4319

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure the router interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Enable PIM. See "PIM Overview" on page 4089.

This example uses the following software components:

- Junos OS Release 7.4 or later

Overview

By default, PIM uses the **inet.0** routing table as its RPF routing table. PIM uses an RPF routing table to resolve its RPF neighbor for a particular multicast source address and to resolve the RPF neighbor for the rendezvous point (RP) address. PIM can optionally use **inet.2** as its RPF routing table. The **inet.2** routing table is dedicated to this purpose.

PIM uses a single routing table for its RPF check, this ensures that the route with the longest matching prefix is chosen as the RPF route.

If multicast routes are exchanged by Multiprotocol Border Gateway Protocol MP-BGP or multitopology IS-IS, they are placed in **inet.2** by default.

Using **inet.2** as the RPF routing table enables you to have a control plane for multicast, which is independent of the normal unicast routing table. You might want to use **inet.2** as the RPF routing table for any of the following reasons:

- If you use traffic engineering or have an interior gateway protocol (IGP) configured for shortcuts, the router has label-switched paths (LSPs) installed as the next hops in **inet.2**. By applying policy, you can have the router install the routes with non-MPLS next-hops in the **inet.2** routing table.
- If you have an MPLS network that does not support multicast traffic over LSP tunnels, you need to configure the router to use a routing table other than **inet.0**. You can have the **inet.2** routing table populated with native IGP, BGP, and interface routes that can be used for RPF.

To populate the PIM RPF table, you use rib groups. A rib group is defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level. The rib group is applied to the PIM protocol by including the **rib-group** statement at the **[edit pim]** hierarchy level. A rib group is most frequently used to place routes in multiple routing tables.

When you configure rib groups for PIM, keep the following in mind:

- The **import-rib** statement copies routes from the protocol to the routing table.
- The **export-rib** statement has no effect on PIM.
- Only the first rib routing table specified in the **import-rib** statement is used by PIM for RPF checks.

Copyright © 2016, Juniper Networks, Inc.

You can also configure IS-IS or OSPF to populate **inet.2** with routes that have regular IP next hops. This allows RPF to work properly even when MPLS is configured for traffic engineering, or when IS-IS or OSPF are configured to use “shortcuts” for local traffic.

You can also configure the PIM protocol to use a rib group for RPF checks under a virtual private network (VPN) routing instance. In this case the rib group is still defined at the **[edit routing-options]** hierarchy level.

Configuration

Configuring a PIM RPF Routing Table Group Using Interface Routes

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-options rib-groups mcast-rpf-rib import-rib inet.2
set protocols pim rib-group mcast-rpf-rib
set routing-options interface-routes rib-group inet if-rib
set routing-options rib-groups if-rib import-rib [ inet.0 inet.2 ]
```

Step-by-Step Procedure

In this example, the network administrator has decided to use the **inet.2** routing table for RPF checks. In this process, local routes are copied into this table by using an interface rib group.

To define an interface routing table group and use it to populate **inet.2** for RPF checks:

1. Use the **show multicast rpf** command to verify that the multicast RPF table is not populated with routes.

   ```
   user@host> show multicast rpf
   instance is not running
   ```

2. Create a multicast routing table group named **mcast-rpf-rib**.

   Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the **import-rib** statement).

   Include the **import-rib** statement and specify the **inet.2** routing table at the **[edit routing-options rib-groups]** hierarchy level.

   ```
   [edit routing-options rib-groups]
   user@host# set mcast-rpf-rib import-rib inet.2
   ```

3. Configure PIM to use the **mcast-rpf-rib** rib group.

   The rib group for PIM can be applied globally or in a routing instance. In this example, the global configuration is shown.

   Include the **rib-group** statement and specify the **mcast-rpf-rib** rib group at the **[edit protocols pim]** hierarchy level.

   ```
   [edit protocols pim]
   user@host# set rib-group mcast-rpf-rib
   ```

4. Create an interface rib group named **if-rib**.

Copyright © 2016, Juniper Networks, Inc.

Include the **rib-group** statement and specify the **inet** address family at the **[edit routing-options interface-routes]** hierarchy level.

```
[edit routing-options interface-routes]
user@host# set rib-group inet if-rib
```

5. Configure the **if-rib** rib group to import routes from the **inet.0** and **inet.2** routing tables.

   Include the **import-rib** statement and specify the **inet.0** and **inet.2** routing tables at the **[edit routing-options rib-groups]** hierarchy level.

```
[edit routing-options rib-groups]
user@host# set if-rib import-rib [ inet.0 inet.2 ]
```

6. Commit the configuration.

```
user@host# commit
```

Verifying Multicast RPF Table

Purpose Verify that the multicast RPF table is now populated with routes.

Action Use the **show multicast rpf** command.

```
user@host> show multicast rpf
Multicast RPF table: inet.2 , 10 entries

10.0.24.12/30
    Protocol: Direct
    Interface: fe-0/1/2.0

10.0.24.13/32
    Protocol: Local

10.0.27.12/30
    Protocol: Direct
    Interface: fe-0/1/3.0

10.0.27.13/32
    Protocol: Local

10.0.224.8/30
    Protocol: Direct
    Interface: ge-1/3/3.0

10.0.224.9/32
    Protocol: Local

127.0.0.1/32
    Inactive

192.168.2.1/32
    Protocol: Direct
    Interface: lo0.0

192.168.187.0/25
    Protocol: Direct
    Interface: fxp0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
192.168.187.12/32
    Protocol: Local
```

Meaning The first line of the sample output shows that the **inet.2** table is being used and that there are 10 routes in the table. The remainder of the sample output lists the routes that populate the **inet.2** routing table.

Example: Configuring RPF Policies

A multicast RPF policy disables RPF checks for a particular multicast (S,G) pair. You usually disable RPF checks on egress routing devices of a point-to-multipoint label-switched path (LSP), because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

This example shows how to configure an RPF check policy named **disable-RPF-on-PE**. The **disable-RPF-on-PE** policy disables RPF checks on packets arriving for group 228.0.0.0/8 or from source address 196.168.25.6.

- Requirements on page 4321
- Overview on page 4321
- Configuration on page 4322
- Verification on page 4323

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.

Overview

An RPF policy behaves like an import policy. If no policy term matches the input packet, the default action is to accept (that is, to perform the RPF check). The **route-filter** statement filters group addresses, and the **source-address-filter** statement filters source addresses.

This example shows how to configure each condition as a separate policy and references both policies in the **rpf-check-policy** statement. This allows you to associate groups in one policy and sources in the other.

NOTE: **Be careful when disabling RPF checks on multicast traffic. If you disable RPF checks in some configurations, multicast loops can result.**

Changes to an RPF check policy take effect immediately:

Copyright © 2016, Juniper Networks, Inc.

- If no policy was previously configured, the policy takes effect immediately.
- If the policy name is changed, the new policy takes effect immediately and any packets no longer filtered are subjected to the RPF check.
- If the policy is deleted, all packets formerly filtered are subjected to the RPF check.
- If the underlying policy is changed, but retains the same name, the new conditions take effect immediately and any packets no longer filtered are subjected to the RPF check.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement disable-RPF-from-group term first from route-filter
  228.0.0.0/8 orlonger
set policy-options policy-statement disable-RPF-from-group term first then reject
set policy-options policy-statement disable-RPF-from-source term first from
  source-address-filter 192.168.25.6/32 exact
set policy-options policy-statement disable-RPF-from-source term first then reject
set routing-options multicast rpf-check-policy [ disable-RPF-from-group
  disable-RPF-from-source ]
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure a policy for group addresses.

   ```
   [edit policy-options]
   user@host# set policy-statement disable-RPF-for-group term first from route-filter
     228.0.0.0/8 orlonger
   user@host# set policy-statement disable-RPF-for-group term first then reject
   ```

2. Configure a policy for a source address.

   ```
   [edit policy-options]
   user@host# set policy-statement disable-RPF-for-source term first from
     source-address-filter 192.168.25.6/32 exact
   user@host# set policy-statement disable-RPF-for-source term first then reject
   ```

3. Apply the policies.

   ```
   [edit routing-options]
   user@host# set multicast rpf-check-policy [ disable-RPF-for-group
     disable-RPF-for-source ]
   ```

4. If you are done configuring the device, commit the configuration.

   ```
   user@host# commit
   ```


Copyright © 2016, Juniper Networks, Inc.

Results

Confirm your configuration by entering the **show policy-options** and **show routing-options** commands.

```
user@host# show policy-options
policy-statement disable-RPF-from-group {
   term first {
      from {
         route-filter 228.0.0.0/8 orlonger;
      }
      then reject;
   }
}
policy-statement disable-RPF-from-source {
   term first {
      from {
         source-address-filter 192.168.25.6/32 exact;
      }
      then reject;
   }
}
```

```
user@host# show routing-options
multicast {
   rpf-check-policy [ disable-RPF-from-group disable-RPF-from-source ];
}
```

Verification

To verify the configuration, run the show multicast rpf command.

Example: Configuring PIM RPF Selection

This example shows how to configure and verify the multicast PIM RPF next-hop neighbor selection for a group or (S,G) pair.

- Requirements on page 4323
- Overview on page 4324
- Configuration on page 4325
- Verification on page 4327

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Security Devices*.
- Make sure that the RPF next-hop neighbor you want to specify is operating.

Copyright © 2016, Juniper Networks, Inc.

Overview

Multicast PIM RPF neighbor selection allows you to specify the RPF neighbor (next hop) and source address for a single group or multiple groups using a prefix list. RPF neighbor selection can only be configured for VPN routing and forwarding (VRF) instances.

If you have multiple service VRFs through which a receiver VRF can learn the same source or rendevous point (RP) address, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows. However, if RPF neighbor selection is configured, RPF checks are based on your configuration instead of the unicast routing protocols.

You can use this static RPF selection as a building block for particular applications. For example, an extranet. Suppose you want to split the multicast flows among parallel PIM links or assign one multicast flow to a specific PIM link. With static RPF selection configured, the router sends join and prune messages based on the configuration.

You can use wildcards to designate the source address. Whether or not you use wildcards affects how the PIM joins work:

- If you configure only a source prefix for a group, all (*,G) joins are sent to the next-hop neighbor selected by the unicast protocol, while (S,G) joins are sent to the next-hop neighbor specified for the source.
- If you configure only a wildcard source for a group, all (*,G) and (S,G) joins are sent to the upstream interface pointing to the wildcard source next-hop neighbor.
- If you configure both a source prefix and a wildcard source for a group, all (S,G) joins are sent to the next-hop neighbor defined for the source prefix, while (*,G) joins are sent to the next-hop neighbor specified for the wildcard source.

Figure 197 shows the topology used in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 197: PIM RPF Selection

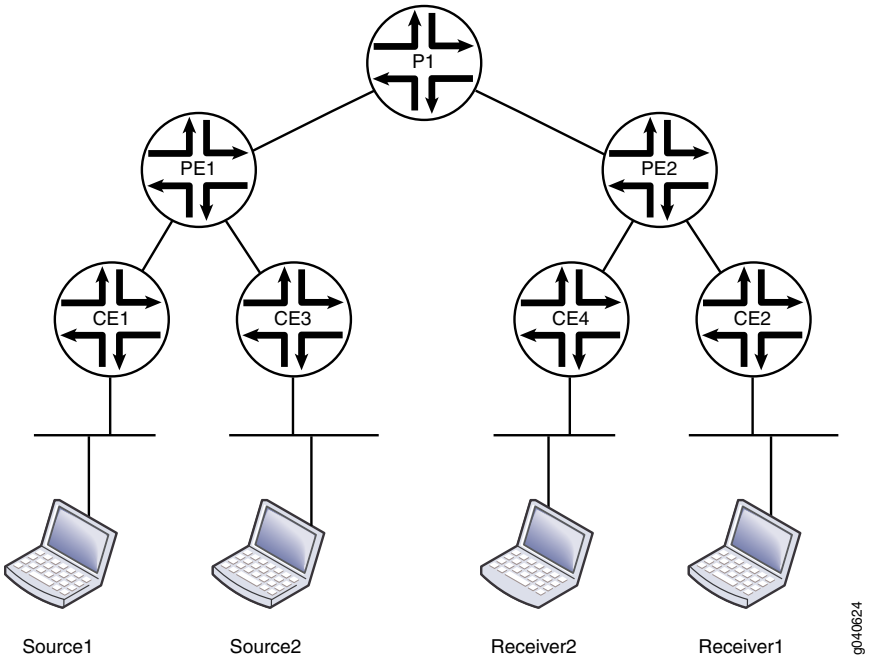

In this example, the RPF selection is configured on the receiver provider edge router (PE2).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instance vpn-a protocols pim rpf-selection group 225.5.0.0/16 wildcard-source
    next-hop 10.12.5.2
set routing-instance vpn-a protocols pim rpf-selection prefix-list group12 wildcard-source
    next-hop 10.12.31.2
set routing-instance vpn-a protocols pim rpf-selection prefix-list group34 source
    22.1.12.0/24 next-hop 10.12.32.2
set policy-options prefix-list group12 225.1.1.0/24
set policy-options prefix-list group12 225.2.0.0/16
set policy-options prefix-list group34 225.3.3.3/32
set policy-options prefix-list group34 225.4.4.0/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM RPF selection:

1. On PE2, configure RFP selection in a routing instance.

```
[edit routing-instance vpn-a protocols pim]
user@host# set rpf-selection group 225.5.0.0/16 wildcard-source next-hop 10.12.5.2
user@host# set rpf-selection prefix-list group12 wildcard-source next-hop 10.12.31.2
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set rpf-selection prefix-list group34 source 22.1.12.0/24 next-hop
  10.12.32.2
user@host# exit
```

2. On PE2, configure the policy.

```
[edit policy-options]
set prefix-list group12 225.1.1.0/24
set prefix-list group12 225.2.0.0/16
set prefix-list group34 225.3.3.3/32
set prefix-list group34 225.4.4.0/24
```

3. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results From configuration mode, confirm your configuration by entering the **show policy-options** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
prefix-list group12 {
  225.1.1.0/24;
  225.2.0.0/16;
}
prefix-list group34 {
  225.3.3.3/32;
  225.4.4.0/24;
}
```

```
user@host# show routing-instances
vpn-a{
  protocols {
    pim {
      rpf-selection {
        group 225.5.0.0/16 {
          wildcard-source {
            next-hop 10.12.5.2;
          }
        }
        prefix-list group12 {
          wildcard-source {
            next-hop 10.12.31.2;
          }
        }
        prefix-list group34 {
          source 22.1.12.0/24 {
            next-hop 10.12.32.2;
          }
        }
      }
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration, run the following commands, checking the upstream interface and the upstream neighbor:

- show pim join **extensive**
- show multicast route

Related Documentation

- Example: Configuring Ingress PE Redundancy on page 4381

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 194

Enabling Multicast Between Layer 2 and Layer 3 Devices Using Snooping

- Example: Configuring IGMP Snooping on page 4329
- Example: Configuring Multicast Snooping on page 4343

Example: Configuring IGMP Snooping

- Understanding Multicast Snooping on page 4329
- Understanding IGMP Snooping on page 4330
- IGMP Snooping Interfaces and Forwarding on page 4331
- IGMP Snooping and Proxies on page 4331
- Multicast-Router Interfaces and IGMP Snooping Proxy Mode on page 4332
- Host-Side Interfaces and IGMP Snooping Proxy Mode on page 4333
- IGMP Snooping and Bridge Domains on page 4333
- Configuring IGMP Snooping on page 4333
- Configuring VLAN-Specific IGMP Snooping Parameters on page 4334
- Example: Configuring IGMP Snooping on page 4335
- Configuring IGMP Snooping Trace Operations on page 4341

Understanding Multicast Snooping

Network devices such as routers operate mainly at the packet level, or Layer 3. Other network devices such as bridges or LAN switches operate mainly at the frame level, or Layer 2. Multicasting functions mainly at the packet level, Layer 3, but there is a way to map Layer 3 IP multicast group addresses to Layer 2 MAC multicast group addresses at the frame level.

Routers can handle both Layer 2 and Layer 3 addressing information because the frame and its addresses must be processed to access the encapsulated packet inside. Routers can run Layer 3 multicast protocols such as PIM or IGMP and determine where to forward multicast content or when a host on an interface joins or leaves a group. However, bridges and LAN switches, as Layer 2 devices, are not supposed to have access to the multicast information inside the packets that their frames carry.

Copyright © 2016, Juniper Networks, Inc.

How then are bridges and other Layer 2 devices to determine when a device on an interface joins or leaves a multicast tree, or whether a host on an attached LAN wants to receive the content of a particular multicast group?

The answer is for the Layer 2 device to implement multicast snooping. Multicast snooping is a general term and applies to the process of a Layer 2 device "snooping" at the Layer 3 packet content to determine which actions are taken to process or forward a frame. There are more specific forms of snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to normally "forbidden" Layer 3 (packet) information. Snooping makes multicasting more efficient in these devices.

Understanding IGMP Snooping

Snooping is a general way for Layer 2 devices, such as Juniper Networks MX Series Ethernet Services Routers, to implement a series of procedures to "snoop" at the Layer 3 packet content to determine which actions are to be taken to process or forward a frame. More specific forms of snooping, such as Internet Group Membership Protocol (IGMP) snooping or Protocol Independent Multicast (PIM) snooping, are used with multicast.

Layer 2 devices (LAN switches or bridges) handle multicast packets and the frames that contain them much in the same way the Layer 3 devices (routers) handle broadcasts. So, a Layer 2 switch processes an arriving frame having a multicast destination media access control (MAC) address by forwarding a copy of the packet (frame) onto each of the other network interfaces of the switch that are in a forwarding state.

However, this approach (sending multicast frames everywhere the device can) is not the most efficient use of network bandwidth, particularly for IPTV applications. IGMP snooping functions by "snooping" at the IGMP packets received by the switch interfaces and building a multicast database similar to that a multicast router builds in a Layer 3 network. Using this database, the switch can forward multicast traffic only onto downstream interfaces with interested receivers, and this technique allows more efficient use of network bandwidth.

You configure IGMP snooping for each bridge on the router. A bridge instance without qualified learning has just one learning domain. For a bridge instance with qualified learning, snooping will function separately within each learning domain in the bridge. That is, IGMP snooping and multicast forwarding will proceed independently in each learning domain in the bridge.

This discussion focuses on bridge instances without qualified learning (those forming one learning domain on the device). Therefore, all the interfaces mentioned are logical interfaces of the bridge or VPLS instance.

Several related concepts are important when discussing IGMP snooping:

- Bridge or VPLS instance interfaces are either multicast-router interfaces or host-side interfaces.
- IGMP snooping supports proxy mode or without-proxy mode.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When integrated routing and bridging (IRB) is used, if the router is an IGMP querier, any leave message received on any Layer 2 interface will cause a group-specific query on all Layer 2 interfaces (as a result of this practice, some corresponding reports might be received on all Layer 2 interfaces). However, if some of the Layer 2 interfaces are also router (Layer 3) interfaces, reports and leaves from other Layer 2 interfaces will not be forwarded on those interfaces.

If an IRB interface is used as an outgoing interface in a multicast forwarding cache entry (as determined by the routing process), then the output interface list is expanded into a subset of the Layer 2 interface in the corresponding bridge. The subset is based on the snooped multicast membership information, according to the multicast forwarding cache entry installed by the snooping process for the bridge.

If no snooping is configured, the IRB output interface list is expanded to all Layer 2 interfaces in the bridge.

The Junos OS does not support IGMP snooping in a VPLS configuration on a virtual switch. This configuration is disallowed in the CLI.

IGMP Snooping Interfaces and Forwarding

IGMP snooping divides the device interfaces into multicast-router interfaces and host-side interfaces. A multicast-router interface is an interface in the direction of a multicasting router. An interface on the bridge is considered a multicast-router interface if it meets at least one of the following criteria:

- It is statically configured as a multicast-router interface in the bridge instance.
- IGMP queries are being received on the interface.

All other interfaces that are not multicast-router interfaces are considered host-side interfaces.

Any multicast traffic received on a bridge interface with IGMP snooping configured will be forwarded according to following rules:

- Any IGMP packet is sent to the Routing Engine for snooping processing.
- Other multicast traffic with destination address 224.0.0/24 is flooded onto all other interfaces of the bridge.
- Other multicast traffic is sent to all the multicast-router interfaces but only to those host-side interfaces that have hosts interested in receiving that multicast group.

IGMP Snooping and Proxies

Without a proxy arrangement, IGMP snooping does not generate or introduce queries and reports. It will only “snoop” reports received from all of its interfaces (including multicast-router interfaces) to build its state and group (S,G) database.

Copyright © 2016, Juniper Networks, Inc.

Without a proxy, IGMP messages are processed as follows:

- Query—All general and group-specific IGMP query messages received on a multicast-router interface are forwarded to all other interfaces (both multicast-router interfaces and host-side interfaces) on the bridge.
- Report—IGMP reports received on any interface of the bridge are forwarded toward other multicast-router interfaces. The receiving interface is added as an interface for that group if a multicast routing entry exists for this group. Also, a group timer is set for the group on that interface. If this timer expires (that is, there was no report for this group during the IGMP group timer period), then the interface is removed as an interface for that group.
- Leave—Any IGMP leave message received on any interface of the bridge. The Leave Group message reduces the time it takes for the multicast router to stop forwarding multicast traffic when there are no longer any members in the host group.

Proxy snooping reduces the number of IGMP reports sent toward an IGMP router.

NOTE: With proxy snooping configured, an IGMP router is not able to perform host tracking.

As proxy for its host-side interfaces, IGMP snooping in proxy mode replies to the queries it receives from an IGMP router on a multicast-router interface. On the host-side interfaces, IGMP snooping in proxy mode behaves as an IGMP router and sends general and group-specific queries on those interfaces.

NOTE: Only group-specific queries are generated by IGMP snooping directly. General queries received from the multicast-router interfaces are flooded to host-side interfaces.

All the queries generated by IGMP snooping are sent using 0.0.0.0 as the source address. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as the source address unless there is a configured source address to use.

Proxy mode functions differently on multicast-router interfaces than it does on host-side interfaces.

Multicast-Router Interfaces and IGMP Snooping Proxy Mode

On multicast-router interfaces, in response to IGMP queries, IGMP snooping in proxy mode sends reports containing aggregate information on groups learned on all host-side interfaces of the bridge.

Besides replying to queries, IGMP snooping in proxy mode forwards all queries, reports, and leaves received on a multicast-router interface to other multicast-router interfaces. IGMP snooping keeps the membership information learned on this interface but does not send a group-specific query for leave messages received on this interface. It simply

Copyright © 2016, Juniper Networks, Inc.

times out the groups learned on this interface if there are no reports for the same group within the timer duration.

NOTE: For the hosts on all the multicast-router interfaces, it is the IGMP router, not the IGMP snooping proxy, that generates general and group-specific queries.

Host-Side Interfaces and IGMP Snooping Proxy Mode

No reports are sent on host-side interfaces by IGMP snooping in proxy mode. IGMP snooping processes reports received on these interfaces and sends group-specific queries onto host-side interfaces when it receives a leave message on the interface. Host-side interfaces do not generate periodic general queries, but forwards or floods general queries received from multicast-router interfaces.

If a group is removed from a host-side interface and this was the last host-side interface for that group, a leave is sent to the multicast-router interfaces. If a group report is received on a host-side interface and this was the first host-side interface for that group, a report is sent to all multicast-router interfaces.

IGMP Snooping and Bridge Domains

IGMP snooping on a VLAN is only allowed for the legacy **vlan-id all** case. In other cases, there is a specific bridge domain configuration that determines the VLAN-specific configuration for IGMP snooping.

Configuring IGMP Snooping

To configure Internet Group Management Protocol (IGMP) snooping, include the **igmp-snooping** statement:

```
igmp-snooping {
    immediate-leave;
    interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
            group ip-address {
                source ip-address;
            }
        }
    }
    proxy {
        source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    vlan vlan-id {
```

Copyright © 2016, Juniper Networks, Inc.

```
        immediate-leave;
        interface interface-name {
          group-limit limit;
          host-only-interface;
          immediate-leave;
          multicast-router-interface;
          static {
            group ip-address {
              source ip-address;
            }
          }
        }
        proxy {
          source-address ip-address;
        }
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
      }
    }
```

You can include this statement at the following hierarchy levels:

- **[edit bridge-domains *bridge-domain-name* protocols]**
- **[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]**

By default, IGMP snooping is not enabled. Statements configured at the VLAN level apply only to that particular VLAN.

Configuring VLAN-Specific IGMP Snooping Parameters

All of the IGMP snooping statements configured with the **igmp-snooping** statement, with the exception of the **traceoptions** statement, can be qualified with the same statement at the VLAN level. To configure IGMP snooping parameters at the VLAN level, include the **vlan** statement:

```
vlan vlan-id;
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
    query-response-interval seconds;
    robust-count number;
}
```

You can include this statement at the following hierarchy levels:

- **[edit bridge-domains *bridge-domain-name* protocols igmp-snooping]**
- **[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]**

Example: Configuring IGMP Snooping

This example shows how to configure IGMP snooping. IGMP snooping can reduce unnecessary traffic from IP multicast applications.

- Requirements on page 4335
- Overview and Topology on page 4335
- Configuration on page 4339
- Verification on page 4341

Requirements

This example uses the following hardware components:

- One MX Series router
- One Layer 3 device functioning as a multicast router

Before you begin:

- Configure the interfaces. See the *Interfaces Feature Guide for Security Devices*.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure a multicast protocol. This feature works with the following multicast protocols:
  - DVMRP
  - PIM-DM
  - PIM-SM
  - PIM-SSM

Overview and Topology

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make

Copyright © 2016, Juniper Networks, Inc.

intelligent decisions and to forward multicast traffic to only the intended destination hosts.

This example includes the following statements:

- **proxy**—Enables the Layer 2 device to actively filter IGMP packets to reduce load on the multicast router. Joins and leaves heading upstream to the multicast router are filtered so that the multicast router has a single entry for the group, regardless of how many active listeners have joined the group. When a listener leaves a group but other listeners remain in the group, the leave message is filtered because the multicast router does not need this information. The status of the group remains the same from the router's point of view.
- **immediate-leave**—When only one IGMP host is connected, the **immediate-leave** statement enables the multicast router to immediately remove the group membership from the interface and suppress the sending of any group-specific queries for the multicast group.

  When you configure this feature on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same interface, and one host sends a leave message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

  When IGMP snooping is enabled on a router running IGMP version 3 (IGMPv3) snooping, after the router receives a report with the type BLOCK_OLD_SOURCES, the router suppresses the sending of group-and-source queries but relies on the Junos OS host-tracking mechanism to determine whether or not it removes a particular source group membership from the interface.
- **query-interval**—Enables you to change the number of IGMP messages sent on the subnet by configuring the interval at which the IGMP querier router sends general host-query messages to solicit membership information.

  By default, the query interval is 125 seconds. You can configure any value in the range 1 through 1024 seconds.
- **query-last-member-interval**—Enables you to change the amount of time it takes a device to detect the loss of the last member of a group.

  The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages.

  By default, the last-member query interval is 1 second. You can configure any value in the range 0.1 through 0.9 seconds, and then 1-second intervals from 1 through 1024 seconds.
- **query-response-interval**—Configures how long the router waits to receive a response from its host-query messages.

  By default, the query response interval is 10 seconds. You can configure any value in the range 1 through 1024 seconds. This interval should be less than the interval set in the **query-interval** statement.

Copyright © 2016, Juniper Networks, Inc.

- **robust-count**—Provides fine-tuning to allow for expected packet loss on a subnet. It is basically the number of intervals to wait before timing out a group. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.

  By default, the robust count is 2. You can configure any value in the range 2 through 10 intervals.

- **group-limit**—Configures a limit for the number of multicast groups (or [S,G] channels in IGMPv3) that can join an interface. After this limit is reached, new reports are ignored and all related flows are discarded, not flooded.

  By default, there is no limit to the number of groups that can join an interface. You can configure a limit in the range 0 through a 32-bit number.

- **host-only-interface**—Configure an IGMP snooping interface to be an exclusively host-side interface. On a host-side interface, received IGMP queries are dropped.

  By default, an interface can face either other multicast routers or hosts.

- **multicast-router-interface**—Configures an IGMP snooping interface to be an exclusively router-facing interface.

  By default, an interface can face either other multicast routers or hosts.

- **static**—Configures an IGMP snooping interface with multicast groups statically.

  By default, the router learns about multicast groups on the interface dynamically.

Figure 198 shows networks without IGMP snooping. Suppose host A is an IP multicast sender and hosts B and C are multicast receivers. The router forwards IP multicast traffic only to those segments with registered receivers (hosts B and C). However, the Layer 2 devices flood the traffic to all hosts on all interfaces.

Copyright © 2016, Juniper Networks, Inc.

Figure 198: Networks Without IGMP Snooping Configured

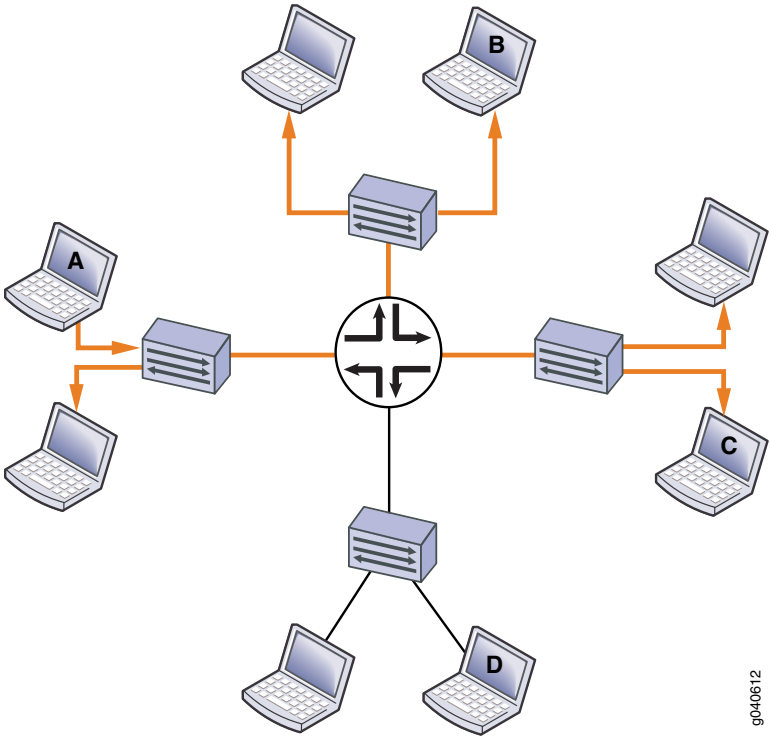

Figure 199 shows the same networks with IGMP snooping configured. The Layer 2 devices forward multicast traffic to registered receivers only.

Copyright © 2016, Juniper Networks, Inc.

Figure 199: Networks with IGMP Snooping Configured

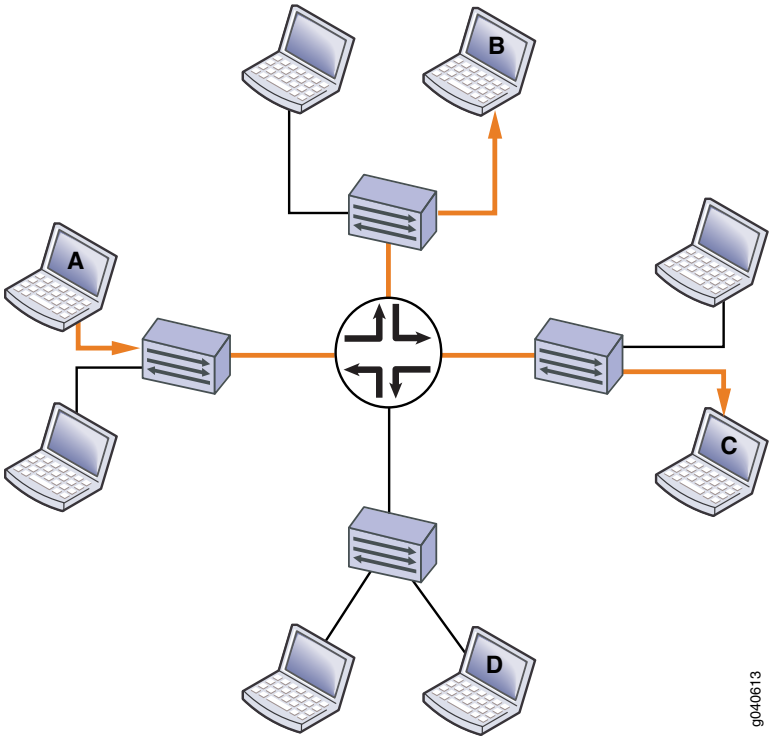

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set bridge-domains domain1 domain-type bridge
set bridge-domains domain1 interface ge-0/0/1.1
set bridge-domains domain1 interface ge-0/0/2.1
set bridge-domains domain1 interface ge-0/0/3.1
set bridge-domains domain1 protocols igmp-snooping query-interval 200
set bridge-domains domain1 protocols igmp-snooping query-response-interval 0.4
set bridge-domains domain1 protocols igmp-snooping query-last-member-interval 0.1
set bridge-domains domain1 protocols igmp-snooping robust-count 4
set bridge-domains domain1 protocols igmp-snooping immediate-leave
set bridge-domains domain1 protocols igmp-snooping proxy
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/1.1
  host-only-interface
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/1.1 group-limit
  50
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/3.1 static group
  225.100.100.100
set bridge-domains domain1 protocols igmp-snooping interface ge-0/0/2.1
  multicast-router-interface
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IGMP snooping:

1. Configure the bridge domain.

```
[edit bridge-domains domain1]
user@host# set domain-type bridge
user@host# set interface ge-0/0/1.1
user@host# set interface ge-0/0/2.1
user@host# set interface ge-0/0/3.1
```

2. Enable IGMP snooping and configure the router to serve as a proxy.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping proxy
```

3. Configure the limit for the number of multicast groups allowed on the **ge-0/0/1.1** interface to 50.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/1.1group-limit 50
```

4. Configure the router to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping immediate-leave
```

5. Statically configure IGMP group membership on a port.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/3.1 static group
  225.100.100.100
```

6. Configure an interface to be an exclusively router-facing interface (to receive multicast traffic).

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/2.1
  multicast-router-interface
```

7. Configure an interface to be an exclusively host-facing interface (to drop IGMP query messages).

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snooping interface ge-0/0/1.1 host-only-interface
```

8. Configure the IGMP message intervals and robustness count.

```
[edit bridge-domains domain1]
user@host# set protocols igmp-snoopingrobust-count 4
user@host# set protocols igmp-snooping query-last-member-interval 0.1
user@host# set protocols igmp-snooping query-interval 200
user@host# set protocols igmp-snooping query-response-interval 0.4
```


Copyright © 2016, Juniper Networks, Inc.

9. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results Confirm your configuration by entering the **show bridge-domains** command.

```
user@host# show bridge-domains
domain1 {
  domain-type bridge;
  interface ge-0/0/1.1;
  interface ge-0/0/2.1;
  interface ge-0/0/3.1;
  protocols {
    igmp-snooping {
      query-interval 200;
      query-response-interval 0.4;
      query-last-member-interval 0.1;
      robust-count 4;
      immediate-leave;
      proxy;
      interface ge-0/0/1.1 {
        host-only-interface;
        group-limit 50;
      }
      interface ge-0/0/3.1 {
        static {
          group 225.100.100.100;
        }
      }
      interface ge-0/0/2.1 {
        multicast-router-interface;
      }
    }
  }
}
```

Verification

To verify the configuration, run the following commands:

- show igmp snooping interface
- show igmp snooping membership
- show igmp snooping statistics

Configuring IGMP Snooping Trace Operations

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy

Copyright © 2016, Juniper Networks, Inc.

actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.
client-notification	Trace notifications.
general	Trace general flow.
group	Trace group operations.
host-notification	Trace host notifications.
leave	Trace leave group messages (IGMPv2 only).
normal	Trace normal events.
packets	Trace all IGMP packets.
policy	Trace policy processing.
query	Trace IGMP membership query messages.
report	Trace membership report messages.
route	Trace routing information.
state	Trace state transitions.
task	Trace routing protocol task processing.
timer	Trace timer processing.

You can configure tracing operations for IGMP snooping globally or in a routing instance. The following example shows the global configuration.

To configure tracing operations for IGMP snooping:

1. Configure the filename for the trace file.

   [edit bridge-domains domain1 protocols igmp-snooping traceoptions]
   user@host# **set file igmp-snoop-trace**

2. (Optional) Configure the maximum number of trace files.

   [edit bridge-domains domain1 protocols igmp-snooping traceoptions]
   user@host# **set file files 5**

3. (Optional) Configure the maximum size of each trace file.

   [edit bridge-domains domain1 protocols igmp-snooping traceoptions]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set file size 1m
```

4. (Optional) Enable unrestricted file access.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set file world-readable
```

5. Configure tracing flags. Suppose you are troubleshooting issues with a policy related to received packets on a particular logical interface with an IP address of 192.168.0.1. The following example shows how to flag all policy events for received packets associated with the IP address.

```
[edit bridge-domains domain1 protocols igmp-snooping traceoptions]
user@host# set flag policy receive | match 192.168.0.1
```

6. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-snoop-trace
```

Related Documentation

- Understanding Multicast Snooping on page 4329

Example: Configuring Multicast Snooping

- Understanding Multicast Snooping on page 4343
- Understanding Multicast Snooping and VPLS Root Protection on page 4344
- Configuring Multicast Snooping on page 4344
- Example: Configuring Multicast Snooping on page 4345
- Enabling Bulk Updates for Multicast Snooping on page 4350
- Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces on page 4351

Understanding Multicast Snooping

Network devices such as routers operate mainly at the packet level, or Layer 3. Other network devices such as bridges or LAN switches operate mainly at the frame level, or Layer 2. Multicasting functions mainly at the packet level, Layer 3, but there is a way to map Layer 3 IP multicast group addresses to Layer 2 MAC multicast group addresses at the frame level.

Routers can handle both Layer 2 and Layer 3 addressing information because the frame and its addresses must be processed to access the encapsulated packet inside. Routers can run Layer 3 multicast protocols such as PIM or IGMP and determine where to forward multicast content or when a host on an interface joins or leaves a group. However, bridges and LAN switches, as Layer 2 devices, are not supposed to have access to the multicast information inside the packets that their frames carry.

How then are bridges and other Layer 2 devices to determine when a device on an interface joins or leaves a multicast tree, or whether a host on an attached LAN wants to receive the content of a particular multicast group?

Copyright © 2016, Juniper Networks, Inc.

The answer is for the Layer 2 device to implement multicast snooping. Multicast snooping is a general term and applies to the process of a Layer 2 device "snooping" at the Layer 3 packet content to determine which actions are taken to process or forward a frame. There are more specific forms of snooping, such as IGMP snooping or PIM snooping. In all cases, snooping involves a device configured to function at Layer 2 having access to normally "forbidden" Layer 3 (packet) information. Snooping makes multicasting more efficient in these devices.

Understanding Multicast Snooping and VPLS Root Protection

Snooping occurs when a Layer 2 protocol such as a spanning-tree protocol is aware of the operational details of a Layer 3 protocol such as the Internet Group Management Protocol (IGMP) or other multicast protocol. Snooping is necessary when Layer 2 devices such as VLAN switches must be aware of Layer 3 information such as the media access control (MAC) addresses of members of a multicast group.

VPLS root protection is a spanning-tree protocol process in which only one interface in a multihomed environment is actively forwarding spanning-tree protocol frames. This protects the root of the spanning tree against bridging loops, but also prevents both devices in the multihomed topology from snooped information, such as IGMP membership reports.

For example, consider a collection of multicast-capable hosts connected to two customer edge (CE) routers (CE1 and CE2) which are connected to each other (a CE1–CE2 link is configured) and multihomed to two provider edge (PE) routers (PE1 and PE2, respectively). The active PE only receives forwarded spanning-tree protocol information on the active PE-CE link, due to root protection operation. As long as the CE1–CE2 link is operational, this is not a problem. However, if the link between CE1 and CE2 fails, and the other PE becomes the active spanning-tree protocol link, no multicast snooping information is available on the new active PE. The new active PE will not forward multicast traffic to the CE and the hosts serviced by this CE router.

The service outage is corrected once the hosts send new group membership IGMP reports to the CE routers. However, the service outage can be avoided if multicast snooping information is available to both PEs in spite of normal spanning-tree protocol root protection operation.

NOTE: **You can configure multicast snooping to ignore messages about spanning tree topology changes for the virtual-switch routing-instance type only.**

Configuring Multicast Snooping

To configure the general multicast snooping parameters for MX Series routers, include the **multicast-snooping-options** statement:

```
multicast-snooping-options {
  flood-groups [ ip-addresses ];
  forwarding-cache {
    threshold suppress value <reuse value>;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    graceful-restart <restart-duration seconds>;
    ignore-stp-topology-change;
    multichassis-lag-replicate-state;
    nexthop-hold-time milliseconds;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
```

You can include this statement at the following hierarchy levels:

- **[edit bridge-domains *bridge-domain-name*]**
- **[edit routing-instances *routing-instance-name*]**
- **[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]**

By default, multicast snooping is disabled. You can enable multicast snooping in VPLS or virtual switch instance types in the instance hierarchy or in one or more bridge domains.

If there are multiple bridge domains configured under a VPLS or virtual switch instance, the multicast snooping options configured at the instance level apply to all the bridge domains. Multicast snooping options configured at the bridge domain level only apply to that particular bridge domain. The options configured at the bridge domain take precedence over the options configured at the instance level.

NOTE: The **ignore-stp-topology-change** statement is supported for the **virtual-switch** routing instance type only and is not supported under the **[edit logical-systems]** hierarchy.

NOTE: The **nexthop-hold-time** statement is supported only at the **[edit routing-instances *routing-instance-name*]** hierarchy, and only for an instance type of **virtual-switch** or **vpls**.

Example: Configuring Multicast Snooping

This example shows how to configure multicast snooping in a bridge or VPLS routing-instance scenario.

- Requirements on page 4346
- Overview and Topology on page 4346
- Configuration on page 4348
- Verification on page 4350

Copyright © 2016, Juniper Networks, Inc.

Requirements

This example uses the following hardware components:

- One MX Series router
- One Layer 3 device functioning as a multicast router

Before you begin:

- Configure the interfaces.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure a multicast protocol. This feature works with the following multicast protocols:
  - DVMRP
  - PIM-DM
  - PIM-SM
  - PIM-SSM

Overview and Topology

IGMP snooping prevents Layer 2 devices from indiscriminately flooding multicast traffic out all interfaces. The settings that you configure for multicast snooping help manage the behavior of IGMP snooping.

You can configure multicast snooping options on the default master instance and on individual bridge or VPLS instances. The default master instance configuration is global and applies to all individual bridge or VPLS instances in the logical router. The configuration for the individual instances overrides the global configuration.

This example includes the following statements:

- **flood-groups**—Enables you to list multicast group addresses for which traffic must be flooded. This setting if useful for making sure that IGMP snooping does not prevent necessary multicast flooding. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms. For example, OSPF uses 224.0.0.5 for all OSPF routers.
- **forwarding-cache**—Specifies how forwarding entries are aged out and how the number of entries is controlled.

  You can configure threshold values on the forwarding cache to suppress (suspend) snooping when the cache entries reach a certain maximum and reuse the cache when the number falls to another threshold value. By default, no threshold values are enabled on the router.

Copyright © 2016, Juniper Networks, Inc.

The suppress threshold suppresses new multicast forwarding cache entries. An optional reuse threshold specifies the point at which the router begins to create new multicast forwarding cache entries. The range for both thresholds is from 1 through 200,000. If configured, the reuse value must be less than the suppression value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

- **graceful-restart**—Configures the time after which routes learned before a restart are replaced with routes relearned. If graceful restart for multicast snooping is disabled, snooping information is lost after a Routing Engine restart.

  By default, the graceful restart duration is 180 seconds (3 minutes). You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

- **ignore-stp-topology-change**—Configures the MX Series router to ignore messages about the spanning-tree topology state change.

  By default the IGMP snooping process on an MX Series router detects interface state changes made by any of the spanning tree protocols (STPs).

  In a VPLS multihoming environment where two PE routers are connected to two interconnected CE routers and STP root protection is enabled on the PE routers, one of the PE router interfaces is in forwarding state and the other is in blocking state.

  If the link interconnecting the two CE routers fails, the PE router interface in blocking state transitions to the forwarding state.

  The PE router interface does not wait to receive membership reports in response to the next general or group-specific query. Instead, the IGMP snooping process sends a general query message toward the CE router. The hosts connected to the CE router reply with reports for all groups they are interested in.

  When the link interconnecting the two CE routers is restored, the original spanning-tree state on both PE routers is restored. The forwarding PE receives a spanning-tree topology change message and sends a general query message toward the CE router to immediately reconstruct the group membership state.

NOTE: **The ignore-stp-topology-change statement is supported for the virtual-switch routing instance type only.**

Figure 200 shows a VPLS multihoming topology in which a customer network has two CE devices with a link between them. Each CE is connected to one PE.

Copyright © 2016, Juniper Networks, Inc.

Figure 200: VPLS Multihoming Topology

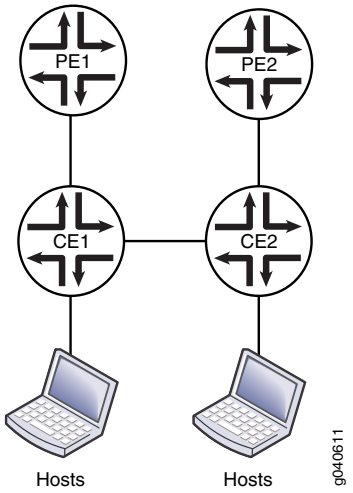

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set bridge-domains domain1 multicast-snooping-options forwarding-cache threshold
  suppress 100
set bridge-domains domain1 multicast-snooping-options forwarding-cache threshold
  reuse 50
set bridge-domains domain1 multicast-snooping-options graceful-restart restart-duration
  120
set routing-instances ce1 instance-type virtual-switch
set routing-instances ce1 bridge-domains domain1 domain-type bridge
set routing-instances ce1 bridge-domains domain1 vlan-id 100
set routing-instances ce1 bridge-domains domain1 interface ge-0/3/9.0
set routing-instances ce1 bridge-domains domain1 interface ge-0/0/6.0
set routing-instances ce1 bridge-domains domain1 multicast-snooping-options
  flood-groups 224.0.0.5
set routing-instances ce1 bridge-domains domain1 multicast-snooping-options
  ignore-stp-topology-change
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure IGMP snooping:

1. Configure multicast snooping settings in the master routing instance.

```
[edit bridge-domains domain1]
user@host# set multicast-snooping-options forwarding-cache threshold suppress
  100 reuse 50
user@host# set multicast-snooping-options graceful-restart 120
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure the routing instance.

```
[edit routing-instances ce1]
user@host# set instance-type virtual-switch
```

3. Configure the bridge domain in the routing instance.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set domain-type bridge
user@host# set interface ge-0/0/6.0
user@host# set interface ge-0/3/9.0
user@host# set vlan-id 100
```

4. Configure flood groups.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set multicast-snooping-options flood-groups 224.0.0.5
```

5. Configure the router to ignore messages about spanning-tree topology state changes.

```
[edit routing-instances ce1 bridge-domains domain1]
user@host# set multicast-snooping-options ignore-stp-topology-change
```

6. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results Confirm your configuration by entering the **show bridge-domains** and **show routing-instances** commands.

```
user@host# show bridge-domains
domain1 {
  multicast-snooping-options {
    forwarding-cache {
      threshold {
        suppress 100;
        reuse 50;
      }
    }
    graceful-restart {
      restart-duration 120;
    }
  }
}
```

```
user@host# show routing-instances
ce1 {
  instance-type virtual-switch;
  bridge-domains {
    domain1 {
      domain-type bridge;
      vlan-id 100;
      interface ge-0/3/9.0; ## 'ge-0/3/9.0' is not defined
      interface ge-0/0/6.0; ## 'ge-0/0/6.0' is not defined
      multicast-snooping-options {
        flood-groups 224.0.0.5;
        ignore-stp-topology-change;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

Verification

To verify the configuration, run the following commands:

- show igmp snooping interface
- show igmp snooping membership
- show igmp snooping statistics
- show multicast snooping route
- show multicast snooping statistics
- show route table

Enabling Bulk Updates for Multicast Snooping

Whenever an individual interface joins or leaves a multicast group, a new next hop entry is installed in the routing table and the forwarding table. You can use the **nexthop-hold-time** statement to specify a time, from 1 through 1000 milliseconds (ms), during which outgoing interface changes are accumulated and then updated in bulk to the routing table and forwarding table. Bulk updating reduces the processing time and memory overhead required to process join and leave messages. This is useful for applications such as Internet Potocol television (IPTV), in which users changing channels can create thousands of interfaces joining or leaving a group in a short period. In IPTV scenarios, typically there is a relatively small and controlled number of streams and a high number of outgoing interfaces. Using bulk updates can reduce the join delay.

In this example, you configure a hold-time of 20 milliseconds for **instance-type virtual-switch**, using the **nexthop-hold-time** statement:

1. Enable the **nexthop-hold-time** statement by configuring it under **multicast-snooping-options**, using 20 milliseconds for the time value.

   ```
   [edit routing-instances vs]
   multicast-snooping-options {
     nexthop-hold-time 20;
   }
   ```

2. Use the **show multicast snooping route** command to verify that the bulk updates feature is turned on.

   ```
   user@host> show multicast snooping route instance vs
   Nexthop Bulking: ON
   Family: INET
   Group: 224.0.0.0
   ```

You can include the **nexthop-hold-time** statement only for routing-instance types of **virtual-switch** or **vpls** at the following hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

- **[edit routing-instances** *routing-instance-name* **multicast-snooping-options]**

If the **nexthop-hold-time** statement is deleted from the router configuration, bulk updates are disabled.

Enabling Multicast Snooping for Multichassis Link Aggregation Group Interfaces

Include the **multichassis-lag-replicate-state** statement at the **[edit multicast-snooping-options]** hierarchy level to enable IGMP snooping and state replication for multichassis link aggregation group (MC-LAG) interfaces.

```
[edit]
multicast-snooping-options {
  multichassis-lag-replicate-state;
}
```

Replicating join and leave messages between links of a dual-link MC-LAG interface enables faster recovery of membership information for MC-LAG interfaces that experience service interruption.

Without state replication, if a dual-link MC-LAG interface experiences a service interruption (for example, if an active link switches to standby), the membership information for the interface is recovered by generating an IGMP query to the network. This method can take from 1 through 10 seconds to complete, which might be too long for some applications.

When state replication is provided for MC-LAG interfaces, IGMP join or leave messages received on an MC-LAG device are replicated from the active MC-LAG link to the standby link through an Interchassis Communication Protocol (ICCP) connection. The standby link processes the messages as if they were received from the corresponding active MC-LAG link, except it does not add itself as a next hop and it does not flood the message to the network. After a failover, the multicast membership status of the link can be recovered within a few seconds or less by retrieving the replicated messages.

This example enables state replication for MC-LAG interfaces in a bridge domain named bridge1:

1. Enable state replication for MC-LAG interfaces.

   ```
   user@host# set multicast-snooping-options multicast-lag-replicate-state
   ```

   After you commit the configuration, multicast snooping automatically identifies the active link during initialization or after failover, and replicates data between the active and standby links without administrator intervention.

2. Use the **show igmp snooping interface** command to display the state for MC-LAG interfaces.

   ```
   user@host> show igmp snooping interface

   Instance: bridge-domain bridge1
   Learning-Domain: default
   Interface: ae0.1
     State: Up Groups: 1
     mc-lag state: standby
     Immediate leave: Off
   ```

Copyright © 2016, Juniper Networks, Inc.

```
    Router interface: no
Interface: ge-0/1/3.100
    State: Up Groups: 1
    Immediate leave: Off
    Router interface: no
Interface: ae1.2
    State: Up Groups: 1
    mc-lag state: standby
    Immediate leave: Off
    Router interface: no
```

NOTE: You can use the **show igmp snooping membership** command to display group membership information for the links of MC-LAG interfaces.

If you delete the **multicast-lag-replicate-state** statement or the configuration of IGMP snooping, replication between MC-LAG links stops within the hierarchy level from which the configuration was deleted. Then, multicast membership is recovered as needed by generating standard IGMP queries over the network.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 195

Configuring Multicast Routing Options

- Examples: Configuring Bandwidth Management on page 4353
- Examples: Configuring the Multicast Forwarding Cache on page 4374
- Example: Configuring Ingress PE Redundancy on page 4381

Examples: Configuring Bandwidth Management

- Understanding Bandwidth Management for Multicast on page 4353
- Bandwidth Management and PIM Graceful Restart on page 4354
- Bandwidth Management and Source Redundancy on page 4354
- Logical Systems and Bandwidth Oversubscription on page 4354
- Example: Defining Interface Bandwidth Maximums on page 4355
- Example: Configuring Multicast with Subscriber VLANs on page 4358
- Configuring Multicast Routing over IP Demux Interfaces on page 4371
- Classifying Packets by Egress Interface on page 4372

Understanding Bandwidth Management for Multicast

Bandwidth management enables you to control the multicast flows that leave a multicast interface. This control enables you to better manage your multicast traffic and reduce or eliminate the chances of interface oversubscription or congestion.

Bandwidth management ensures that multicast traffic oversubscription does not occur on an interface. When managing multicast bandwidth, you define the maximum amount of multicast bandwidth that an individual interface can use as well as the bandwidth individual multicast flows use.

For example, the routing software cannot add a flow to an interface if doing so exceeds the allowed bandwidth for that interface. Under these circumstances, the interface is rejected. This rejection, however, does not prevent a multicast protocol (for example, PIM) from sending a join message upstream. Traffic continues to arrive on the router, even though the router is not sending the flow from the expected outgoing interfaces.

You can configure the flow bandwidth statically by specifying a bandwidth value for the flow in bits per second, or you can enable the flow bandwidth to be measured and adaptively changed. When using the adaptive bandwidth option, the routing software

Copyright © 2016, Juniper Networks, Inc.

queries the statistics for the flows to be measured at 5-second intervals and calculates the bandwidth based on the queries. The routing software uses the maximum value measured within the last minute (that is, the last 12 measuring points) as the flow bandwidth.

For more information, see the following sections:

- Bandwidth Management and PIM Graceful Restart on page 4354
- Bandwidth Management and Source Redundancy on page 4354
- Logical Systems and Bandwidth Oversubscription on page 4354

Bandwidth Management and PIM Graceful Restart

When using PIM graceful restart, after the routing process restarts on the Routing Engine, previously admitted interfaces are always readmitted and the available bandwidth is adjusted on the interfaces. When using the adaptive bandwidth option, the bandwidth measurement is initially based on the configured or default starting bandwidth, which might be inaccurate during the first minute. This means that new flows might be incorrectly rejected or admitted temporarily. You can correct this problem by issuing the **clear multicast bandwidth-admission** operational command.

If PIM graceful restart is not configured, after the routing process restarts, previously admitted or rejected interfaces might be rejected or admitted in an unpredictable manner.

Bandwidth Management and Source Redundancy

When using source redundancy, multiple sources (for example, s1 and s2) might exist for the same destination group (g). However, only one of the sources can actively transmit at any time. In this case, multiple forwarding entries—(s1,g) and (s2,g)—are created after each goes through the admission process.

With redundant sources, unlike unrelated entries, an OIF that is already admitted for one entry—for example, (s1,g)—is automatically admitted for other redundancy entries—for example, (s2,g). The remaining bandwidth on the interface is deducted each time an outbound interface is added, even though only one sender actively transmits. By measuring bandwidth, the bandwidth deducted for the inactive entries is credited back when the router detects no traffic is being transmitted.

For more information about defining redundant sources, see "Example: Configuring a Multicast Flow Map" on page 4377.

Logical Systems and Bandwidth Oversubscription

You can manage bandwidth at both the physical and logical interface level. However, if more than one logical system shares the same physical interface, the interface might become oversubscribed. Oversubscription occurs if the total bandwidth of all separately configured maximum bandwidth values for the interfaces on each logical system exceeds the bandwidth of the physical interface.

When displaying interface bandwidth information, a negative available bandwidth value indicates oversubscription on the interface.

Copyright © 2016, Juniper Networks, Inc.

Interface bandwidth can become oversubscribed when the configured maximum bandwidth decreases or when some flow bandwidths increase because of a configuration change or an actual increase in the traffic rate.

Interface bandwidth can become available again if one of the following occurs:

- The configured maximum bandwidth increases.
- Some flows are no longer transmitted from interfaces, and bandwidth reserves for them are now available to other flows.
- Some flow bandwidths decrease because of a configuration change or an actual decrease in the traffic rate.

Interfaces that are rejected for a flow because of insufficient bandwidth are not automatically readmitted, even when bandwidth becomes available again. Rejected interfaces have an opportunity to be readmitted when one of the following occurs:

- The multicast routing protocol updates the forwarding entry for the flow after receiving a join, leave, or prune message or after a topology change occurs.
- The multicast routing protocol updates the forwarding entry for the flow due to configuration changes.
- You manually reapply bandwidth management to a specific flow or to all flows using the **clear multicast bandwidth-admission** operational command.

In addition, even if previously available bandwidth is no longer available, already admitted interfaces are not removed until one of the following occurs:

- The multicast routing protocol explicitly removes the interfaces after receiving a leave or prune message or after a topology change occurs.
- You manually reapply bandwidth management to a specific flow or to all flows using the **clear multicast bandwidth-admission** operational command.

Example: Defining Interface Bandwidth Maximums

This example shows you how to configure the maximum bandwidth for a physical or logical interface.

- Requirements on page 4355
- Overview on page 4356
- Configuration on page 4356
- Verification on page 4358

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.

Copyright © 2016, Juniper Networks, Inc.

- Configure a multicast protocol. This feature works with the following multicast protocols:
  - DVMRP
  - PIM-DM
  - PIM-SM
  - PIM-SSM

Overview

The maximum bandwidth setting applies admission control either against the configured interface bandwidth or against the native speed of the underlying interface (when there is no configured bandwidth for the interface).

If you configure several logical interfaces (for example, to support VLANs or PVCs) on the same underlying physical interface, and no bandwidth is configured for the logical interfaces, it is assumed that the logical interfaces all have the same bandwidth as the underlying interface. This can cause oversubscription. To prevent oversubscription, configure bandwidth for the logical interfaces, or configure admission control at the physical interface level.

You only need to define the maximum bandwidth for an interface on which you want to apply bandwidth management. An interface that does not have a defined maximum bandwidth transmits all multicast flows as determined by the multicast protocol that is running on the interface (for example, PIM).

If you specify **maximum-bandwidth** without including a bits-per-second value, admission control is enabled based on the bandwidth configured for the interface. In the following example, admission control is enabled for logical interface unit **200**, and the maximum bandwidth is 20 Mbps. If the bandwidth is not configured on the interface, the maximum bandwidth is the link speed.

```
routing-options {
   multicast {
      interface fe-0/2/0.200 {
         maximum-bandwidth;
      }
interfaces {
   fe-0/2/0 {
      unit 200 {
         bandwidth 20m;
      }
   }
}
```

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces fe-0/2/0 unit 200 bandwidth 20m
set routing-options multicast interface fe-0/2/0.200 maximum-bandwidth
set routing-options multicast interface fe-0/2/1 maximum-bandwidth 60m
set routing-options multicast interface fe-0/2/1.200 maximum-bandwidth 10m
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a bandwidth maximum:

1. Configure the a logical interface bandwidth.

   ```
   [edit interfaces]
   user@host# set fe-0/2/0 unit 200 bandwidth 20m
   ```

2. Enable admission control on the logical interface.

   ```
   [edit routing-options]
   user@host# set multicast interface fe-0/2/0.200 maximum-bandwidth
   ```

3. On a physical interface, enable admission control and set the maximum bandwidth to 60 Mbps.

   ```
   [edit routing-options]
   user@host# set multicast interface fe-0/2/1 maximum-bandwidth 60m
   ```

4. For a logical interface on the same physical interface shown in Step 3, set a smaller maximum bandwidth.

   ```
   [edit routing-options]
   user@host# set multicast interface fe-0/2/1.200 maximum-bandwidth 10m
   ```

Results

Confirm your configuration by entering the **show interfaces** and **show routing-options** commands.

```
user@host# show interfaces
fe-0/2/0 {
  unit 200 {
    bandwidth 20m;
  }
}
user@host# show routing-options
multicast {
  interface fe-0/2/0.200 {
    maximum-bandwidth;
  }
  interface fe-0/2/1 {
    maximum-bandwidth 60m;
  }
  interface fe-0/2/1.200 {
    maximum-bandwidth 10m;
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration, run the show multicast interface command.

Example: Configuring Multicast with Subscriber VLANs

This example shows how to configure an MX Series router to function as a broadband service router (BSR).

- Requirements on page 4358
- Overview and Topology on page 4358
- Configuration on page 4362
- Verification on page 4370

Requirements

This example uses the following hardware components:

- One MX Series router or EX Series switch with a PIC that supports traffic control profile queuing
- One DSLAM

Before you begin:

- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure PIM and IGMP or MLD on the interfaces.

Overview and Topology

When multiple BSR interfaces receive IGMP and MLD join and leave requests for the same multicast stream, the BSR sends a copy of the multicast stream on each interface. Both the multicast control packets (IGMP and MLD) and the multicast data packets flow on the same BSR interface, along with the unicast data. Because all per-customer traffic has its own interface on the BSR, per-customer accounting, call admission control (CAC), and quality-of-service (QoS) adjustment are supported. The QoS bandwidth used by multicast reduces the unicast bandwidth.

Multiple interfaces on the BSR might connect to a shared device (for example, a DSLAM). The BSR sends the same multicast stream multiple times to the shared device, thus wasting bandwidth. It is more efficient to send the multicast stream one time to the DSLAM and replicate the multicast streams in the DSLAM. There are two approaches that you can use.

The first approach is to continue to send unicast data on the per-customer interfaces, but have the DSLAM route all the per-customer IGMP and MLD join and leave requests to the BSR on a single dedicated interface (a multicast VLAN). The DSLAM receives the multicast streams from the BSR on the dedicated interface with no unnecessary replication and performs the necessary replication to the customers. Because all multicast control and data packets use only one interface, only one copy of a stream is sent even

Copyright © 2016, Juniper Networks, Inc.

if there are multiple requests. This approach is called reverse outgoing interface (OIF) mapping. Reverse OIF mapping enables the BSR to propagate the multicast state of the shared interface to the customer interfaces, which enables per-customer accounting and QoS adjustment to work. When a customer changes the TV channel, the router gateway (RG) sends an IGMP or MLD join and leave messages to the DSLAM. The DSLAM transparently passes the request to the BSR through the multicast VLAN. The BSR maps the IGMP or MLD request to one of the subscriber VLANs based on the IP source address or the source MAC address. When the subscriber VLAN is found, QoS adjustment and accounting are perfomed on that VLAN or interface.

The second approach is for the DSLAM to continue to send unicast data and all the per-customer IGMP and MLD join and leave requests to the BSR on the individual customer interfaces, but to have the multicast streams arrive on a single dedicated interface. If multiple customers request the same multicast stream, the BSR sends one copy of the data on the dedicated interface. The DSLAM receives the multicast streams from the BSR on the dedicated interface and performs the necessary replication to the customers. Because the multicast control packets use many customer interfaces, configuration on the BSR must specify how to map each customer's multicast data packets to the single dedicated output interface. QoS adjustment is supported on the customer interfaces. CAC is supported on the shared interface. This second approach is called multicast OIF mapping.

OIF mapping and reverse OIF mapping are not supported on the same customer interface or shared interface. This example shows how to configure the two different approaches. Both approaches support QoS adjustment, and both approaches support MLD/IPv6. The reverse OIF mapping example focuses on IGMP/IPv4 and enables QoS adjustment. The OIF mapping example focuses on MLD/IPv6 and disables QoS adjustment.

The first approach (reverse OIF mapping) includes the following statements:

- **flow-map**—Defines a flow map that controls the bandwidth for each flow.
- **maximum-bandwidth**—Enables CAC.
- **reverse-oif-mapping**—Enables the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD join or leave request that it receives over the multicast VLAN.

  After the subscriber VLAN is identified, the routing device immediately adjusts the QoS (in this case, the bandwidth) on that VLAN based on the addition or removal of a subscriber.

  The routing device uses IGMP and MLD join or leave reports to obtain the subscriber VLAN information. This means that the connecting equipment (for example, the DSLAM) must forward all IGMP and MLD reports to the routing device for this feature to function properly. Using report suppression or an IGMP proxy can result in reverse OIF mapping not working properly.
- **subscriber-leave-timer**—Introduces a delay to the QoS update. After receiving an IGMP or MLD leave request, this statement defines a time delay (between 1 and 30 seconds) that the routing device waits before updating the QoS for the remaining subscriber interfaces. You might use this delay to decrease how often the routing device adjusts

Copyright © 2016, Juniper Networks, Inc.

the overall QoS bandwidth on the VLAN when a subscriber sends rapid leave and join messages (for example, when changing channels in an IPTV network).

- **traffic-control-profile**—Configures a shaping rate on the logical interface. The configured shaping rate must be configured as an absolute value, not as a percentage.

The second approach (OIF mapping) includes the following statements:

- **map-to-interface**—In a policy statement, enables you to build the OIF map.

  The OIF map is a routing policy statement that can contain multiple terms. When creating OIF maps, keep the following in mind:

  - If you specify a physical interface (for example, **ge-0/0/0**), a ".0" is appended to the interface to create a logical interface (for example, **ge-0/0/0.0**).
  - Configure a routing policy for each logical system. You cannot configure routing policies dynamically.
  - The interface must also have IGMP, MLD, or PIM configured.
  - You cannot map to a mapped interface.
  - We recommend that you configure policy statements for IGMP and MLD separately.
  - Specify either a logical interface or the keyword **self**. The **self** keyword specifies that multicast data packets be sent on the same interface as the control packets and that no mapping occur. If no term matches, then no multicast data packets are sent.

- **no-qos-adjust**—Disables QoS adjustment.

  QoS adjustment decreases the available bandwidth on the client interface by the amount of bandwidth consumed by the multicast streams that are mapped from the client interface to the shared interface. This action always occurs unless it is explicitly disabled.

  If you disable QoS adjustment, available bandwidth is not reduced on the customer interface when multicast streams are added to the shared interface.

  NOTE: You can dynamically disable QoS adjustment for IGMP and MLD interfaces using dynamic profiles.

- **oif-map**—Associate a map with an IGMP or MLD interface. The OIF map is then applied to all IGMP or MLD requests received on the configured interface. In this example, subscriber VLANs 1 and 2 have MLD configured, and each VLAN points to an OIF map that directs some traffic to **ge-2/3/9.4000**, some traffic to **ge-2/3/9.4001**, and some traffic to **self**.

  NOTE: You can dynamically associate OIF maps with IGMP interfaces using dynamic profiles.

- **passive**—Defines either IGMP or MLD to use passive mode.

Copyright © 2016, Juniper Networks, Inc.

The OIF map interface should not typically pass IGMP or MLD control traffic and should be configured as passive. However, the OIF map implementation does support running IGMP or MLD on an interface (control and data) in addition to mapping data streams to the same interface. In this case, you should configure IGMP or MLD normally (that is, not in passive mode) on the mapped interface. In this example, the OIF map interfaces (**ge-2/3/9.4000** and **ge-2/3/9.4001**) are configured as MLD passive.

By default, specifying the **passive** statement means that no general queries, group-specific queries, or group-source-specific queries are sent over the interface and that all received control traffic is ignored by the interface. However, you can selectively activate up to two out of the three available options for the **passive** statement while keeping the other functions passive (inactive).

These options include the following:

- **send-general-query**—When specified, the interface sends general queries.
- **send-group-query**—When specified, the interface sends group-specific and group-source-specific queries.
- **allow-receive**—When specified, the interface receives control traffic.

Figure 201 shows the scenario.

In both approaches, if multiple customers request the same multicast stream, the BSR sends one copy of the stream on the shared multicast VLAN interface. The DSLAM receives the multicast stream from the BSR on the shared interface and performs the necessary replication to the customers.

In the first approach (reverse OIF mapping), the DSLAM uses the per-customer subscriber VLANs for unicast data only. IGMP and MLD join and leave requests are sent on the multicast VLAN.

In the second approach (OIF mapping), the DSLAM uses the per-customer subscriber VLANs for unicast data and for IGMP and MLD join and leave requests. The multicast VLAN is used only for multicast streams, not for join and leave requests.

Copyright © 2016, Juniper Networks, Inc.

Figure 201: Multicast with Subscriber VLANs

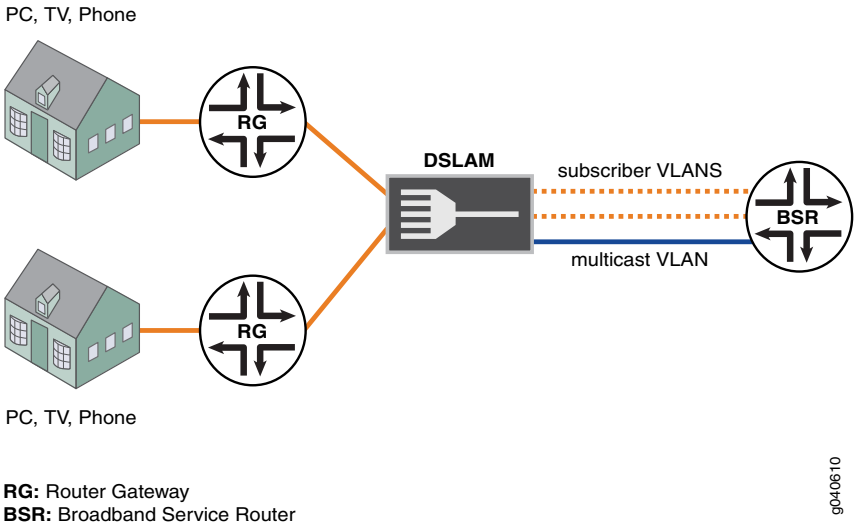

RG: Router Gateway
BSR: Broadband Service Router

Configuration

Configuring a Reverse OIF Map

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode. .

```
set class-of-service traffic-control-profiles tcp-ifl shaping-rate 20m
set class-of-service interfaces ge-2/2/0 shaping-rate 240m
set class-of-service interfaces ge-2/2/0 unit 50 output-traffic-control-profile tcp-ifl
set class-of-service interfaces ge-2/2/0 unit 51 output-traffic-control-profile tcp-ifl
set interfaces ge-2/0/0 unit 0 family inet address 30.0.0.2/24
set interfaces ge-2/2/0 hierarchical-scheduler
set interfaces ge-2/2/0 vlan-tagging
set interfaces ge-2/2/0 unit 10 vlan-id 10
set interfaces ge-2/2/0 unit 10 family inet address 40.0.0.2/24
set interfaces ge-2/2/0 unit 50 vlan-id 50
set interfaces ge-2/2/0 unit 50 family inet address 50.0.0.2/24
set interfaces ge-2/2/0 unit 51 vlan-id 51
set interfaces ge-2/2/0 unit 51 family inet address 50.0.1.2/24
set policy-options policy-statement all-mcast-groups from source-address-filter
  30.0.0.0/8 orlonger
set policy-options policy-statement all-mcast-groups then accept
set protocols igmp interface all
set protocols igmp interface fxp0.0 disable
set protocols pim rp local address 20.0.0.2
set protocols pim interface all
set protocols pim interface fxp0.0 disable
set protocols pim interface ge-2/2/0.10 disable
set routing-options multicast flow-map map1 policy all-mcast-groups
set routing-options multicast flow-map map1 bandwidth 10m
set routing-options multicast flow-map map1 bandwidth adaptive
set routing-options multicast interface ge-2/2/0.10 maximum-bandwidth 500m
```


Copyright © 2016, Juniper Networks, Inc.

```
set routing-options multicast interface ge-2/2/0.10 reverse-oif-mapping
set routing-options multicast interface ge-2/2/0.10 subscriber-leave-timer 20
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure reverse OIF mapping:

1. Configure a logical interface for unicast data traffic.

   ```
   [edit interfaces ge-2/0/0]
   user@host# set unit 0 family inet address 30.0.0.2/24
   ```

2. Configure a logical interface for subscriber control traffic.

   ```
   [edit interfaces ge-2/2/0]
   user@host# set hierarchical-scheduler
   user@host# set vlan-tagging
   user@host# set unit 10 vlan-id 10
   user@host# set unit 10 family inet address 40.0.0.2/24
   ```

3. Configure two logical interfaces on which QoS adjustments are made.

   ```
   [edit interfaces ge-2/2/0]
   user@host# set unit 50 vlan-id 50
   user@host# set unit 50 family inet address 50.0.0.2/24
   user@host# set unit 51 vlan-id 51
   user@host# set unit 51 family inet address 50.0.1.2/24
   ```

4. Configure a policy.

   ```
   [edit policy-options policy-statement all-mcast-groups]
   user@host# set from source-address-filter 30.0.0.0/8 orlonger
   user@host# set then accept
   ```

5. Enable a flow map that references the policy.

   ```
   [edit routing-options multicast]
   user@host# set flow-map map1 policy all-mcast-groups
   user@host# set flow-map map1 bandwidth 10m adaptive
   ```

6. Enable OIF mapping on the logical interface that receives subscriber control traffic.

   ```
   [edit routing-options multicast]
   user@host# set interface ge-2/2/0.10 maximum-bandwidth 500m
   user@host# set interface ge-2/2/0.10 reverse-oif-mapping
   user@host# set interface ge-2/2/0.10 subscriber-leave-timer 20
   ```

7. Configure PIM and IGMP.

   ```
   [edit protocols]
   user@host# set igmp interface all
   user@host# set igmp interface fxp0.0 disable
   user@host# set pim rp local address 20.0.0.2
   user@host# set pim interface all
   user@host# set pim interface fxp0.0 disable
   user@host# set pim interface ge-2/2/0.10 disable
   ```

Copyright © 2016, Juniper Networks, Inc.

8. Configure the hierarchical scheduler by configuring a shaping rate for the physical interface and a slower shaping rate for the logical interfaces on which QoS adjustments are made.

```
[edit class-of-service interfaces ge-2/2/0]
user@host# set shaping-rate 240m
user@host# set unit 50 output-traffic-control-profile tcp-ifl
user@host# set unit 51 output-traffic-control-profile tcp-ifl

[edit class-of-service traffic-control-profiles tcp-30m-no-smap]
user@host# set shaping-rate 20m
```

Results From configuration mode, confirm your configuration by entering the **show class-of-service**, **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show class-of-service
traffic-control-profiles {
    tcp-ifl {
        shaping-rate 20m;
    }
}
interfaces {
    ge-2/2/0 {
        shaping-rate 240m;
        unit 50 {
            output-traffic-control-profile tcp-ifl;
        }
        unit 51 {
            output-traffic-control-profile tcp-ifl;
        }
    }
}
```

```
user@host# show interfaces
ge-2/0/0 {
    unit 0 {
        family inet {
            address 30.0.0.2/24;
        }
    }
}
ge-2/2/0 {
    hierarchical-scheduler;
    vlan-tagging;
    unit 10 {
        vlan-id 10;
        family inet {
            address 40.0.0.2/24;
        }
    }
    unit 50 {
        vlan-id 50;
        family inet {
```


Copyright © 2016, Juniper Networks, Inc.

```
        address 50.0.0.2/24;
      }
    }
    unit 51 {
      vlan-id 51;
      family inet {
        address 50.0.1.2/24;
      }
    }
}
```

```
user@host# show policy-options
policy-statement all-mcast-groups {
  from {
    source-address-filter 30.0.0.0/8 orlonger;
  }
  then accept;
}
```

```
user@host# show protocols
igmp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
pim {
  rp {
    local {
      address 20.0.0.2;
    }
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface ge-2/2/0.10 {
    disable;
  }
}
```

```
user@host# show routing-options
multicast {
  flow-map map1 {
    policy all-mcast-groups;
    bandwidth 10m adaptive;
  }
  interface ge-2/2/0.10 {
    maximum-bandwidth 500m;
    reverse-oif-mapping;
    subscriber-leave-timer 20;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring an OIF Map

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-2/3/8 unit 0 family inet6 address C300:0101::/24
set interfaces ge-2/3/9 vlan-tagging
set interfaces ge-2/3/9 unit 1 vlan-id 1
set interfaces ge-2/3/9 unit 1 family inet6 address C400:0101::/24
set interfaces ge-2/3/9 unit 2 vlan-id 2
set interfaces ge-2/3/9 unit 2 family inet6 address C400:0201::/24
set interfaces ge-2/3/9 unit 4000 vlan-id 4000
set interfaces ge-2/3/9 unit 4000 family inet6 address C40F:A001::/24
set interfaces ge-2/3/9 unit 4001 vlan-id 4001
set interfaces ge-2/3/9 unit 4001 family inet6 address C40F:A101::/24
set policy-options policy-statement g539-v6 term g539-4000 from route-filter
  FF05:0101:0000::/39 orlonger
set policy-options policy-statement g539-v6 term g539-4000 then map-to-interface
  ge-2/3/9.4000
set policy-options policy-statement g539-v6 term g539-4000 then accept
set policy-options policy-statement g539-v6 term g539-4001 from route-filter
  FF05:0101:0200::/39 orlonger
set policy-options policy-statement g539-v6 term g539-4001 then map-to-interface
  ge-2/3/9.4001
set policy-options policy-statement g539-v6 term g539-4001 then accept
set policy-options policy-statement g539-v6 term self from route-filter
  FF05:0101:0700::/40 orlonger
set policy-options policy-statement g539-v6 term self then map-to-interface self
set policy-options policy-statement g539-v6 term self then accept
set policy-options policy-statement g539-v6-all term g539 from route-filter 0::/0 orlonger
set policy-options policy-statement g539-v6-all term g539 then map-to-interface
  ge-2/3/9.4000
set policy-options policy-statement g539-v6-all term g539 then accept
set protocols mld interface fxp0.0 disable
set protocols mld interface ge-2/3/9.4000 passive
set protocols mld interface ge-2/3/9.4001 passive
set protocols mld interface ge-2/3/9.1 version 1
set protocols mld interface ge-2/3/9.1 oif-map g539-v6
set protocols mld interface ge-2/3/9.2 version 2
set protocols mld interface ge-2/3/9.2 oif-map g539-v6
set protocols pim rp local address 20.0.0.4
set protocols pim rp local family inet6 address C000::1
set protocols pim interface ge-2/3/8.0 mode sparse
set protocols pim interface ge-2/3/8.0 version 2
set routing-options multicast interface ge-2/3/9.1 no-qos-adjust
set routing-options multicast interface ge-2/3/9.2 no-qos-adjust
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure reverse OIF mapping:

1. Configure a logical interface for unicast data traffic.

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces ge-2/3/8 ]
user@host# set unit 0 family inet6 address C300:0101::/24
```

2. Configure logical interfaces for subscriber VLANs.

```
[edit interfaces ge-2/3/9]
user@host# set vlan-tagging
user@host# set unit 1 vlan-id 1
user@host# set unit 1 family inet6 address C400:0101::/24
user@host# set unit 2 vlan-id 2
user@host# set unit 2 family inet6 address C400:0201::/24 lo0 unit 0 family inet6
  address C000::1/128
user@host# set unit 2 family inet6 address C400:0201::/24
```

3. Configure two map-to logical interfaces.

```
[edit interfaces ge-2/2/0]
user@host# set unit 4000 vlan-id 4000
user@host# set unit 4000 family inet6 address C40F:A001::/24
user@host# set unit 4001 vlan-id 4001
user@host# set unit 4001 family inet6 address C40F:A101::/24
```

4. Configure the OIF map.

```
[edit policy-options policy-statement g539-v6]
user@host# set term g539-4000 from route-filter FF05:0101:0000::/39 orlonger
user@host# set then map-to-interface ge-2/3/9.4000
user@host# set then accept
user@host# set term g539-4001 from route-filter FF05:0101:0200::/39 orlonger
user@host# set then map-to-interface ge-2/3/9.4001
user@host# set then accept
user@host# set term self from route-filter FF05:0101:0700::/40 orlonger
user@host# set then map-to-interface self
user@host# set then accept


[edit policy-options policy-statement g539-v6-all]
user@host# set term g539 from route-filter 0::/0 orlonger
user@host# set then map-to-interface ge-2/3/9.4000
user@host# set then accept
```

5. Disable QoS adjustment on the subscriber VLANs.

```
[edit routing-options multicast]
user@host# set interface ge-2/3/9.1 no-qos-adjust
user@host# set interface ge-2/3/9.2 no-qos-adjust
```

6. Configure PIM and MLD. Point the MLD subscriber VLANs to the OIF map.

```
[edit protocols]
user@host# set pim rp local address 20.0.0.4
user@host# set pim rp local family inet6 address C000::1 #C000::1 is the address
  of lo0
user@host# set pim interface ge-2/3/8.0 mode sparse
user@host# set pim interface ge-2/3/8.0 version 2
user@host# set mld interface fxp0.0 disable
user@host# set interface ge-2/3/9.4000 passive
user@host# set interface ge-2/3/9.4001 passive
user@host# set interface ge-2/3/9.1 version 1
user@host# set interface ge-2/3/9.1 oif-map g539-v6
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface ge-2/3/9.2 version 2
user@host# set interface ge-2/3/9.2 oif-map g539-v6
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-2/3/8 {
  unit 0 {
    family inet6 {
      address C300:0101::/24;
    }
  }
}
ge-2/3/9 {
  vlan-tagging;
  unit 1 {
    vlan-id 1;
    family inet6 {
      address C400:0101::/24;
    }
  }
  unit 2 {
    vlan-id 2;
    family inet6 {
      address C400:0201::/24;
    }
  }
  unit 4000 {
    vlan-id 4000;
    family inet6 {
      address C40F:A001::/24;
    }
  }
  unit 4001 {
    vlan-id 4001;
    family inet6 {
      address C40F:A101::/24;
    }
  }
}
```

```
user@host# show policy-options
policy-statement g539-v6 {
  term g539-4000 {
    from {
      route-filter FF05:0101:0000::/39 orlonger;
    }
    then {
      map-to-interface ge-2/3/9.4000;
      accept;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    term g539-4001 {
      from {
        route-filter FF05:0101:0200::/39 orlonger;
      }
      then {
        map-to-interface ge-2/3/9.4001;
        accept;
      }
    }
    term self {
      from {
        route-filter FF05:0101:0700::/40 orlonger;
      }
      then {
        map-to-interface self;
        accept;
      }
    }
}
policy-statement g539-v6-all {
    term g539 {
      from {
        route-filter 0::/0 orlonger;
      }
      then {
        map-to-interface ge-2/3/9.4000;
        accept;
      }
    }
}
```

```
user@host# show protocols
mld {
    interface fxp0.0 {
      disable;
    }
    interface ge-2/3/9.4000 {
      passive;
    }
    interface ge-2/3/9.4001 {
      passive;
    }
    interface ge-2/3/9.1 {
      version 1;
      oif-map g539-v6;
    }
    interface ge-2/3/9.2 {
      version 2;
      oif-map g539-v6;
    }
}
pim {
    rp {
      local {
        address 20.0.0.4;
        family inet6 {
```

Copyright © 2016, Juniper Networks, Inc.

```
            address C000::1;
          }
        }
      }
      interface ge-2/3/8.0 {
        mode sparse;
        version 2;
      }
    }
```

```
user@host# show routing-options
multicast {
  interface ge-2/3/9.1 no-qos-adjust;
  interface ge-2/3/9.2 no-qos-adjust;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To verify the configuration, run the following commands:

- **show igmp statistics**
- **show class-of-service interface**
- **show interfaces statistics**
- **show mld statistics**
- **show multicast interface**
- **show policy**

Copyright © 2016, Juniper Networks, Inc.

Configuring Multicast Routing over IP Demux Interfaces

In a subscriber management network, fields in packets sent from IP demux interfaces are intended to correspond to a specific client that resides on the other side of an aggregation device (for example, a Multiservice Access Node [MSAN]). However, packets sent from a Broadband Services Router (BSR) to an MSAN do not identify the demux interface. Once it obtains a packet, it is up to the MSAN device to determine which client receives the packet.

Depending on the intelligence of the MSAN device, determining which client receives the packet can occur in an inefficient manner. For example, when it receives IGMP control traffic, an MSAN might forward the control traffic to all clients instead of the one intended client. In addition, once a data stream destination is established, though an MSAN can use IGMP snooping to determine which hosts reside in a particular group and limit data streams to only that group, the MSAN still must send multiple copies of the data stream to each group member, even if that data stream is intended for only one client in the group.

Various multicast features, when combined, enable you to avoid the inefficiencies mentioned above. These features include the following:

- The ability to configure the IP demux interface **family** statement to use **inet** for either the numbered or unnumbered primary interface.
- The ability to configure IGMP on the primary interface to send general queries for all clients. The demux configuration prevents the primary IGMP interface from receiving any client IGMP control packets. Instead, all IGMP control packets go to the demux interfaces. However, to guarantee that no joins occur on the primary interface:
  - For static IGMP interfaces—Include the **passive send-general-query** statement in the IGMP configuration at the **[edit protocols igmp interface *interface-name*]** hierarchy level.
  - For dynamic IGMP demux interfaces—Include the **passive send-general-query** statement at the **[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]** hierarchy level.
- The ability to map all multicast groups to the primary interface as follows:
  - For static IGMP interfaces—Include the **oif-map** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.
  - For dynamic IGMP demux interfaces—Include the **oif-map** statement at the **[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]** hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

Using the **oif-map** statement, you can map the same IGMP group to the same output interface and send only one copy of the multicast stream from the interface.

- The ability to configure IGMP on each demux interface. To prevent duplicate general queries:
  - For static IGMP interfaces—Include the **passive allow-receive send-group-query** statement at the **[edit protocols igmp interface *interface-name*]** hierarchy level.
  - For dynamic demux interfaces—Include the **passive allow-receive send-group-query** statement at the **[edit dynamic-profiles *profile-name* protocols igmp interface *interface-name*]** hierarchy level.

NOTE: To send only one copy of each group, regardless of how many customers join, use the oif-map statement as previously mentioned.

Classifying Packets by Egress Interface

For Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers with the Intelligent Queuing (IQ), IQ2, Enhanced IQ (IQE), Multiservices link services intelligent queuing (LSQ) interfaces, or ATM2 PICs, you can classify unicast and multicast packets based on the egress interface. For unicast traffic, you can also use a multifield filter, but only egress interface classification applies to multicast traffic as well as unicast traffic. If you configure egress classification of an interface, you cannot perform Differentiated Services code point (DSCP) rewrites on the interface. By default, the system will not perform any classification based on the egress interface.

To enable packet classification by the egress interface, you first configure a forwarding class map and one or more queue numbers for the egress interface at the **[edit class-of-service forwarding-classes-interface-specific *forwarding-class-map-name*]** hierarchy level:

```
[edit class-of-service]
forwarding-classes-interface-specific  forwarding-class-map-name {
   class class-name queue-num queue-number [ restricted-queue queue-number ];
}
```

For T Series routers that are restricted to only four queues, you can control the queue assignment with the **restricted-queue** option, or you can allow the system to automatically determine the queue in a modular fashion. For example, a map assigning packets to queue 6 would map to queue 2 on a four-queue system.

NOTE: If you configure an output forwarding class map associating a forwarding class with a queue number, this map is not supported on multiservices link services intelligent queuing (lsq-) interfaces.

Copyright © 2016, Juniper Networks, Inc.

Once the forwarding class map has been configured, you apply the map to the logical interface by using the **output-forwarding-class-map** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
output-forwarding-class-map forwarding-class-map-name;
```

All parameters relating to the queues and forwarding class must be configured as well. For more information about configuring forwarding classes and queues, see Configuring Forwarding Classes.

This example shows how to configure an interface-specific forwarding-class map named **FCMAP1** that restricts queues 5 and 6 to different queues on four-queue systems and then applies **FCMAP1** to **unit 0** of interface **ge-6/0/0**:

```
[edit class-of-service]
forwarding-classes-interface-specific FCMAP1 {
  class FC1 queue-num 6 restricted-queue 3;
  class FC2 queue-num 5 restricted-queue 2;
  class FC3 queue-num 3;
  class FC4 queue-num 0;
  class FC3 queue-num 0;
  class FC4 queue-num 1;
}

[edit class-of-service]
interfaces {
  ge-6/0/0 unit 0 {
    output-forwarding-class-map FCMAP1;
  }
}
```

Note that without the **restricted-queue** option in **FCMAP1**, the example would assign **FC1** and **FC2** to queues 2 and 1, respectively, on a system restricted to four queues.

Use the **show class-of-service forwarding-class *forwarding-class-map-name*** command to display the forwarding-class map queue configuration:

```
user@host> show class-of-service forwarding-class FCMAP2

Forwarding class                          ID      Queue  Restricted queue
  FC1                                      0       6          3
  FC2                                      1       5          2
  FC3                                      2       3          3
  FC4                                      3       0          0
  FC5                                      4       0          0
  FC6                                      5       1          1
  FC7                                      6       6          2
  FC8                                      7       7          3
```

Use the **show class-of-service interface *interface-name*** command to display the forwarding-class maps (and other information) assigned to a logical interface:

```
user@host> show class-of-service interface ge-6/0/0

Physical interface: ge-6/0/0, Index: 128
Queues supported: 8, Queues in use: 8
```

Copyright © 2016, Juniper Networks, Inc.

```
 Scheduler map: <default>, Index: 2
 Input scheduler map: <default>, Index: 3
 Chassis scheduler map: <default-chassis>, Index: 4

Logical interface: ge-6/0/0.0, Index: 67
 Object                  Name            Type                  Index
 Scheduler-map           sch-map1        Output                 6998
 Scheduler-map           sch-map1        Input                  6998
 Classifier              dot1p           ieee8021p              4906
 forwarding-class-map    FCMAP1          Output                 1221

Logical interface: ge-6/0/0.1, Index 68
 Object                  Name            Type                  Index
 Scheduler-map           <default>       Output                    2
 Scheduler-map           <default>       Input                     3

Logical interface: ge-6/0/0.32767, Index 69
 Object                  Name            Type                  Index
 Scheduler-map           <default>       Output                    2
 Scheduler-map           <default>       Input                     3
```

Related Documentation

- Examples: Configuring the Multicast Forwarding Cache on page 4374

Examples: Configuring the Multicast Forwarding Cache

- Understanding the Multicast Forwarding Cache on page 4374
- Example: Configuring the Multicast Forwarding Cache on page 4374
- Example: Configuring a Multicast Flow Map on page 4377

Understanding the Multicast Forwarding Cache

IP multicast protocols can create numerous entries in the multicast forwarding cache. If the forwarding cache fills up with entries that prevent the addition of higher-priority entries, applications and protocols might not function properly. You can manage the multicast forwarding cache properties by limiting the size of the cache and by controlling the length of time that entries remain in the cache. By managing timeout values, you can give preference to more important forwarding cache entries while removing other less important entries.

Example: Configuring the Multicast Forwarding Cache

When a routing device receives multicast traffic, it places the (S,G) route information in the multicast forwarding cache, **inet.1**. This example shows how to configure multicast forwarding cache limits to prevent the cache from filling up with entries.

- Requirements on page 4375
- Overview on page 4375
- Configuration on page 4375
- Verification on page 4376

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure a multicast protocol. This feature works with the following multicast protocols:
  - DVMRP
  - PIM-DM
  - PIM-SM
  - PIM-SSM

Overview

This example includes the following statements:

- **forwarding-cache**—Specifies how forwarding entries are aged out and how the number of entries is controlled.
- **timeout**—Specifies an idle period after which entries are aged out and removed from **inet.1**. You can specify a timeout in the range from 1 through 720 minutes.
- **threshold**—Enables you to specify threshold values on the forwarding cache to suppress (suspend) entries from being added when the cache entries reach a certain maximum and begin adding entries to the cache when the number falls to another threshold value. By default, no threshold values are enabled on the routing device.

  The suppress threshold suspends the addition of new multicast forwarding cache entries. If you do not specify a suppress value, multicast forwarding cache entries are created as necessary. If you specify a suppress threshold, you can optionally specify a reuse threshold, which sets the point at which the device resumes adding new multicast forwarding cache entries. During suspension, forwarding cache entries time out. After a certain number of entries time out, the reuse threshold is reached, and new entries are added. The range for both thresholds is from 1 through 200,000. If configured, the reuse value must be less than the suppression value. If you do not specify a reuse value, the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set routing-options multicast forwarding-cache threshold suppress 150000
set routing-options multicast forwarding-cache threshold reuse 34
set routing-options multicast forwarding-cache timeout 60
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the multicast forwarding cache:

1. Configure the maximum size of the forwarding cache.

   ```
   [edit routing-options multicast forwarding-cache]
   user@host# set threshold suppress 150000
   ```

2. Configure the amount of time (in minutes) entries can remain idle before being removed.

   ```
   [edit routing-options multicast forwarding-cache]
   user@host# set timeout 60
   ```

3. Configure the size of the forwarding cache when suppression stops and new entries can be added.

   ```
   [edit routing-options multicast forwarding-cache]
   user@host# set threshold reuse 70000
   ```

Results

Confirm your configuration by entering the **show routing-options** command.

```
user@host# show routing-options
multicast {
  forwarding-cache {
    threshold {
      suppress 150000;
      reuse 70000;
    }
    timeout 60;
  }
}
```

Verification

To verify the configuration, run the **show multicast route extensive** command.

```
user@host> show multicast route extensive
Family: INET
Group: 232.0.0.1
    Source: 11.11.11.11/32
    Upstream interface: fe-0/2/0.200
    Downstream interface list:
        fe-0/2/1.210
    Downstream interface list rejected by CAC:
        fe-0/2/1.220
    Session description: Source specific multicast
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 337
```


Copyright © 2016, Juniper Networks, Inc.

```
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 60 minutes
Wrong incoming interface notifications: 0
```

Example: Configuring a Multicast Flow Map

This example shows how to configure a flow map to prevent certain forwarding cache entries from aging out, thus allowing for faster failover from one source to another. Flow maps enable you to configure bandwidth variables and multicast forwarding cache timeout values for entries defined by the flow map policy.

- Requirements on page 4377
- Overview on page 4377
- Configuration on page 4379
- Verification on page 4380

Requirements

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol. See the *Junos OS Routing Protocols Library for Security Devices*.
- Configure a multicast protocol. This feature works with the following multicast protocols:
  - DVMRP
  - PIM-DM
  - PIM-SM
  - PIM-SSM

Overview

Flow maps are typically used for fast multicast source failover when there are multiple sources for the same group. For example, when one video source is actively sending the traffic, the forwarding states for other video sources are timed out after a few minutes. Later, when a new source starts sending the traffic again, it takes time to install a new forwarding state for the new source if the forwarding state is not already there. This switchover delay is worsened when there are many video streams. Using flow maps with longer timeout values or permanent cache entries helps reduce this switchover delay.

NOTE: The permanent forwarding state must exist on all routing devices in the path for fast source switchover to function properly.

Copyright © 2016, Juniper Networks, Inc.

This example includes the following statements:

- **bandwidth**—Specifies the bandwidth for each flow that is defined by a flow map to ensure that an interface is not oversubscribed for multicast traffic. If adding one more flow would cause overall bandwidth to exceed the allowed bandwidth for the interface, the request is rejected. A rejected request means that traffic might not be delivered out of some or all of the expected outgoing interfaces. You can define the bandwidth associated with multicast flows that match a flow map by specifying a bandwidth in bits per second or by specifying that the bandwidth is measured and adaptively modified.

  When you use the **adaptive** option, the bandwidth adjusts based on measurements made at 5-second intervals. The flow uses the maximum bandwidth value from the last 12 measured values (1 minute).

  When you configure a bandwidth value with the **adaptive** option, the bandwidth value acts as the starting bandwidth for the flow. The bandwidth then changes based on subsequent measured bandwidth values. If you do not specify a bandwidth value with the **adaptive** option, the starting bandwidth defaults to 2 megabits per second (Mbps).

  For example, the **bandwidth 2m adaptive** statement is equivalent to the **bandwidth adaptive** statement because they both use the same starting bandwidth (2 Mbps, the default). If the actual flow bandwidth is 4 Mbps, the measured flow bandwidth changes to 4 Mbps after reaching the first measuring point (5 seconds). However, if the actual flow bandwidth rate is 1 Mbps, the measured flow bandwidth remains at 2 Mbps for the first 12 measurement cycles (1 minute) and then changes to the measured 1 Mbps value.

- **flow-map**—Defines a flow map that controls the forwarding cache timeout of specified source and group addresses, controls the bandwidth for each flow, and specifies redundant sources. If a flow can match multiple flow maps, the first flow map applies.

- **forwarding-cache**—Enables you to configure the forwarding cache properties of entries defined by a flow map. You can specify a timeout of **never** to make the forwarding entries permanent, or you can specify a timeout in the range from 1 through 720 minutes. If you set the value to **never**, you can specify the **non-discard-entry-only** option to make an exception for entries that are in the pruned state. In other words, the **never non-discard-entry-only** statement allows entries in the pruned state to time out, while entries in the forwarding state never time out.

- **policy**—Specifies source and group addresses to which the flow map applies. This example creates a flow map policy called **policyForFlow1**. The policy matches the source address using the **source-address-filter** statement and matches the group address using the **prefix-list-filter** statement.

  NOTE: The addresses must match the configured policy for flow mapping to occur.

- **redundant-sources**—Specify redundant (backup) sources for flows identified by a flow map.

Copyright © 2016, Juniper Networks, Inc.

Outbound interfaces that are admitted for one of the forwarding entries are automatically admitted for any other entries identified by the redundant source configuration.

In this example, forwarding entries (10.11.11.11, g1) and (10.11.11.12, g1) match the flow map **flowMap1**. In this case, if a particular outbound interface is admitted for entry (10.11.11.11, g1), it is automatically admitted for entry (10.11.11.12, g1), even if there is no longer enough remaining bandwidth available after creating entry (10.11.11.11, g1). The interface is added because only one of the two sources can send traffic at any time.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options prefix-list permanentEntries1 232.1.1.0/24
set policy-options policy-statement policyForFlow1 from source-address-filter 11.11.11.11/32
    exact
set policy-options policy-statement policyForFlow1 from prefix-list-filter
    permanentEntries1 orlonger
set policy-options policy-statement policyForFlow1 then accept
set routing-options multicast flow-map flowMap1 policy policyForFlow1
set routing-options multicast flow-map flowMap1 bandwidth 2m
set routing-options multicast flow-map flowMap1 bandwidth adaptive
set routing-options multicast flow-map flowMap1 redundant-sources 10.11.11.11
set routing-options multicast flow-map flowMap1 redundant-sources 10.11.11.12
set routing-options multicast flow-map flowMap1 forwarding-cache timeout never
    non-discard-entry-only
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a flow map:

1. Configure the flow map policy.

   ```
   [edit policy-options]
   user@host# set prefix-list permanentEntries1 232.1.1.0/24
   user@host# set policy policyForFlow1 from source-address-filter 11.11.11.11/32 exact
   user@host# set policy policyForFlow1 from prefix-list-filter permanentEntries1
     orlonger
   user@host# set policy policyForFlow1 then accept
   ```

2. Apply the flow map policy.

   ```
   [edit routing-options]
   user@host# set multicast flow-map flowMap1 policy policyForFlow1
   ```

3. Configure permanent forwarding entries (that is, entries that never time out), and enable entries in the pruned state to time out.

   ```
   [edit routing-options]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set multicast flow-map flowMap1 forwarding-cache timeout never
  non-discard-entry-only
```

4. Configure the flow map bandwidth to be adaptive with a default starting bandwidth of 2 Mbps.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 bandwidth 2m adaptive
```

5. Specify backup sources.

```
[edit routing-options]
user@host# set multicast flow-map flowMap1 redundant-sources [ 10.11.11.11 10.11.11.12
  ]
```

6. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy-options** and **show routing-options** commands.

```
user@host# show policy-options
prefix-list permanentEntries1 {
   232.1.1.0/24;
}
policy-statement policyForFlow1 {
   from {
      source-address-filter 11.11.11.11/32 exact;
      prefix-list-filter permanentEntries1 orlonger;
   }
   then accept;
}
```

```
user@host# show routing-options
multicast {
   flow-map flowMap1 {
      policy policyForFlow1;
      bandwidth 2m adaptive;
      redundant-sources [ 10.11.11.11 10.11.11.12 ];
      forwarding-cache {
         timeout never non-discard-entry-only;
      }
   }
}
```

Verification

To verify the configuration, run the following commands:

- show multicast flow-map
- show multicast route **extensive**

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Examples: Configuring Bandwidth Management on page 4353

Example: Configuring Ingress PE Redundancy

- Understanding Ingress PE Redundancy on page 4381
- Example: Configuring Ingress PE Redundancy on page 4381

Understanding Ingress PE Redundancy

In many network topologies, point-to-multipoint label-switched paths (LSPs) are used to distribute multicast traffic over a virtual private network (VPN). When traffic engineering is added to the provider edge (PE) routers, a popular deployment option has been to use traffic-engineered point-to-multipoint LSPs at the origin PE. In these network deployments, the PE is a single point of failure. Network operators have previously provided redundancy by broadcasting duplicate streams of multicast traffic from multiple PEs, a practice which at least doubles the bandwidth required for each stream.

Ingress PE redundancy eliminates the bandwidth duplication requirement by configuring one or more ingress PEs as a group. Within a group, one PE is designated as the primary PE and one or more others become backup PEs for the configured traffic stream. The solution depends on a full mesh of point-to-point (P2P) LSPs among the primary and backup PEs. Also, you must configure a full set of point-to-multipoint LSPs at the backup PEs, even though these point-to-multipoint LSPs at the backup PEs are not sending any traffic or using any bandwidth. The P2P LSPs are configured with bidirectional forwarding detection (BFD). When BFD detects a failure on the primary PE, a new designated forwarder is elected for the stream.

Example: Configuring Ingress PE Redundancy

This example shows how to configure one PE as part of a backup PE group to enable ingress PE redundancy for multicast traffic streams.

- Requirements on page 4381
- Overview on page 4382
- Configuration on page 4383
- Verification on page 4386

Requirements

Before you begin:

- Configure the router interfaces.
- Configure a full mesh of P2P LSPs between the PEs in the backup group.

Copyright © 2016, Juniper Networks, Inc.

Overview

Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution. When point-to-multipoint LSPs are used for multicast traffic, the PE device can become a single point of failure. One way to provide redundancy is by broadcasting duplicate streams from multiple PEs, thus doubling the bandwidth requirements for each stream. This feature implements redundancy between two or more PEs by designating a primary and one or more backup PEs for each configured stream. The solution depends on the configuration of a full mesh of P2P LSPs between the primary and backup PEs. These LSPs are configured with Bidirectional Forwarding Detection (BFD) running on top of them. BFD is used on the backup PEs to detect failure on the primary PE routing device and to elect a new designated forwarder for the stream.

A full mesh is required so that each member of the group can make an independent decision about the health of the other PEs and determine the designated forwarder for the group. The key concept in a backup PE group is that of a designated PE. A designated PE is a PE that forwards data on the static route. All other PEs in the backup PE group do not forward any data on the static route. This allows you to have one designated forwarder. If the designated forwarder fails, another PE takes over as the designated forwarder, thus allowing the traffic flow to continue uninterrupted.

Each PE in the backup PE group makes its own local decision regarding the designated forwarder. Thus, there is no inter-PE communication regarding designated forwarder. A PE computes the designated forwarder based on the IP address of all PEs and the connectivity status of other PEs. Connectivity status is determined based on the state of the BFD session on the P2P LSP to a PE.

A PE chosen is as the designated forwarder if it satisfies the following conditions:

- The PE is in the UP state. Either it is the local PE, or the BFD session on the P2P LSP to that PE is in the UP state.
- The PE has the lowest IP address among all PEs that are in the UP state.

Because all PEs have P2P LSPs to each other, each PE can determine the UP state of each other PE, and all PEs converge to the same designated forwarder.

If the designated forwarder PE fails, then all other PEs lose connectivity with the designated forwarder, and their BFD session ends. Consequently, other PEs then choose another designated forwarder. The new forwarder starts forwarding traffic. Thus, the traffic loss is limited to the failure detection time, which is the BFD session detection time.

When a PE that was the designated forwarder fails and then resumes operating, all other PEs recognize this fact, rerun the designated forwarder algorithm, and choose the PE as the designated forwarder. Consequently, the backup designated forwarder stops forwarding traffic. Thus, traffic switches back to the most eligible designated forwarder.

This example includes the following statements:

Copyright © 2016, Juniper Networks, Inc.

- **associate-backup-pe-groups**—Monitors the health of the routing device at the other end of the LSP. You can configure multiple backup PE groups that contain the same routing device's address. Failure of this LSP indicates to all of these groups that the destination PE routing device is down. So, the **associate-backup-pe-groups** statement is not tied to any specific group but applies to all groups that are monitoring the health of the LSP to the remote address.

  If there are multiple LSPs with the **associate-backup-pe-groups** statement to the same destination PE, then the local routing device picks the first LSP to that PE for detection purposes.

  We do not recommend configuring multiple LSPs to the same destination. If you do, make sure that the LSP parameters (for example, liveliness detection) are similar to avoid false failure notification even when the remote PE is up.

- **backup-pe-group**—Configures ingress PE redundancy for multicast traffic streams.
- **bfd-liveness-detection**—Enables BFD for each LSP.
- **label-switched-path**—Configures an LSP. You must configure a full mesh of P2P LSPs between the primary and backup PEs.

  NOTE: We recommend that you configure the P2P LSPs with fast reroute and node link protection so that link failures do not result in the LSP failure. For the purpose of PE redundancy, a failure in the P2P LSP is treated as a PE failure. Redundancy in the inter-PE path is also encouraged.

- **p2mp-lsp-next-hop**—Enables you to associate a backup PE group with a static route.
- **static**—Applies the backup group to a static route on the PE. This ensures that the static route is active (installed in the forwarding table) when the local PE is the designated forwarder for the configured backup PE group.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set policy-options policy-statement no-rpf from route-filter 225.1.1.1/32 exact
set policy-options policy-statement no-rpf then reject
set protocols mpls label-switched-path backup_PE1 to 10.255.16.61
set protocols mpls label-switched-path backup_PE1 oam bfd-liveness-detection
  minimum-interval 500
set protocols mpls label-switched-path backup_PE1 oam bfd-liveness-detection multiplier
  3
set protocols mpls label-switched-path backup_PE1 associate-backup-pe-groups
set protocols mpls label-switched-path dest1 to 10.255.16.57
set protocols mpls label-switched-path dest1 p2mp p2mp-lsp
set protocols mpls label-switched-path dest2 to 10.255.16.55
set protocols mpls label-switched-path dest2 p2mp p2mp-lsp
set protocols mpls interface all
```

Copyright © 2016, Juniper Networks, Inc.

```
set protocols mpls interface fxp0.0 disable
set routing-options static route 1.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
set routing-options static route 1.1.1.1/32 backup-pe-group g1
set routing-options static route 225.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
set routing-options static route 225.1.1.1/32 backup-pe-group g1
set routing-options multicast rpf-check-policy no-rpf
set routing-options multicast interface fe-1/3/3.0 enable
set routing-options multicast backup-pe-group g1 backups 10.255.16.61
set routing-options multicast backup-pe-group g1 local-address 10.255.16.59
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure ingress PE redundancy:

1. Configure the multicast settings.

   ```
   [edit routing-options multicast]
   user@host# set rpf-check-policy no-rpf
   user@host# set interface fe-1/3/3.0 enable
   ```

2. Configure the RPF policy.

   ```
   [edit policy-options policy-statement no-rpf]
   user@host# set from route-filter 225.1.1.1/32 exact
   user@host# set then reject
   ```

3. Configure the backup PE group.

   ```
   [edit routing-options multicast]
   user@host# set backup-pe-group g1 backups 10.255.16.61
   user@host# set backup-pe-group g1 local-address 10.255.16.59
   ```

4. Configure the static routes for the point-to-multipoint LSPs backup PE group.

   ```
   [edit routing-options static]
   user@host# set route 1.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
   user@host# set route 1.1.1.1/32 backup-pe-group g1
   user@host# set route 225.1.1.1/32 p2mp-lsp-next-hop p2mp-lsp
   user@host# set route 225.1.1.1/32 backup-pe-group g1
   ```

5. Configure the MPLS interfaces.

   ```
   [edit protocols mpls]
   user@host# set interface all
   user@host# set interface fxp0.0 disable
   ```

6. Configure the LSP to the redundant router.

   ```
   [edit protocols mpls]
   user@host# set label-switched-path backup_PE1 to 10.255.16.61
   user@host# set label-switched-path backup_PE1 oam bfd-liveness-detection
     minimum-interval 500
   user@host# set label-switched-path backup_PE1 oam bfd-liveness-detection
     multiplier 3
   user@host# set label-switched-path backup_PE1 associate-backup-pe-groups
   ```


Copyright © 2016, Juniper Networks, Inc.

7. Configure LSPs to two traffic destinations.

```
[edit protocols mpls]
user@host# set label-switched-path dest1 to 10.255.16.57
user@host# set label-switched-path dest1 p2mp p2mp-lsp
user@host# set label-switched-path dest2 to 10.255.16.55
user@host# set label-switched-path dest2 p2mp p2mp-lsp
```

8. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

Results

Confirm your configuration by entering the **show policy**, **show protocols**, and **show routing-options** commands.

```
user@host# show policy
policy-statement no-rpf {
  from {
    route-filter 225.1.1.1/32 exact;
  }
  then reject;
}
```

```
user@host# show protocols
mpls {
  label-switched-path backup_PE1 {
    to 10.255.16.61;
    oam {
      bfd-liveness-detection {
        minimum-interval 500;
        multiplier 3;
      }
    }
    associate-backup-pe-groups;
  }
  label-switched-path dest1 {
    to 10.255.16.57;
    p2mp p2mp-lsp;
  }
  label-switched-path dest2 {
    to 10.255.16.55;
    p2mp p2mp-lsp;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

```
user@host# show routing-options
static {
  route 1.1.1.1/32 {
    p2mp-lsp-next-hop p2mp-lsp;
    backup-pe-group g1;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    route 225.1.1.1/32 {
      p2mp-lsp-next-hop p2mp-lsp;
      backup-pe-group g1;
    }
  }
  multicast {
    rpf-check-policy no-rpf;
    interface fe-1/3/3.0 enable;
    backup-pe-group g1 {
      backups 10.255.16.61;
      local-address 10.255.16.59;
    }
  }
```

Verification

To verify the configuration, run the following commands:

- **show mpls lsp**
- **show multicast backup-pe-groups**
- **show multicast rpf**

Related Documentation

- Examples: Configuring Bandwidth Management on page 4353
- Examples: Configuring the Multicast Forwarding Cache on page 4374

Copyright © 2016, Juniper Networks, Inc.

PART 65

Configuration Statements and Operational Commands

- Configuration Statements on page 4389
- Operational Commands on page 4631

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 196

Configuration Statements

- accept-remote-source on page 4396
- accounting (Protocols IGMP) on page 4396
- accounting (Protocols IGMP AMT Interface) on page 4397
- accounting (Protocols IGMP Interface) on page 4397
- accounting (Protocols MLD) on page 4398
- accounting (Protocols MLD Interface) on page 4398
- active-source-limit on page 4399
- address (Anycast RPs) on page 4400
- address (Bidirectional Rendezvous Points) on page 4401
- address (Local RPs) on page 4402
- address (Static RPs) on page 4403
- algorithm on page 4404
- amt (IGMP) on page 4405
- amt (Protocols) on page 4406
- anycast-pim on page 4407
- anycast-prefix on page 4408
- asm-override-ssm on page 4409
- assert-timeout on page 4410
- authentication on page 4411
- authentication-key on page 4412
- auto-rp on page 4413
- backoff-period on page 4414
- backup-pe-group on page 4415
- backups on page 4416
- bandwidth on page 4417
- bfd-liveness-detection on page 4418
- bidirectional (Interface) on page 4419
- bidirectional (RP) on page 4420

Copyright © 2016, Juniper Networks, Inc.

- bootstrap on page 4421
- bootstrap-export on page 4422
- bootstrap-import on page 4423
- bootstrap-priority on page 4424
- data-encapsulation on page 4425
- default-peer on page 4426
- defaults on page 4427
- dense-groups on page 4428
- detection-time (BFD for PIM) on page 4429
- df-election on page 4430
- disable (PIM) on page 4431
- disable (PIM Graceful Restart) on page 4432
- disable (Protocols DVMRP) on page 4432
- disable (Protocols IGMP) on page 4433
- disable (Protocols MLD) on page 4433
- disable (Protocols MSDP) on page 4434
- disable (Protocols SAP) on page 4435
- dr-election-on-p2p on page 4435
- dr-register-policy on page 4436
- dvmrp on page 4437
- embedded-rp on page 4438
- exclude (Protocols IGMP) on page 4438
- exclude (Protocols MLD) on page 4439
- export (Protocols DVMRP) on page 4439
- export (Protocols MSDP) on page 4440
- export (Protocols PIM) on page 4441
- export (Bootstrap) on page 4442
- family (Bootstrap) on page 4443
- family (Local RP) on page 4444
- family (Protocols AMT Relay) on page 4445
- family (Protocols PIM) on page 4446
- flood-groups on page 4447
- flow-map on page 4448
- forwarding-cache (Bridge Domains) on page 4449
- forwarding-cache (Flow Maps) on page 4449
- forwarding-cache (Multicast) on page 4450
- graceful-restart (Multicast Snooping) on page 4451

Copyright © 2016, Juniper Networks, Inc.

- graceful-restart (Protocols PIM) on page 4452
- group (Bridge Domains) on page 4453
- group (Protocols IGMP) on page 4454
- group (Protocols MLD) on page 4455
- group (Protocols MSDP) on page 4456
- group (RPF Selection) on page 4457
- group-count on page 4458
- group-count (Protocols MLD) on page 4458
- group-increment (Protocols IGMP) on page 4459
- group-increment (Protocols MLD) on page 4459
- group-limit (Protocols IGMP) on page 4460
- group-limit (IGMP and MLD Snooping) on page 4461
- group-limit (Protocols MLD) on page 4462
- group-policy (Protocols IGMP) on page 4462
- group-policy (Protocols IGMP AMT Interface) on page 4463
- group-policy (Protocols MLD) on page 4463
- group-ranges on page 4464
- hello-interval on page 4465
- hold-time (Protocols DVMRP) on page 4465
- hold-time (Protocols PIM) on page 4466
- host-only-interface on page 4467
- igmp on page 4468
- igmp-snooping on page 4470
- ignore-stp-topology-change on page 4471
- immediate-leave (Bridge Domains) on page 4472
- immediate-leave (Protocols IGMP) on page 4474
- immediate-leave (Protocols MLD) on page 4475
- import (Protocols DVMRP) on page 4476
- import (Protocols MSDP) on page 4477
- import (Protocols PIM) on page 4478
- import (Protocols PIM Bootstrap) on page 4479
- infinity on page 4480
- inet on page 4481
- interface (Bridge Domains) on page 4482
- interface (Protocols DVMRP) on page 4483
- interface (Protocols IGMP) on page 4484
- interface (Protocols MLD) on page 4485

Copyright © 2016, Juniper Networks, Inc.

- interface (Protocols PIM) on page 4486
- interface (Routing Options) on page 4487
- interface (Scoping) on page 4488
- join-load-balance on page 4489
- join-prune-timeout on page 4490
- key-chain on page 4490
- listen on page 4491
- local on page 4492
- local-address (Protocols AMT) on page 4493
- local-address (Protocols MSDP) on page 4494
- local-address (Protocols PIM) on page 4495
- local-address (Routing Options) on page 4496
- loose-check on page 4497
- mapping-agent-election on page 4498
- maximum (MSDP Active Source Messages) on page 4499
- maximum-bandwidth on page 4500
- maximum-rps on page 4501
- maximum-transmit-rate (Protocols IGMP) on page 4501
- metric (Protocols DVMRP) on page 4502
- minimum-interval (PIM BFD Liveness Detection) on page 4502
- minimum-interval (PIM BFD Transmit Interval) on page 4503
- minimum-receive-interval on page 4504
- mld on page 4505
- mode (Protocols DVMRP) on page 4506
- mode (Protocols MSDP) on page 4507
- mode (Protocols PIM) on page 4508
- msdp on page 4509
- multicast (Dynamic Profiles Routing Options) on page 4511
- multicast-router-interface on page 4513
- multicast-snooping-options on page 4514
- multichassis-lag-replicate-state on page 4515
- multiplier on page 4516
- neighbor-policy on page 4516
- nexthop-hold-time on page 4517
- next-hop (PIM RPF Selection) on page 4517
- no-adaptation (PIM BFD Liveness Detection) on page 4518
- no-bidirectional-mode on page 4519

Copyright © 2016, Juniper Networks, Inc.

- no-multicast-echo on page 4520
- no-qos-adjust on page 4521
- offer-period on page 4522
- oif-map (IGMP Interface) on page 4523
- oif-map (MLD Interface) on page 4523
- override (PIM Static RP) on page 4524
- override-interval on page 4525
- passive (IGMP) on page 4526
- passive (MLD) on page 4527
- peer on page 4528
- pgm on page 4529
- pim on page 4530
- pim-to-igmp-proxy on page 4534
- pim-to-mld-proxy on page 4535
- policy (Flow Maps) on page 4536
- policy (SSM Maps) on page 4536
- prefix-list (PIM RPF Selection) on page 4537
- priority (Bootstrap) on page 4538
- priority (PIM Interfaces) on page 4539
- priority (PIM RPs) on page 4540
- propagation-delay on page 4541
- proxy on page 4542
- query-interval (Bridge Domains) on page 4543
- query-interval (Protocols IGMP) on page 4544
- query-interval (Protocols IGMP AMT) on page 4545
- query-interval (Protocols MLD) on page 4545
- query-last-member-interval (Bridge Domains) on page 4546
- query-last-member-interval (Protocols IGMP) on page 4547
- query-last-member-interval (Protocols MLD) on page 4547
- query-response-interval (Bridge Domains) on page 4548
- query-response-interval (Protocols IGMP) on page 4549
- query-response-interval (Protocols IGMP AMT) on page 4550
- query-response-interval (Protocols MLD) on page 4551
- redundant-sources on page 4552
- relay (AMT Protocol) on page 4553
- relay (IGMP) on page 4554
- reset-tracking-bit on page 4555

Copyright © 2016, Juniper Networks, Inc.

- restart-duration on page 4556
- reverse-oif-mapping on page 4557
- rib-group (Protocol DVMRP) on page 4557
- rib-group (Protocols MSDP) on page 4558
- rib-group (Protocols PIM) on page 4559
- robust-count (Bridge Domains) on page 4560
- robust-count (Protocols IGMP) on page 4561
- robust-count (Protocols IGMP AMT) on page 4561
- robust-count (Protocols MLD) on page 4562
- robustness-count on page 4563
- rp on page 4564
- rp-register-policy on page 4566
- rp-set on page 4567
- rpf-check-policy on page 4568
- rpf-selection on page 4569
- sap on page 4570
- scope on page 4571
- scope-policy on page 4572
- secret-key-timeout on page 4573
- source (Bridge Domains) on page 4573
- source (PIM RPF Selection) on page 4574
- source (Protocols IGMP) on page 4575
- source (Protocols MLD) on page 4575
- source (Protocols MSDP) on page 4576
- source (Routing Instances) on page 4577
- source-address on page 4577
- source-count (Protocols IGMP) on page 4578
- source-count (Protocols MLD) on page 4578
- source-increment (Protocols IGMP) on page 4579
- source-increment (Protocols MLD) on page 4579
- spt-threshold on page 4580
- ssm-groups on page 4581
- ssm-map (Protocols IGMP AMT) on page 4582
- ssm-map (Protocols MLD) on page 4582
- ssm-map (Multicast Routing Options) on page 4583
- ssm-map (Protocols IGMP) on page 4583
- ssm-map-policy (IGMP) on page 4584

Copyright © 2016, Juniper Networks, Inc.

- ssm-map-policy (MLD) on page 4584
- static (Bridge Domains) on page 4585
- static (Protocols MLD) on page 4586
- static (Protocols IGMP) on page 4587
- static (Protocols PIM) on page 4588
- subscriber-leave-timer on page 4589
- threshold (Bridge Domains) on page 4590
- threshold (Protocols MSDP) on page 4591
- threshold (PIM BFD Detection Time) on page 4592
- threshold (PIM BFD Transmit Interval) on page 4593
- timeout (Flow Maps) on page 4594
- timeout (Multicast) on page 4595
- traceoptions (Multicast Snooping Options) on page 4596
- traceoptions (Protocols AMT) on page 4598
- traceoptions (Protocols DVMRP) on page 4601
- traceoptions (Protocols IGMP) on page 4604
- traceoptions (Protocols IGMP Snooping) on page 4607
- traceoptions (Protocols MLD) on page 4609
- traceoptions (Protocols MSDP) on page 4612
- traceoptions (Protocols PGM) on page 4615
- traceoptions (Protocols PIM) on page 4617
- transmit-interval (PIM BFD Liveness Detection) on page 4620
- tunnel-devices (Routing Instances) on page 4621
- tunnel-limit (Protocols AMT) on page 4622
- upstream-interface on page 4623
- version (BFD) on page 4624
- version (PIM) on page 4625
- version (Protocols IGMP) on page 4626
- version (Protocols IGMP AMT) on page 4626
- version (Protocols MLD) on page 4627
- vlan (Bridge Domains) on page 4628
- vpn-group-address on page 4629
- wildcard-source (PIM RPF Selection) on page 4629

Copyright © 2016, Juniper Networks, Inc.

accept-remote-source

Syntax accept-remote-source;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.

Description Accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This statement enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Interface to Accept Traffic from a Remote Source on page 4262

accounting (Protocols IGMP)

Syntax accounting;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable the collection of IGMP join and leave event statistics on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events on page 4056

Copyright © 2016, Juniper Networks, Inc.

accounting (Protocols IGMP AMT Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable or disable the collection of IGMP join and leave event statistics for an Automatic Multicast Tunneling (AMT) interface.

Default Disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

accounting (Protocols IGMP Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable or disable the collection of IGMP join and leave event statistics for an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events on page 4056

Copyright © 2016, Juniper Networks, Inc.

accounting (Protocols MLD)

Syntax accounting;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced in Junos OS Release 9.1.

Description Enable the collection of MLD join and leave event statistics on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Recording MLD Join and Leave Events on page 4082

accounting (Protocols MLD Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.1.

Description Enable or disable the collection of MLD join and leave event statistics for an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Recording MLD Join and Leave Events on page 4082

Copyright © 2016, Juniper Networks, Inc.

active-source-limit

Syntax

```
active-source-limit {
    maximum number;
    threshold number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* protocols msdp source *ip-address/prefix-length*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*],
[edit protocols msdp],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit protocols msdp source *ip-address/prefix-length*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp source *ip-address/prefix-length*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Limit the number of active source messages the routing device accepts.

Default

If you do not include this statement, the router accepts any number of MSDP active source messages.

Options

The options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262

Copyright © 2016, Juniper Networks, Inc.

address (Anycast RPs)

Syntax address *address* <forward-msdp-sa>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit protocols pim rp local (inet | inet6) anycast-pim rp-set],
[edit routing-instances *routing-instance-name* protocols pim rp local (inet | inet6) anycast-pim rp-set]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.

Options ***address***—RP address in an RP set.

forward-msdp-sa—(Optional) Forward MSDP SAs to this address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

address (Bidirectional Rendezvous Points)

Syntax

```
address address {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bidirectional],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bidirectional],
[edit protocols pim rp bidirectional],
[edit routing-instances *routing-instance-name* protocols pim rp bidirectional]

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Configure bidirectional rendezvous point (RP) addresses. The address can be a loopback interface address, an address of a link interface, or an address that is not assigned to an interface but belongs to a subnet that is reachable by the bidirectional PIM routers in the network.

Options

address—Bidirectional RP address.

Default: 232.0.0.0/8

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

address (Local RPs)

Syntax address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the local rendezvous point (RP) address.

Options ***address***—Local RP address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Local PIM RPs on page 4136

Copyright © 2016, Juniper Networks, Inc.

address (Static RPs)

Syntax

```
address address {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  override;
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp static],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp static],
[edit protocols pim static],
[edit routing-instances *routing-instance-name* protocols pim rp static]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

Options

address—Static RP address.

Default: 224.0.0.0/4

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4137

Copyright © 2016, Juniper Networks, Inc.

algorithm

Syntax algorithm *algorithm-name*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the algorithm to use for BFD authentication.

Options ***algorithm-name***—Name of algorithm to use for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional Forwarding Detection Authentication for PIM on page 4219
- Configuring BFD Authentication for PIM on page 4222
- authentication on page 4411

Copyright © 2016, Juniper Networks, Inc.

amt (IGMP)

Syntax

```
amt {
  relay {
    defaults {
      (accounting (Protocols IGMP AMT Interface) | no-accounting (Protocols IGMP AMT
        Interface));
      group-policy [ policy-names ];
      query-interval seconds;
      query-response-interval seconds;
      robust-count number;
      ssm-map ssm-map-name;
      version version;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp],
[edit protocols igmp],
[edit routing-instances *routing-instance-name* protocols igmp]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure Automatic Multicast Tunneling (AMT) relay attributes.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

Copyright © 2016, Juniper Networks, Inc.

amt (Protocols)

Syntax

```
amt {
  relay {
    accounting;
    family {
      inet {
        anycast-prefix ip-prefix</prefix-length>;
        local-address ip-address;
      }
    }
    secret-key-timeout minutes;
    tunnel-limit number;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Enable Automatic Multicast Tunneling (AMT) on the router or switch. You must also configure the local address and anycast prefix for AMT to function.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

anycast-pim

Syntax

```
anycast-pim {
  rp-set {
    address address <forward-msdp-sa>;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties for anycast RP using PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast With or Without MSDP on page 4142

Copyright © 2016, Juniper Networks, Inc.

anycast-prefix

Syntax anycast-prefix *ip-prefix/<prefix-length>*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols amt relay family inet],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay family inet],
[edit protocols amt relay family inet],
[edit routing-instances *routing-instance-name* protocols amt relay family inet]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify an IP address prefix to use for the Automatic Multicast Tunneling (AMT) relay anycast address. The prefix is advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways. The IP address that the prefix is derived from can be configured on any interface in the system. Typically, the router's **lo0.0** loopback address prefix is used for configuring the AMT anycast prefix in the default routing instance, and the router's **lo0.*n*** loopback address prefix is used for configuring the AMT anycast prefix in VPN routing instances. However, the anycast address can be either the primary or secondary **lo0.0** loopback address.

Default None. The anycast prefix must be configured.

Options ***ip-prefix/<prefix-length>***—IP address prefix.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

asm-override-ssm

Syntax asm-override-ssm;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Enable the routing device to accept any-source multicast join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast groups.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4187

Copyright © 2016, Juniper Networks, Inc.

assert-timeout

Syntax assert-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds.

Options ***seconds***—Time for routing device to wait before another assert message cycle.

Range: 5 through 210 seconds

Default: 180 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the PIM Assert Timeout on page 4175

 Copyright © 2016, Juniper Networks, Inc.

authentication

Syntax

```
authentication {
  algorithm algorithm-name;
  key-chain key-chain-name;
  loose-check;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD Authentication for PIM on page 4222
- Configuring BFD for PIM on page 4221
- Understanding Bidirectional Forwarding Detection Authentication for PIM on page 4219
- bfd-liveness-detection on page 4418
- key-chain on page 4490
- loose-check on page 4497

Copyright © 2016, Juniper Networks, Inc.

authentication-key

Syntax authentication-key *peer-key*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.

Default If you do not include this statement, the router accepts any valid MSDP messages from the peer address.

Options *peer-key*—MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (“ ”).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP in a Routing Instance on page 4254

Copyright © 2016, Juniper Networks, Inc.

auto-rp

Syntax

```
auto-rp {
  (announce | discovery | mapping);
  (mapping-agent-election | no-mapping-agent-election);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced in Junos OS Release 7.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure automatic RP announcement and discovery.

Options

announce—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.

discovery—Configure the routing device to listen only for mapping packets.

mapping—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Auto-RP on page 4152

Copyright © 2016, Juniper Networks, Inc.

backoff-period

Syntax backoff-period *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name* bidirectional df-election],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election],
[edit protocols pim interface *interface-name* bidirectional df-election],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election]

Release Information Statement introduced in Junos OS Release 12.1.

Description Configure the designated forwarder (DF) election backoff period for bidirectional PIM. The **backoff-period** statement configures the period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

NOTE: Junos OS checks rendezvous point (RP) unicast reachability before accepting incoming DF messages. DF messages for unreachable rendezvous points are ignored. This is needed to prevent the following example scenario. Routers A and B are downstream routers on the same LAN, and both are supposed to send DF election messages with an infinite metric on their upstream interfaces (reverse-path forwarding [RPF] interfaces). Router A has a higher IP address than Router B. When both routers lose the path to the RP, both send an Offer message with the infinite metric onto the LAN. Router A wins the election because it has a higher IP address, and Router B backs off as a result. After three Offer messages, according to RFC 5015, Router A looks up the RP and finds no path to the RP. As a result, Router A transitions to the Lose state and sends nothing. On the other hand, after backing off for an interval of 3 x the Offer period, Router B does not receive any messages, and resumes the DF election by sending a new Offer message. Hence, the pattern repeats indefinitely.

Options *milliseconds*—Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

Range: 100 through 65,535 milliseconds

Default: 1000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

backup-pe-group

Syntax

```
backup-pe-group group-name {
  backups [ addresses ];
  local-address address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

group-name—Name of the group for PE backups.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ingress PE Redundancy on page 4381

Copyright © 2016, Juniper Networks, Inc.

backups

Syntax backups [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options ***addresses***—Addresses of other PEs in the backup group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Ingress PE Redundancy on page 4381

Copyright © 2016, Juniper Networks, Inc.

bandwidth

Syntax bandwidth (*bps* | adaptive);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit logical-systems *logical-system-name* routing-options multicast flow-map],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit routing-options multicast flow-map]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the bandwidth property for multicast flow maps.

Options **adaptive**—Specify that the bandwidth is measured for the flows that are matched by the flow map.

bps—Bandwidth, in bits per second, for the flow map.

Range: 0 through any amount of bandwidth

Default: 2 Mbps

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Multicast Flow Map on page 4377

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (0 | 1 | automatic);
}
```

Hierarchy Level

[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.1.
authentication option introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- Configuring BFD Authentication for PIM on page 4222

Copyright © 2016, Juniper Networks, Inc.

bidirectional (Interface)

Syntax

```
bidirectional {
  df-election {
    backoff-period milliseconds;
    offer-period milliseconds;
    robustness-count number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Configure parameters for bidirectional PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

bidirectional (RP)

Syntax

```
bidirectional {
  address address {
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    priority number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Configure the routing device's rendezvous-point (RP) properties for bidirectional PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

bootstrap

Syntax

```
bootstrap {
  family (inet | inet6) {
    export [ policy-names ];
    import [ policy-names ];
    priority number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure parameters to control bootstrap routers and messages.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149

Copyright © 2016, Juniper Networks, Inc.

bootstrap-export

Syntax bootstrap-export [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more export policies to control outgoing PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- bootstrap-import on page 4423

Copyright © 2016, Juniper Networks, Inc.

bootstrap-import

Syntax bootstrap-import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more import policies to control incoming PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- **bootstrap-export on page 4422**

Copyright © 2016, Juniper Networks, Inc.

bootstrap-priority

Syntax bootstrap-priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.

Options
number—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.

Range: 0 through 255

Default: 0

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Bootstrap Properties for IPv4 on page 4147

Copyright © 2016, Juniper Networks, Inc.

data-encapsulation

Syntax data-encapsulation (disable | enable);

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.

Default If you do not include this statement, the RP encapsulates multicast data.

Options **disable**—(Optional) Do not use MSDP data encapsulation.

enable—Use MSDP data encapsulation.

Default: **enable**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262

Copyright © 2016, Juniper Networks, Inc.

default-peer

Syntax default-peer;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262

Copyright © 2016, Juniper Networks, Inc.

defaults

Syntax

```
defaults {
  (accounting (Protocols IGMP AMT Interface) | no-accounting (Protocols IGMP AMT
    Interface));
  group-policy [ policy-names ];
  query-interval seconds;
  query-response-interval seconds;
  robust-count number;
  ssm-map ssm-map-name;
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* statement-name protocols igmp amt relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* statement-name protocols igmp amt relay],
[edit protocols igmp amt relay],
[edit routing-instances *routing-instance-name* statement-name protocols igmp amt relay]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure default IGMP attributes for all Automatic Multicast Tunneling (AMT) interfaces.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

dense-groups

Syntax

```
dense-groups {
  addresses;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure which groups are operating in dense mode.

Options

addresses—Address of groups operating in dense mode.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Sparse-Dense Mode Properties on page 4111

Copyright © 2016, Juniper Networks, Inc.

detection-time (BFD for PIM)

Syntax

```
detection-time {
  threshold milliseconds;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418
- threshold on page 4592

Copyright © 2016, Juniper Networks, Inc.

df-election

Syntax

```
df-election {
  backoff-period milliseconds;
  offer-period milliseconds;
  robustness-count number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim interface *interface-name* bidirectional],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional],
[edit protocols pim interface *interface-name* bidirectional],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional]

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Optionally, configure the designated forwarder (DF) election parameters for bidirectional PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

disable (PIM)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim family (inet | inet6)],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit protocols pim],
[edit protocols pim family (inet | inet6)],
[edit protocols pim interface *interface-name*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before Junos OS Release 7.4.
disable statement extended to the **[family]** hierarchy level in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Explicitly disable PIM at the protocol, interface or family hierarchy levels.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling PIM on page 4100
- family (Protocols PIM) on page 4446

Copyright © 2016, Juniper Networks, Inc.

disable (PIM Graceful Restart)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim graceful-restart],
[edit protocols pim graceful-restart],
[edit routing-instances *routing-instance-name* protocols pim graceful-restart]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Explicitly disable PIM sparse mode graceful restart.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Sparse Mode Graceful Restart on page 4238

disable (Protocols DVMRP)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols dvmrp],
[edit logical-systems *logical-system-name* protocols dvmrp interface *interface-name*],
[edit protocols dvmrp],
[edit protocols dvmrp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Explicitly disable DVMRP on the system or on an interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring DVMRP to Announce Unicast Routes on page 4292

Copyright © 2016, Juniper Networks, Inc.

disable (Protocols IGMP)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Disable IGMP on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling IGMP on page 4060

disable (Protocols MLD)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*], [edit protocols mld interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Disable MLD on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling MLD on page 4085

Copyright © 2016, Juniper Networks, Inc.

disable (Protocols MSDP)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Explicitly disable MSDP.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling MSDP on page 4269

Copyright © 2016, Juniper Networks, Inc.

disable (Protocols SAP)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols sap],
[edit protocols sap]

Release Information Statement introduced before Junos OS Release 7.4.

Description Explicitly disable SAP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Session Announcement Protocol on page 4271

dr-election-on-p2p

Syntax dr-election-on-p2p;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable PIM designated router (DR) election on point-to-point (P2P) links.

Default No PIM DR election is performed on point-to-point links.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Designated Router Election on Point-to-Point Links on page 4104

Copyright © 2016, Juniper Networks, Inc.

dr-register-policy

Syntax dr-register-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to control outgoing PIM register messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Register Message Filters on a PIM RP and DR on page 4164
- rp-register-policy on page 4566

Copyright © 2016, Juniper Networks, Inc.

dvmrp

Syntax

```
dvmrp {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  interface interface-name {
    disable;
    hold-time seconds;
    metric metric;
    mode (forwarding | unicast-routing);
  }
  rib-group group-name;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable DVMRP on the router.

Default DVMRP is disabled on the router.

Options The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring DVMRP on page 4288

Copyright © 2016, Juniper Networks, Inc.

embedded-rp

Syntax

```
embedded-rp {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  maximum-rps limit;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties for embedded IP version 6 (IPv6) RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Embedded RP for IPv6 on page 4158

exclude (Protocols IGMP)

Syntax

```
exclude;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

Copyright © 2016, Juniper Networks, Inc.

exclude (Protocols MLD)

Syntax exclude;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. By default, the group operates in include mode.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Enabling MLD Static Group Membership on page 4075

export (Protocols DVMRP)

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp],
[edit protocols dvmrp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply one or more policies to routes being exported from the routing table into DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, the routing table exports into DVMRP only the routes that it learned from DVMRP and direct routes.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- import (Protocols DVMRP) on page 4476
- Example: Configuring DVMRP to Announce Unicast Routes on page 4292

Copyright © 2016, Juniper Networks, Inc.

export (Protocols MSDP)

Syntax

export [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more policies to routes being exported from the routing table into MSDP.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP in a Routing Instance on page 4254
- import on page 4477

Copyright © 2016, Juniper Networks, Inc.

export (Protocols PIM)

Syntax

export [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Filtering Outgoing PIM Join Messages on page 4162

Copyright © 2016, Juniper Networks, Inc.

export (Bootstrap)

Syntax

export [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)],
[edit protocols pim rp bootstrap family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap family (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- import (Protocols PIM Bootstrap) on page 4479

Copyright © 2016, Juniper Networks, Inc.

family (Bootstrap)

Syntax

```
family (inet | inet6) {
  export [ policy-names ];
  import [ policy-names ];
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap],
[edit protocols pim rp bootstrap],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure which IP protocol type bootstrap properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149

Copyright © 2016, Juniper Networks, Inc.

family (Local RP)

Syntax

```
family (inet | inet6) {
  disable;
  address address;
  anycast-pim {
    local-address address;
    rp-set {
      address address <forward-msdp-sa>;
    }
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  override;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp local],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local],
[edit protocols pim rp local],
[edit routing-instances *routing-instance-name* protocols pim rp local]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure which IP protocol type local RP properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 4136

Copyright © 2016, Juniper Networks, Inc.

family (Protocols AMT Relay)

Syntax

```
family {
  inet {
    anycast-prefix ip-prefix/<prefix-length>;
    local-address ip-address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols amt relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay],
[edit protocols amt relay],
[edit routing-instances *routing-instance-name* protocols amt relay]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure the protocol address family for Automatic Multicast Tunneling (AMT) relay functions. Only the **inet** family for IPv4 protocol addresses is supported.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

family (Protocols PIM)

Syntax

```
family (inet | inet6) {
    disable;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Disable the PIM protocol for the specified family.

Options

inet—Disable the PIM protocol for the IP version 4 (IPv4) address family.

inet6—Disable the PIM protocol for the IP version 6 (IPv6) address family.

Related Documentation

- Disabling PIM on page 4100
- disable (PIM Graceful Restart) on page 4432
- disable (PIM) on page 4431

Copyright © 2016, Juniper Networks, Inc.

flood-groups

Syntax flood-groups [*ip-addresses*];

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Establish a list of flood group addresses for multicast snooping.

Options ***ip-addresses***—List of IP addresses subject to flooding.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast Snooping on page 4345

Copyright © 2016, Juniper Networks, Inc.

flow-map

Syntax

```
flow-map flow-map-name {
  bandwidth (bps | adaptive);
  forwarding-cache {
    timeout (never non-discard-entry-only | minutes);
  }
  policy [ policy-names ];
  redundant-sources [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure multicast flow maps.

Options

flow-map-name—Name of the flow-map.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Multicast Flow Map on page 4377

Copyright © 2016, Juniper Networks, Inc.

forwarding-cache (Bridge Domains)

Syntax

```
forwarding-cache {
    threshold suppress value <reuse value>;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Establish multicast snooping forwarding cache parameter values.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast Snooping on page 4345

forwarding-cache (Flow Maps)

Syntax

```
forwarding-cache {
    timeout (minutes | never non-discard-entry-only );
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure multicast forwarding cache properties for the flow map.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Multicast Flow Map on page 4377

Copyright © 2016, Juniper Networks, Inc.

forwarding-cache (Multicast)

Syntax

```
forwarding-cache {
    threshold suppress value <reuse value>;
    timeout minutes;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits and timeout values.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the Multicast Forwarding Cache on page 4374

Copyright © 2016, Juniper Networks, Inc.

graceful-restart (Multicast Snooping)

Syntax graceful-restart <restart-duration *seconds*>;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Establish the graceful restart duration for multicast snooping. You can set this value between 0 and 300 seconds. If you set the duration to 0, graceful restart is effectively disabled. Set this value slightly larger than the IGMP query response interval.

Default 180 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast Snooping on page 4345
- query-response-interval (Bridge Domains) on page 4548

Copyright © 2016, Juniper Networks, Inc.

graceful-restart (Protocols PIM)

Syntax

```
graceful-restart {
    disable;
    no-bidirectional-mode;
    restart-duration seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure PIM sparse mode graceful restart.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Sparse Mode Graceful Restart on page 4238

Copyright © 2016, Juniper Networks, Inc.

group (Bridge Domains)

Syntax

```
group ip-address {
    source-address ip-address;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name* static],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name* static]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure the IGMP multicast group address that receives data on an interface and (optionally) a source address for certain packets.

Options

ip-address—Group address.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

group (Protocols IGMP)

Syntax

```
group multicast-group-address {
  exclude;
  group-count number;
  group-increment increment;
  source ip-address {
    source-count number;
    source-increment increment;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static], [edit protocols igmp interface *interface-name* static]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

NOTE: You must specify a unique address for each group.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

Copyright © 2016, Juniper Networks, Inc.

group (Protocols MLD)

Syntax

```
group multicast-group-address {
  exclude;
  group-count number;
  group-increment increment;
  source ip-address {
    source-count number;
    source-increment increment;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static], [edit protocols mld interface *interface-name* static]

Release Information Statement introduced before Junos OS Release 7.4.

Description The MLD multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

Options ***multicast-group-address***—Address of the group.

NOTE: You must specify a unique address for each group.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

group (Protocols MSDP)

Syntax

```
group group-name {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  mode (mesh-group | standard);
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  peer address; {
    disable;
    active-source-limit {
      maximum number;
      threshold number;
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the peer statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

Options

group-name—Name of the MSDP group.

The remaining statements are explained separately.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring MSDP in a Routing Instance on page 4254

group (RPF Selection)

Syntax

```
group group-address{
  sourcesource-address{
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level
[edit routing-instances *routing-instance-name* edit protocols pim rpf-selection]

Release Information
Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Configure the PIM group address for which you configure RPF selectiongroup (RPF Selection).

Default
By default, PIM RPF selection is not configured.

Options
group-address—PIM group address for which you configure RPF selection.

Required Privilege Level
view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

group-count

Syntax group-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the number of static groups to be created.

Options *number*—Number of static groups.

Default:

Range: 1 through 512

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling IGMP Static Group Membership on page 4050

group-count (Protocols MLD)

Syntax group-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the number of static groups to be created.

Options *number*—Number of static groups.

Default: 1

Range: 1 through 512

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

group-increment (Protocols IGMP)

Syntax group-increment *increment*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options *increment*—Number of times the address should be incremented.

Default: 0.0.0.1

Range: 0.0.0.1 through 255.255.255.255

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling IGMP Static Group Membership on page 4050

group-increment (Protocols MLD)

Syntax group-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options *increment*—Number of times the address should be incremented.

Default: ::1

Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

group-limit (Protocols IGMP)

Syntax group-limit *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

Default By default, there is no limit to the number of multicast groups that can join the interface.

Options *limit*—group limit value for the interface.

Range: 1 through 32767

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4057

Copyright © 2016, Juniper Networks, Inc.

group-limit (IGMP and MLD Snooping)

Syntax group-limit *limit*;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

Default By default, there is no limit to the number of multicast groups joining an interface.

Options ***limit***—a 32-bit number for the limit on the interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335

group-limit (Protocols MLD)

Syntax group-limit *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure a limit for the number of multicast groups (or [S,G] channels in MLDv2) allowed on a logical interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.

Default By default, there is no limit to the number of multicast groups that can join the interface.

Options *limit*—group value limit for the interface.
Range: 1 through 32767

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Examples: Configuring MLD on page 4063

group-policy (Protocols IGMP)

Syntax group-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4045

Copyright © 2016, Juniper Networks, Inc.

group-policy (Protocols IGMP AMT Interface)

Syntax group-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description When this statement is enabled on the Automatic Multicast Tunneling (AMT) interfaces running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Options ***policy-names***—Name of the policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

group-policy (Protocols MLD)

Syntax group-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.1.

Description When a router running MLD version 1 or version 2 (MLDv1 or MLDv2), receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Filtering Unwanted MLD Reports at the MLD Interface Level on page 4072

Copyright © 2016, Juniper Networks, Inc.

group-ranges

Syntax

```
group-ranges {
  destination-ip-prefix</prefix-length>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* protocols pim rp embedded-rp],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp embedded-rp],
[edit protocols pim rp local family (inet | inet6)],
[edit protocols pim rp static address *address*],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp static address *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 13.3 for the PTX5000 router.

Description

Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).

Default

The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).

Options

destination-ip-prefix</prefix-length>—Addresses or address ranges for which this routing device can be an RP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 4136
- Configuring PIM Embedded RP for IPv6 on page 4158
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify how often the routing device sends PIM hello packets out of an interface.

Options *seconds*—Length of time between PIM hello packets.

Range: 0 through 255

Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- hold-time (Protocols PIM) on page 4466
- Modifying the PIM Hello Interval on page 4096

hold-time (Protocols DVMRP)

Syntax hold-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp interface *interface-name*],
[edit protocols dvmrp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the time period for which a neighbor is to consider the sending router (this router) to be operative (up).

Options *seconds*—Hold time.

Range: 1 through 255

Default: 35 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring DVMRP on page 4288

Copyright © 2016, Juniper Networks, Inc.

hold-time (Protocols PIM)

Syntax hold-time *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

Description Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).

Options *seconds*—Hold time.

Range: 0 through 255

Default: 150 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Local PIM RPs on page 4136
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

host-only-interface

Syntax host-only-interface;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.

Default The interface can either be a host-side or multicast-router interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335
- multicast-router-interface on page 4513

Copyright © 2016, Juniper Networks, Inc.

igmp

Syntax

```
igmp {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-limit limit;
    group-policy [ policy-names ];
    immediate-leave;
    oif-map map-name;
    passive;
    promiscuous-mode;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Enable IGMP on the router. IGMP must be enabled for the router to receive multicast packets.

The remaining statements are explained separately.

Default IGMP is disabled on the router. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling IGMP on page 4041

igmp-snooping

Syntax

```
igmp-snooping {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  vlan vlan-id {
    immediate-leave;
    interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      multicast-router-interface;
      static {
        group ip-address {
          source ip-address;
        }
      }
    }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]
[edit routing-instances *routing-instance-name* protocols]
[edit protocols]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Enable IGMP snooping on the router.

Copyright © 2016, Juniper Networks, Inc.

Default IGMP snooping is disabled on the router.

Options The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Understanding IGMP Snooping on page 4330

ignore-stp-topology-change

Syntax ignore-stp-topology-change;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Ignore messages about spanning tree topology changes. This statement is supported for the **virtual-switch** routing instance type only.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast Snooping on page 4345

Copyright © 2016, Juniper Networks, Inc.

immediate-leave (Bridge Domains)

Syntax immediate-leave;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335

immediate-leave (Protocols IGMP)

Syntax immediate-leave;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Specifying Immediate-Leave Host Removal for IGMP on page 4044

Copyright © 2016, Juniper Networks, Inc.

immediate-leave (Protocols MLD)

Syntax immediate-leave;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*], [edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.

Description The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.

NOTE: Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Specifying Immediate-Leave Host Removal for MLD on page 4071

Copyright © 2016, Juniper Networks, Inc.

import (Protocols DVMRP)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp],
[edit protocols dvmrp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply one or more policies to routes being imported into the routing table from DVMRP. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, DVMRP shares with the routing table only those routes that were learned from DVMRP routers.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- export (Protocols DVMRP) on page 4439
- Example: Configuring DVMRP to Announce Unicast Routes on page 4292

Copyright © 2016, Juniper Networks, Inc.

import (Protocols MSDP)

Syntax

import [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more policies to routes being imported into the routing table from MSDP.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP in a Routing Instance on page 4254
- export (Protocols MSDP) on page 4440

Copyright © 2016, Juniper Networks, Inc.

import (Protocols PIM)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply one or more policies to routes being imported into the routing table from PIM. Use the **import** statement to filter PIM join messages and prevent them from entering the network.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Filtering Incoming PIM Join Messages on page 4163

 Copyright © 2016, Juniper Networks, Inc.

import (Protocols PIM Bootstrap)

Syntax

import [*policy-names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more import policies to control incoming PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- **export (Bootstrap) on page 4442**

Copyright © 2016, Juniper Networks, Inc.

infinity

Syntax infinity [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim spt-threshold],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim spt-threshold],
[edit protocols pim spt-threshold],
[edit routing-instances *routing-instance-name* protocols pim spt-threshold]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the **infinity** statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the PIM SPT Threshold Policy on page 4177

 Copyright © 2016, Juniper Networks, Inc.

inet

Syntax

```
inet {
  anycast-prefix ip-prefix/<prefix-length>;
  local-address ip-address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols amt relay family],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay family],
[edit protocols amt relay family],
[edit routing-instances *routing-instance-name* protocols amt relay family]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Specify the IPv4 local address and anycast prefix for Automatic Multicast Tunneling (AMT) relay functions.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

interface (Bridge Domains)

Syntax

```
interface interface-name {
  group-limit limit;
  host-only-interface;
  multicast-router-interface;
  static {
    group ip-address {
      source ip-address;
    }
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Enable IGMP snooping on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

interface (Protocols DVMRP)

Syntax

```
interface interface-name {
  disable;
  hold-time seconds;
  metric metric;
  mode (forwarding | unicast-routing);
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp],
[edit protocols dvmrp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable DVMRP on an interface and configure interface-specific properties.

Options ***interface-name***—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring DVMRP on page 4288

Copyright © 2016, Juniper Networks, Inc.

interface (Protocols IGMP)

Syntax

```
interface interface-name {
  disable;
  (accounting | no-accounting);
  group-limit limit;
  group-policy [ policy-names ];
  immediate-leave;
  oif-map map-name;
  passive;
  promiscuous-mode;
  ssm-map ssm-map-name;
  ssm-map-policy ssm-map-policy-name;
  static {
    group multicast-group-address {
      exclude;
      group-count number;
      group-increment increment;
      source ip-address {
        source-count number;
        source-increment increment;
      }
    }
  }
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable IGMP on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP on page 4041

Copyright © 2016, Juniper Networks, Inc.

interface (Protocols MLD)

Syntax

```
interface interface-name {
  disable;
  (accounting | no-accounting);
  group-limit limit;
  group-policy [ policy-names ];
  group-threshold value;
  immediate-leave;
  log-interval seconds;
  oif-map [ map-names ];
  passive;
  ssm-map ssm-map-name;
  ssm-map-policy ssm-map-policy-name;
  static {
    group multicast-group-address {
      exclude;
      group-count number
      group-increment increment
      source ip-address {
        source-count number;
        source-increment increment;
      }
    }
  }
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable MLD on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD on page 4067

Copyright © 2016, Juniper Networks, Inc.

interface (Protocols PIM)

Syntax

```
interface (Protocols PIM) (all | interface-name) {
  accept-remote-source;
  disable;
  bfd-liveness-detection {
    authentication {
      algorithmalgorithm-name;
      key-chainkey-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  bidirectional {
    df-election {
      backoff-period milliseconds;
      offer-period milliseconds;
      robustness-count number;
    }
  }
  family (inet | inet6) {
    disable;
  }
  hello-interval seconds;
  mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
  neighbor-policy [ policy-names ];
  override-interval milliseconds;
  priority number;
  propagation-delay milliseconds;
  reset-tracking-bit;
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Enable PIM on an interface and configure interface-specific properties.

Copyright © 2016, Juniper Networks, Inc.

Options *interface-name*—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- PIM on Aggregated Interfaces on page 4098

interface (Routing Options)

Syntax

```
interface interface-names {
  maximum-bandwidth bps;
  no-qos-adjust;
  reverse-oif-mapping {
    no-qos-adjust;
  }
  subscriber-leave-timer seconds;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable multicast traffic on an interface.

TIP: You cannot enable multicast traffic on an interface by using the **routing-options multicast interface** statement and configure PIM on the interface.

Options *interface-name*—Names of the physical or logical interface.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Defining Interface Bandwidth Maximums on page 4355
- Example: Configuring Multicast with Subscriber VLANs on page 4358

Copyright © 2016, Juniper Networks, Inc.

interface (Scoping)

Syntax interface [*interface-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the set of interfaces for multicast scoping.

Options ***interface-names***—Names of the interfaces to scope. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast Snooping on page 4345

 Copyright © 2016, Juniper Networks, Inc.

join-load-balance

Syntax

```
join-load-balance {
  automatic;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable load balancing of PIM join messages across interfaces and routing devices.

Options

automatic—Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Join Load Balancing on page 4119
- clear pim join-distribution on page 4650 in the CLI Explorer

Copyright © 2016, Juniper Networks, Inc.

join-prune-timeout

Syntax join-prune-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.

Options **seconds**—Number of seconds to wait for the periodic join message to arrive.

Range: 210 through 240 seconds

Default: 210 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Modifying the Join State Timeout on page 4122

key-chain

Syntax key-chain *key-chain-name*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the security keychain to use for BFD authentication.

Options ***key-chain-name***—Name of the security keychain to use for BFD authentication. The name is a unique integer between **0** and **63**. This must match one of the keychains in the **authentication-key-chains** statement at the **[edit security]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BFD Authentication for PIM on page 4222
- Understanding Bidirectional Forwarding Detection Authentication for PIM on page 4219
- **authentication on page 4411**

Copyright © 2016, Juniper Networks, Inc.

listen

Syntax listen *address* <port *port*>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols sap],
[edit protocols sap]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify an address and optionally a port on which SAP and SDP listen, in addition to the default SAP address and port on which they always listen, 224.2.127.254:9875. To specify multiple additional addresses or pairs of address and port, include multiple **listen** statements.

Options ***address***—(Optional) Address on which SAP listens for session advertisements.

Default: 224.2.127.254

port *port*—(Optional) Port on which SAP listens for session advertisements.

Default: 9875

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Session Announcement Protocol on page 4271

Copyright © 2016, Juniper Networks, Inc.

local

Syntax

```
local {
  disable;
  address address;
  family (inet | inet6) {
    disable;
    address address;
    anycast-pim {
      local-address address;
        rp-set {
        address address <forward-msdp-sa>;
      }
    }
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  override;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
The remaining statements are explained separately.

Description

Configure the routing device's RP properties.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 4136

Copyright © 2016, Juniper Networks, Inc.

local-address (Protocols AMT)

Syntax local-address *ip-address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols amt relay family inet],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay family inet],
[edit protocols amt relay family inet],
[edit routing-instances *routing-instance-name* protocols amt relay family inet]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the local unique IP address to send in Automatic Multicast Tunneling (AMT) relay advertisement messages, for use as the IP source of AMT control messages, and as the source of the data tunnel encapsulation. The address can be configured on any interface in the system. Typically, the router's **lo0.0** loopback address is used for configuring the AMT local address in the default routing instance, and the router's **lo0.*n*** loopback address is used for configuring the AMT local address in VPN routing instances.

Default None. The local address must be configured.

Options ***ip-address***—Unique unicast IP address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

local-address (Protocols MSDP)

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.

Options ***address***—IP address of the local end of the connection.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring MSDP in a Routing Instance on page 4254

Copyright © 2016, Juniper Networks, Inc.

local-address (Protocols PIM)

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim],
[edit protocols pim rp local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6) anycast-pim]

Release Information
Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.

Options ***address***—Anycast RP IPv4 or IPv6 address, depending on **family** configuration.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM Anycast With or Without MSDP on page 4142

Copyright © 2016, Juniper Networks, Inc.

local-address (Routing Options)

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information
Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options ***address***—Address of local PEs in the backup group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Ingress PE Redundancy on page 4381

Copyright © 2016, Juniper Networks, Inc.

loose-check

Syntax loose-check;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection authentication],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection authentication]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD Authentication for PIM on page 4222
- Understanding Bidirectional Forwarding Detection Authentication for PIM on page 4219
- **authentication on page 4411**

Copyright © 2016, Juniper Networks, Inc.

mapping-agent-election

Syntax (mapping-agent-election | no-mapping-agent-election);

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp auto-rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp auto-rp],
[edit protocols pim rp auto-rp],
[edit routing-instances *routing-instance-name* protocols pim rp auto-rp]

Release Information Statement introduced in Junos OS Release 7.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device mapping announcements as a mapping agent.

Options

mapping-agent-election—Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.

no-mapping-agent-election—Mapping agents always announce mappings and do not perform mapping agent election.

Default: **mapping-agent-election**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Auto-RP on page 4152

Copyright © 2016, Juniper Networks, Inc.

maximum (MSDP Active Source Messages)

Syntax maximum *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp active-source-limit],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances *routing-instance-name* protocols msdp active-source-limit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the maximum number of MSDP active source messages the router accepts.

Options *number*—Maximum number of active source messages.

Range: 1 through 1,000,000

Default: 25,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262
- threshold (Protocols MSDP) on page 4591

Copyright © 2016, Juniper Networks, Inc.

maximum-bandwidth

Syntax maximum-bandwidth *bps*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the multicast bandwidth for the interface.

Options *bps*—Bandwidth rate, in bits per second, for the multicast interface.

Range: 0 through any amount of bandwidth

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Defining Interface Bandwidth Maximums on page 4355

Copyright © 2016, Juniper Networks, Inc.

maximum-rps

Syntax maximum-rps *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp embedded-rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp embedded-rp],
[edit protocols pim rp embedded-rp],
[edit routing-instances *routing-instance-name* protocols pim rp embedded-rp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Limit the number of RPs that the routing device acknowledges.

Options *limit*—Number of RPs.

Range: 1 through 500

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Embedded RP for IPv6 on page 4158

maximum-transmit-rate (Protocols IGMP)

Syntax maximum-transmit-rate *packets-per-second*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Limit the transmission rate of IGMP packets

Options **packets-per-second**—Maximum number of IGMP packets transmitted in one second by the router.

Range: 1 through 10000

Default: 500 packets

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Limiting the Maximum IGMP Message Rate on page 4049

Copyright © 2016, Juniper Networks, Inc.

metric (Protocols DVMRP)

Syntax metric *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp interface *interface-name*],
[edit protocols dvmrp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the DVMRP metric value.

Options *metric*—Metric value.

Range: 1 through 31

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring DVMRP on page 4288

minimum-interval (PIM BFD Liveness Detection)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Options *milliseconds*—Minimum transmit and receive interval.

Range: 1 through 255,000 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221

Copyright © 2016, Juniper Networks, Inc.

minimum-interval (PIM BFD Transmit Interval)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection transmit-interval],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection transmit-interval]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit protocols pim interface *interface-name* bfd-liveness-detection]** hierarchy level.

Options ***milliseconds***—Minimum transmit interval value.

Range: 1 through 255,000

NOTE: The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418
- minimum-interval on page 4502
- threshold on page 4593

Copyright © 2016, Juniper Networks, Inc.

minimum-receive-interval

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit protocols pim interface *interface-name* bfd-liveness-detection]** hierarchy level.

Options ***milliseconds***—Minimum receive interval.

Range: 1 through 255,000 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221

Copyright © 2016, Juniper Networks, Inc.

mld

Syntax

```
mld {
  accounting;
  interface interface-name {
    (accounting | no-accounting);
    disable;
    group-limit limit;
    group-policy [ policy-names ];
    immediate-leave;
    oif-map [ map-names ];
    passive;
    ssm-map ssm-map-name;
    ssm-map-policy ssm-map-policy-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
    version version;
  }
  maximum-transmit-rate packets-per-second;
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable MLD on the router. MLD must be enabled for the router to receive multicast packets.

Default MLD is disabled on the router. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options The statements are explained separately.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling MLD on page 4067

mode (Protocols DVMRP)

Syntax
```
mode (forwarding | unicast-routing);
```

Hierarchy Level
[edit logical-systems *logical-system-name* protocols dvmrp interface *interface-name*],
[edit protocols dvmrp interface *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Configure DVMRP for multicast traffic forwarding or unicast routing.

Options
forwarding—DVMRP performs unicast routing as well as multicast data forwarding.

unicast-routing—DVMRP performs unicast routing only. To forward multicast data, you must configure Protocol Independent Multicast (PIM) on the interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring DVMRP to Announce Unicast Routes on page 4292

Copyright © 2016, Juniper Networks, Inc.

mode (Protocols MSDP)

Syntax mode (mesh-group | standard);

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is **standard**.

Default If you do not include this statement, default flooding is applied.

Options **mesh-group**—Group of peers that are mesh group members.

standard—Use standard MSDP source-active flooding rules.

Default: standard

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262

Copyright © 2016, Juniper Networks, Inc.

mode (Protocols PIM)

Syntax mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
bidirectional-sparse and **bidirectional-sparse-dense** options introduced in Junos OS Release 12.1.

Description Configure the PIM mode on the interface.

Options The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:

- **bidirectional-sparse**—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.
- **bidirectional-sparse-dense**—Use if multicast groups, except those that are specified in the **dense-groups** statement, are operating in bidirectional, sparse, or SSM mode.
- **dense**—Use if all multicast groups are operating in dense mode.
- **sparse**—Use if all multicast groups are operating in sparse mode or SSM mode.
- **sparse-dense**—Use if multicast groups, except those that are specified in the **dense-groups** statement, are operating in sparse mode or SSM mode.

Default: Sparse mode

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Dense Mode Properties on page 4109
- Configuring PIM Sparse-Dense Mode Properties on page 4111
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

msdp

Syntax

```
msdp {
    disable;
    active-source-limit {
        log-interval seconds;
        log-warning value;
        maximum number;
        threshold number;
    }
    data-encapsulation (disable | enable);
    export [ policy-names ];
    group group-name {
        ... group-configuration ...
    }
    hold-time seconds;
    import [ policy-names ];
    local-address address;
    keep-alive seconds;
    peer address {
        ... peer-configuration ...
    }
    rib-group group-name;
    source ip-prefix</prefix-length> {
        active-source-limit {
            maximum number;
            threshold number;
        }
    }
    sa-hold-time seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    group group-name {
        disable;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        mode (mesh-group | standard);
        peer address {
            ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
                just following ...
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
    peer address {
        disable;
        active-source-limit {
            maximum number;
            threshold number;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.4 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.

Default

MSDP is disabled on the router or switch.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP in a Routing Instance on page 4254

Copyright © 2016, Juniper Networks, Inc.

multicast (Dynamic Profiles Routing Options)

Syntax

```
multicast {
    asm-override-ssm;
    backup-pe-group group-name {
        backups [ addresses ];
        local-address address;
    }
    flow-map flow-map-name {
        bandwidth (bps | adaptive);
        forwarding-cache {
            timeout (never non-discard-entry-only | minutes);
        }
        policy [ policy-names ];
        redundant-sources [ addresses ];
    }
    forwarding-cache {
        threshold suppress value <reuse value>;
        timeout minutes;
    }
    interface interface-name {
        maximum-bandwidth bps;
        no-qos-adjust;
        reverse-oif-mapping {
            no-qos-adjust;
        }
        subscriber-leave-timer seconds;
    }
    pim-to-igmp-proxy {
        upstream-interface [ interface-names ];
    }
    pim-to-mld-proxy {
        upstream-interface [ interface-names ];
    }
    rpf-check-policy [ policy-names ];
    scope scope-name {
        interface [ interface-names ];
        prefix destination-prefix;
    }
    scope-policy [ policy-names ];
    ssm-groups [ addresses ];
    ssm-map ssm-map-name {
        policy [ policy-names ];
        source [ addresses ];
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <disable>;
    }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-options],
[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options],

Copyright © 2016, Juniper Networks, Inc.

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

NOTE: You cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the **scope** statement does apply individually to a specific routing instance.

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
interface and **maximum-bandwidth** statements introduced in Junos OS Release 8.3.
interface and **maximum-bandwidth** statements introduced in Junos OS Release 9.0 for EX Series switches.
Statement added to **[edit dynamic-profiles routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options]** hierarchy levels in Junos OS Release 9.6.

Description

Configure multicast routing options properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the Multicast Forwarding Cache on page 4374
- Example: Configuring a Multicast Flow Map on page 4377
- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4187

Copyright © 2016, Juniper Networks, Inc.

multicast-router-interface

Syntax multicast-router-interface;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure an interface as a bridge interface toward other multicast routers.

Default The interface can either be a host-side or multicast-router interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335
- host-only-interface on page 4467

Copyright © 2016, Juniper Networks, Inc.

multicast-snooping-options

Syntax

```
multicast-snooping-options {
  flood-groups [ ip-addresses ];
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  graceful-restart <restart-duration seconds>;
  ignore-stp-topology-change;
  multichassis-lag-replicate-state;
  nexthop-hold-time milliseconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Establish multicast snooping option values.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Multicast Snooping on page 4344
- Enabling Bulk Updates for Multicast Snooping on page 4350
- Example: Configuring Multicast Snooping on page 4345

Copyright © 2016, Juniper Networks, Inc.

multichassis-lag-replicate-state

Syntax multichassis-lag-replicate-state;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options],
[edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 10.2.

Description Provide multicast snooping for multichassis link aggregation group interfaces. Replicate IGMP join and leave messages from the active link to the standby link of a dual-link multichassis link aggregation group interface, enabling faster recovery of membership information after failover.

Default If not included, membership information is recovered using a standard IGMP network query.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Multicast Snooping on page 4344
- multicast-snooping-options on page 4514

Copyright © 2016, Juniper Networks, Inc.

multiplier

Syntax multiplier *number*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Options *number*—Number of hello packets.

Range: 1 through 255

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BFD for PIM on page 4221

neighbor-policy

Syntax neighbor-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply a PIM interface-level policy to filter neighbor IP addresses.

Options *policy-name*—Name of the policy that filters neighbor IP addresses.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Interface-Level PIM Neighbor Policies on page 4161

Copyright © 2016, Juniper Networks, Inc.

nexthop-hold-time

Syntax nexthop-hold-time *milliseconds*;

Hierarchy Level [edit routing-instances *routing-instance-name* multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description Accumulate outgoing interface changes in order to perform bulk updates to the forwarding table and the routing table. Delete the statement to turn off bulk updates.

Options **milliseconds**—Set the hold time duration from 1 through 1000 milliseconds.
Range: 1 through 1000 milliseconds.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling Bulk Updates for Multicast Snooping on page 4350

next-hop (PIM RPF Selection)

Syntax next-hop *next-hop-address*;

Hierarchy Level [edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* wildcard-source],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* wildcard-source]

Release Information Statement introduced in JUNOS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the specific next-hop address for the PIM group source.

Options ***next-hop-address***—Specific next-hop address for the PIM group source.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

no-adaptation (PIM BFD Liveness Detection)

Syntax no-adaptation;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure BFD sessions not to adapt to changing network conditions. We recommend that you *do not* disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418

Copyright © 2016, Juniper Networks, Inc.

no-bidirectional-mode

Syntax no-birectional-mode;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim graceful-restart],
[edit protocols pim graceful-restart],
[edit routing-instances *routing-instance-name* protocols pim graceful-restart]

Release Information Statement introduced in Junos OS Release 12.1.

Description Disable forwarding for bidirectional PIM routes during graceful restart recovery, both in cases of a routing protocol process (rpd) restart and graceful Routing Engine switchover.

Bidirectional PIM accepts packets for a bidirectional route on multiple interfaces. This means that some topologies might develop multicast routing loops if all PIM neighbors are not synchronized with regard to the identity of the designated forwarder (DF) on each link. If one router is forwarding without actively participating in DF elections, particularly after unicast routing changes, multicast routing loops might occur.

If graceful restart for PIM is enabled and the forwarding of packets on bidirectional routes is disallowed (by including the **no-bidirectional-mode** statement in the configuration), PIM behaves conservatively to avoid multicast routing loops during the recovery period. When the routing protocol process (rpd) restarts, all bidirectional routes are deleted. After graceful restart has completed, the routes are re-added, based on the converged unicast and bidirectional PIM state. While graceful restart is active, bidirectional multicast flows drop packets.

Default If graceful restart for PIM is enabled and the bidirectional PIM is enabled, the default graceful restart behavior is to continue forwarding packets on bidirectional routes. If the gracefully restarting router was serving as a DF for some interfaces to rendezvous points, the restarting router sends a DF Winner message with a metric of 0 on each of these RP interfaces. This ensures that a neighbor router does not become the DF due to unicast topology changes that might occur during the graceful restart period. Sending a DF Winner message with a metric of 0 prevents another PIM neighbor from assuming the DF role until after graceful restart completes. When graceful restart completes, the gracefully restarted router sends another DF Winner message with the actual converged unicast metric.

NOTE: Graceful Routing Engine switchover operates independently of the graceful restart behavior. If graceful Routing Engine switchover is configured without graceful restart, all PIM routes for all modes are deleted when the rpd process restarts. If graceful Routing Engine switchover is configured with graceful restart, the behavior is the same as described here, except that the recovery happens on the Routing Engine that assumes mastership.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Sparse Mode Graceful Restart on page 4238
- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

no-multicast-echo

Syntax
no-multicast-echo

Hierarchy Level
[edit system]

Release Information
Statement introduced in Junos OS Release 8.1.

Description
Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

Default
The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Configuring the Junos OS to Disable the Routing Engine Response to Multicast Ping Packets*

Copyright © 2016, Juniper Networks, Inc.

no-qos-adjust

Syntax no-qos-adjust;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name* reverse-oif-mapping]

Release Information
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement added to **[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, and **[edit routing-options multicast interface *interface-name*]** hierarchy levels in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast with Subscriber VLANs on page 4358

Copyright © 2016, Juniper Networks, Inc.

offer-period

Syntax offer-period *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name* bidirectional df-election],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election],
[edit protocols pim interface *interface-name* bidirectional df-election],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election]

Release Information Statement introduced in Junos OS Release 12.1.

Description Configure the designated forwarder (DF) election offer period for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The **offer-period** statement modifies the interval between repeated DF election messages. The robustness-count statement determines the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routers on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.

If a router receives from a neighbor a better offer than its own, the router stops participating in the election for a period of robustness-count * **offer-period**. Eventually, all routers except the best candidate stop sending Offer messages.

Options *milliseconds*—Interval to wait before retransmitting DF Offer and Winner messages.

Range: 100 through 10,000 milliseconds

Default: 100

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205
- **robustness-count on page 4563**

Copyright © 2016, Juniper Networks, Inc.

oif-map (IGMP Interface)

Syntax oif-map *map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast with Subscriber VLANs on page 4358

oif-map (MLD Interface)

Syntax oif-map *map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*], [edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Associate an outgoing interface (OIF) map to an MLD logical interface. The OIF map is a routing policy statement that can contain multiple terms.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Multicast with Subscriber VLANs on page 4358

Copyright © 2016, Juniper Networks, Inc.

override (PIM Static RP)

Syntax override;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp local],
[edit logical-systems *logical-system-name* protocols pim rp local family inet],
[edit logical-systems *logical-system-name* protocols pim rp local family inet6],
[edit logical-systems *logical-system-name* protocols pim rp static address *address*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp local],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp local family inet],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp local family inet6],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp static address *address*],
[edit protocols pim rp local],
[edit protocols pim rp local family inet],
[edit protocols pim rp local family inet6],
[edit protocols pim rp static address *address*],
[edit routing-instances *instance-name* protocols pim rp local],
[edit routing-instances *instance-name* protocols pim rp local family inet],
[edit routing-instances *instance-name* protocols pim rp local family inet6],
[edit routing-instances *instance-name* protocols pim rp static address *address*]

Release Information Statement introduced in Junos OS Release 11.4.

Description When you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for a given group range, and allow dynamic RP mapping for all other groups.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static RP on page 4135
- Configuring PIM Auto-RP on page 4152

Copyright © 2016, Juniper Networks, Inc.

override-interval

Syntax override-interval *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim]
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.

Options This is a random timer with a value in milliseconds.

Range: 0 through maximum override value

Default: 2000 milliseconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling Join Suppression on page 4122
- propagation-delay on page 4541
- reset-tracking-bit on page 4555

Copyright © 2016, Juniper Networks, Inc.

passive (IGMP)

Syntax passive <allow-receive> <send-general-query> <send-group-query>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
allow-receive, **send-general-query**, and **send-group-query** options were added in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the passive statement.

Options

allow-receive—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast with Subscriber VLANs on page 4358
- Enabling IGMP on page 4041

Copyright © 2016, Juniper Networks, Inc.

passive (MLD)

Syntax passive;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.
allow-receive, **send-general-query**, and **send-group-query** options added in Junos OS Release 10.0.

Description Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.

NOTE: You can selectively activate up to two out of the three available options for the passive statement while keeping the other functions passive (inactive). Activating all three options is equivalent to not using the passive statement.

Options **allow-receive**—Enables IGMP to receive control traffic on the interface.

send-general-query—Enables IGMP to send general queries on the interface.

send-group-query—Enables IGMP to send group-specific and group-source-specific queries on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast with Subscriber VLANs on page 4358

Copyright © 2016, Juniper Networks, Inc.

peer

Syntax

```
peer address {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  authentication-key peer-key;
  default-peer;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Define an MSDP peering relationship. An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple **peer** statements.

By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the peer statement.

At least one peer must be configured for MSDP to function. You must configure ***address*** and **local-address**.

Options

address—Name of the MSDP peer.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Example: Configuring MSDP in a Routing Instance on page 4254

pgm

Syntax

```
pgm {
  traceoptions {
    flag flag <flag-modifier>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure PGM globally and set tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

The remaining statement is explained separately.

Default

The default PGM trace options are inherited from the routing protocol **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- PGM Configuration Guidelines on page 4251

Copyright © 2016, Juniper Networks, Inc.

pim

Syntax

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        no-bidirectional-mode;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        family (inet | inet6) {
            disable;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm algorithm-name;
                    key-chain key-chain-name;
                loose-check;
                detection-time {
                    threshold milliseconds;
                }
            minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            version (0 | 1 | automatic);
            }
            accept-remote-source;
            disable;
            bidirectional {
                df-election {
                    backoff-period milliseconds;
                    offer-period milliseconds;
                    robustness-count number;
                }
            }
            family (inet | inet6) {
                disable;
            }
            hello-interval seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
    mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse |
      sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
}
join-load-balance;
join-prune-timeout;
mdt {
    data-mdt-reuse;
    group-range multicast-prefix;
    threshold {
        group group-address {
            source source-address {
                rate threshold-rate;
            }
        }
        tunnel-limit limit;
    }
}
mvpn {
    autodiscovery {
        inet-mdt;
    }
}
nonstop-routing;
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group group-name;
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
            }
        }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-import [ policy-names ];
    bootstrap-export [ policy-names ];
```

Copyright © 2016, Juniper Networks, Inc.

```
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  maximum-rps limit;
}
group-rp-mapping {
  family (inet | inet6) {
    log-interval seconds;
    maximum limit;
    threshold value;
    }
  }
  log-interval seconds;
  maximum limit;
  threshold value;
  }
}
local {
  family (inet | inet6) {
    address address;
    anycast-pim {
      rp-set {
        address address <forward-msdp-sa>;
      }
      disable;
      local-address address;
    }
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
  }
}
register-limit {
  family (inet | inet6) {
    log-interval seconds;
    maximum limit;
    threshold value;
    }
  }
  log-interval seconds;
  maximum limit;
  threshold value;
  }
}
rp-register-policy [ policy-names ];
spt-threshold {
  infinity [ policy-names ];
}
static {
  address address {
```


Copyright © 2016, Juniper Networks, Inc.

```
            override;
            version version;
            group-ranges {
              destination-ip-prefix</prefix-length>;
            }
          }
        }
      }
      rpf-selection {
        group group-address{
          sourcesource-address{
            next-hop next-hop-address;
          }
          wildcard-source {
            next-hop next-hop-address;
          }
        }
        prefix-list prefix-list-addresses {
          source source-address {
            next-hop next-hop-address;
          }
          wildcard-source {
            next-hop next-hop-address;
          }
        }
      sglimit {
        family (inet | inet6) {
          log-interval seconds;
          maximum limit;
          threshold value;
          }
        }
        log-interval seconds;
        maximum limit;
        threshold value;
        }
      }
      traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
      }
      tunnel-devices [ mt-fpc/pic/port ];
    }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.
family statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Copyright © 2016, Juniper Networks, Inc.

Description Enable PIM on the routing device.

The remaining statements are explained separately.

Default PIM is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Dense Mode Properties on page 4109
- Configuring PIM Sparse-Dense Mode Properties on page 4111

pim-to-igmp-proxy

Syntax
```
pim-to-igmp-proxy {
  upstream-interface [ interface-names ];
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-igmp-proxy** statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM-to-IGMP Message Translation on page 4241

Copyright © 2016, Juniper Networks, Inc.

pim-to-mld-proxy

Syntax

```
pim-to-mld-proxy {
  upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-mld-proxy** statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-to-MLD Message Translation on page 4242

policy (Flow Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a flow map policy.

Options ***policy-names***—Name of one or more policies for flow mapping.

Required Privilege Level routing—To view this statement in the configuration.

policy (SSM Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Apply one or more policies to an SSM map.

Options ***policy-names***—Name of one or more policies for SSM mapping.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation

- Example: Configuring SSM Mapping on page 4192

 Copyright © 2016, Juniper Networks, Inc.

prefix-list (PIM RPF Selection)

Syntax

```
prefix-list prefix-list-addresses {
  source source-address {
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level

[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address* wildcard-source],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* source *source-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses* wildcard-source]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.

Options

prefix-list-addresses—List of prefixes (addresses) for multiple PIM groups.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

priority (Bootstrap)

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bootstrap (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)],
[edit protocols pim rp bootstrap (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols pim rp bootstrap (inet | inet6)]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device's likelihood to be elected as the bootstrap router.

Options
number—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.

Range: 0 through a 32-bit number

Default: 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Bootstrap Properties for IPv4 on page 4147
- Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4149
- bootstrap-priority on page 4424

Copyright © 2016, Juniper Networks, Inc.

priority (PIM Interfaces)

Syntax priority *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the routing device's likelihood to be elected as the designated router.

Options *number*—Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.

Range: 0 through 4294967295

Default: 1 (Each routing device has an equal probability of becoming the DR.)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Interface Priority for PIM Designated Router Selection on page 4103

Copyright © 2016, Juniper Networks, Inc.

priority (PIM RPs)

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim rp bidirectional address *address*],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit protocols pim rp bidirectional address *address*],
[edit protocols pim rp local family (inet | inet6)],
[edit routing-instances *instance-name* protocols pim rp bidirectional address *address*],
[edit routing-instances *routing-instance-name* protocols pim rp local family (inet | inet6)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional RP addresses introduced in Junos OS Release 12.1.

Description For PIM-SM, configure this routing device's priority for becoming an RP.

For bidirectional PIM, configure this RP address' priority for becoming an RP.

The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.

Options *number*—Priority for becoming an RP. A lower value corresponds to a higher priority.

Range: 0 through 255

Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs on page 4136
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

propagation-delay

Syntax propagation-delay *milliseconds*;

Hierarchy Level
[edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Set a delay for implementing a PIM prune message on the upstream router on a multicast network for which join suppression has been enabled. The router waits for the prune pending period to detect whether a join message is currently being suppressed by another router.

Options ***milliseconds***—Interval for the prune pending timer, which is the sum of the **propagation-delay** value and the **override-interval** value.

Range: 250 through 2000 milliseconds

Default: 500 milliseconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling Join Suppression on page 4122
- override-interval on page 4525
- reset-tracking-bit on page 4555

Copyright © 2016, Juniper Networks, Inc.

proxy

Syntax

```
proxy {
    source-address ip-address;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure proxy mode and options, including source address. All the queries generated by IGMP snooping are sent using 0.0.0.0 as the source address in order to avoid participating in IGMP querier election. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as the source address unless there is a configured source address to use.

Default

By default, IGMP snooping does not employ proxy mode.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

query-interval (Bridge Domains)

Syntax query-interval *seconds*;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 8.5.

Description Configure the interval for host-query message timeouts.

Options *seconds*—Time interval.

Range: 1 through 1024

Default: 125 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335
- query-last-member-interval (Bridge Domains) on page 4546
- query-response-interval (Bridge Domains) on page 4548

Copyright © 2016, Juniper Networks, Inc.

query-interval (Protocols IGMP)

Syntax query-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how often the querier router sends general host-query messages.

Options *seconds*—Time interval.

Range: 1 through 1024

Default: 125 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Host-Query Message Interval on page 4042
- query-last-member-interval (Protocols IGMP) on page 4547
- query-response-interval (Protocols IGMP) on page 4549

Copyright © 2016, Juniper Networks, Inc.

query-interval (Protocols IGMP AMT)

Syntax query-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify how often the querier router sends IGMP general host-query messages through an Automatic Multicast Tunneling (AMT) interface.

Options ***seconds***—Number of seconds between sending of general host query messages.

Range: 1 through 1024

Default: 125 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

query-interval (Protocols MLD)

Syntax query-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify how often the querier router sends general host-query messages.

Options ***seconds***—Time interval.

Range: 1 through 1024

Default: 125 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the MLD Host-Query Message Interval on page 4068
- query-last-member-interval (Protocols MLD) on page 4547
- query-response-interval (Protocols MLD) on page 4551

Copyright © 2016, Juniper Networks, Inc.

query-last-member-interval (Bridge Domains)

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the interval for group-specific query timeouts.

Options ***seconds***—Time interval, in fractions of a second or seconds.

Range: 0.1 through 0.9, then in 1-second intervals 1 through 1024

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335
- query-interval on page 4543
- query-response-interval on page 4548

Copyright © 2016, Juniper Networks, Inc.

query-last-member-interval (Protocols IGMP)

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how often the querier router sends group-specific query messages.

Options *seconds*—Time interval, in fractions of a second or seconds.

Range: 0.1 through 0.9, then in 1-second intervals 1 through 999999

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Last-Member Query Interval on page 4047
- query-interval (Protocols IGMP) on page 4544
- query-response-interval (Protocols IGMP) on page 4549

query-last-member-interval (Protocols MLD)

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify how often the querier router sends group-specific query messages.

Options *seconds*—Time interval, in fractions of a second or seconds.

Range: 0.1 through 0.9, then in 1-second intervals from 1 through 1024

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the MLD Last-Member Query Interval on page 4070
- query-interval (Protocols MLD) on page 4545
- query-response-interval (Protocols MLD) on page 4551

Copyright © 2016, Juniper Networks, Inc.

query-response-interval (Bridge Domains)

Syntax query-response-interval *seconds*;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how long to wait to receive a response to a specific query message from a host.

Options ***seconds***—Time interval. This interval should be less than the host-query interval.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335
- query-interval (Bridge Domains) on page 4543
- query-last-member-interval (Bridge Domains) on page 4546

Copyright © 2016, Juniper Networks, Inc.

query-response-interval (Protocols IGMP)

Syntax query-response-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify how long the querier router waits to receive a response to a host-query message from a host.

Options ***seconds***—The query response interval must be less than the query interval.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Query Response Interval on page 4043
- query-interval (Protocols IGMP) on page 4544
- query-last-member-interval (Protocols IGMP) on page 4547

Copyright © 2016, Juniper Networks, Inc.

query-response-interval (Protocols IGMP AMT)

Syntax query-response-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify how long the IGMP querier router waits to receive a response to a host query message from a host through an Automatic Multicast Tunneling (AMT) interface. The query response interval must be less than the query interval.

Options ***seconds***—Time to wait to receive a response to a host query message.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

Copyright © 2016, Juniper Networks, Inc.

query-response-interval (Protocols MLD)

Syntax query-response-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify how long the querier router waits to receive a response to a host-query message from a host.

Options ***seconds***—Time interval. This interval must be less than the interval between general host-query messages.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the MLD Query Response Interval on page 4069
- query-interval (Protocols MLD) on page 4545
- query-last-member-interval (Protocols MLD) on page 4547

Copyright © 2016, Juniper Networks, Inc.

redundant-sources

Syntax redundant-sources [*addresses*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a list of redundant sources for multicast flows defined by a flow map.

Options ***addresses***—List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring a Multicast Flow Map on page 4377

Copyright © 2016, Juniper Networks, Inc.

relay (AMT Protocol)

Syntax

```
relay {
  accounting;
  family {
    inet {
      anycast-prefix ip-prefix/<prefix-length>;
      local-address ip-address;
    }
  }
  secret-key-timeout minutes;
  tunnel-limit number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols amt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt],
[edit protocols amt],
[edit routing-instances *routing-instance-name* protocols amt]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure the protocol address family, secret key timeout, and tunnel limit for Automatic Multicast Tunneling (AMT) relay functions.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

relay (IGMP)

Syntax

```
relay {
  defaults {
    (accounting (Protocols IGMP AMT Interface) | no-accounting (Protocols IGMP AMT
      Interface));
    group-policy [ policy-names ];
    query-interval seconds;
    query-response-interval seconds;
    robust-count number;
    ssm-map ssm-map-name;
    version version;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* statement-name protocols igmp amt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* statement-name protocols igmp amt],
[edit protocols igmp amt],
[edit routing-instances *routing-instance-name* statement-name protocols igmp amt]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure default Automatic Multicast Tunneling (AMT) interface attributes.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

Copyright © 2016, Juniper Networks, Inc.

reset-tracking-bit

Syntax reset-tracking-bit;

Hierarchy Level
[edit protocols pim],
[edit protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds (1.1 × periodic through 1.4 × periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Enabling Join Suppression on page 4122
- override-interval on page 4525
- propagation-delay on page 4541

Copyright © 2016, Juniper Networks, Inc.

restart-duration

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim graceful-restart],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim graceful-restart],
[edit protocols pim graceful-restart],
[edit routing-instances *routing-instance-name* protocols pim graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the duration of the graceful restart interval.

Options *seconds*—Time that the routing device waits (in seconds) to complete PIM sparse mode graceful restart.

Range: 30 through 300

Default: 60

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Sparse Mode Graceful Restart on page 4238

Copyright © 2016, Juniper Networks, Inc.

reverse-oif-mapping

Syntax

```
reverse-oif-mapping {
    no-qos-adjust;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
The **no-qos-adjust** statement added in Junos OS Release 9.5.
The **no-qos-adjust** statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast with Subscriber VLANs on page 4358

rib-group (Protocol DVMRP)

Syntax

rib-group *group-name*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols dvmrp],
[edit protocols dvmrp]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Associate a routing table group with DVMRP.

Options

group-name—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring DVMRP on page 4288

Copyright © 2016, Juniper Networks, Inc.

rib-group (Protocols MSDP)

Syntax rib-group *group-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Associate a routing table group with MSDP.

Options *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP in a Routing Instance on page 4254

Copyright © 2016, Juniper Networks, Inc.

rib-group (Protocols PIM)

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Associate a routing table group with PIM.

Options

table-name—Name of the routing table. The name must be one that you defined with the **rib-groups** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Dedicated PIM RPF Routing Table on page 4317

Copyright © 2016, Juniper Networks, Inc.

robust-count (Bridge Domains)

Syntax robust-count *number*;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Provide fine-tuning to allow for expected packet loss on a subnet. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.

Options *number*—Robust interval.

Range: 2 through 10

Default: 2

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

robust-count (Protocols IGMP)

Syntax robust-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.

Options *number*—Robustness variable.

Range: 2 through 10

Default: 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Modifying the IGMP Robustness Variable on page 4047

robust-count (Protocols IGMP AMT)

Syntax robust-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the expected IGMP packet loss on an Automatic Multicast Tunneling (AMT) tunnel. If a tunnel is expected to have packet loss, increase the robust count.

Options *number*—Number of packets that can be lost before the AMT protocol deletes the multicast state.

Range: 2 through 10

Default: 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

Copyright © 2016, Juniper Networks, Inc.

robust-count (Protocols MLD)

Syntax robust-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced before Junos OS Release 7.4.

Description Tune for the expected packet loss on a subnet.

Options *number*—Time interval. This interval must be less than the interval between general host-query messages.

Range: 2 through 10

Default: 2 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Modifying the MLD Robustness Variable on page 4073

Copyright © 2016, Juniper Networks, Inc.

robustness-count

Syntax robustness-count *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name* bidirectional df-election],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election],
[edit protocols pim interface *interface-name* bidirectional df-election],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bidirectional df-election]

Release Information Statement introduced in Junos OS Release 12.1.

Description Configure the designated forwarder (DF) election robustness count for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The **robustness-count** statement sets the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routers on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.

If a router receives from a neighbor a better offer than its own, the router stops participating in the election for a period of **robustness-count** * offer-period. Eventually, all routers except the best candidate stop sending Offer messages.

Options *number*—Number of transmission attempts for DF election messages.

Range: 1 through 10

Default: 3

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Bidirectional PIM on page 4199
- Example: Configuring Bidirectional PIM on page 4205

Copyright © 2016, Juniper Networks, Inc.

rp

Syntax

```
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
            }
        }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    local {
        family (inet | inet6) {
            disable;
            address address;
            anycast-pim {
                local-address address;
                    address address <forward-msdp-sa>;
                    rp-set {
                }
            }
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            override;
            priority number;
        }
    }
    rp-register-policy [ policy-names ];
    static {
        address address {
```


Copyright © 2016, Juniper Networks, Inc.

```
        override;
        version version;
        group-ranges {
          destination-ip-prefix</prefix-length>;
        }
      }
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.

The remaining statements are explained separately.

Default

If you do not include the **rp** statement, the routing device can never become the RP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding PIM Sparse Mode on page 4113

Copyright © 2016, Juniper Networks, Inc.

rp-register-policy

Syntax rp-register-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply one or more policies to control incoming PIM register messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Register Message Filters on a PIM RP and DR on page 4164
- dr-register-policy on page 4436

Copyright © 2016, Juniper Networks, Inc.

rp-set

Syntax

```
rp-set {
    address address <forward-msdp-sa>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim local family (inet | inet6) anycast-pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim],
[edit protocols pim local family (inet | inet6) anycast-pim],
[edit routing-instances *routing-instance-name* protocols pim local family (inet | inet6) anycast-pim]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast With or Without MSDP on page 4142

Copyright © 2016, Juniper Networks, Inc.

rpf-check-policy

Syntax rpf-check-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options ***policy-names***—Name of one or more multicast RPF check policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RPF Policies on page 4321

Copyright © 2016, Juniper Networks, Inc.

rpf-selection

Syntax

```
rpf-selection {
  group group-address {
    sourcesource-address {
      next-hop next-hop-address;
    }
    wildcard-source {
      next-hop next-hop-address;
    }
  }
  prefix-list prefix-list-addresses {
    source source-address {
      next-hop next-hop-address;
    }
    wildcard-source {
      next-hop next-hop-address;
    }
  }
```

Hierarchy Level [edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

The remaining statements are explained separately.

Default If you omit the **rpf-selection** statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

Options ***source-address***—Specific source address for the PIM group.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

sap

Syntax

```
sap {
  disable;
  listen address <port port>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable the router to listen to session directory announcements for multimedia and other multicast sessions.

SAP and SDP always listen on the default SAP address and port, 224.2.127.254:9875. To have SAP listen on additional addresses or pairs of address and port, include a **listen** statement for each address or pair.

Options

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Session Announcement Protocol on page 4271
- listen on page 4491

Copyright © 2016, Juniper Networks, Inc.

scope

Syntax

```
scope scope-name {
  interface [ interface-names ];
  prefix destination-prefix;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure multicast scoping.

Options

scope-name—Name of the multicast scope.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Multicast Snooping on page 4344

Copyright © 2016, Juniper Networks, Inc.

scope-policy

Syntax scope-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]

NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the **scope** statement to a specific routing instance at the **[edit routing-instances *routing-instance-name* routing-options multicast]** or **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]** hierarchy level.

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

Options ***policy-names***—Name of one or more multicast scope policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- scope on page 4571

Copyright © 2016, Juniper Networks, Inc.

secret-key-timeout

Syntax secret-key-timeout *minutes*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols amt relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay],
[edit protocols amt relay],
[edit routing-instances *routing-instance-name* protocols amt relay]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the period in minutes after which the local opaque secret key used in the Automatic Multicast Tunneling (AMT) Message Authentication Code (MAC) times out and is regenerated.

Default 60 minutes

Options ***minutes***—Number of minutes to wait before generating a new MAC opaque secret key.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the AMT Protocol on page 4277

source (Bridge Domains)

Syntax source *ip-address*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name* static group],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name* static group]

Release Information Statement introduced in Junos OS Release 8.5.

Description Statically define multicast group source addresses on an interface.

Options ***ip-address***—IP address to use as the source for the group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

source (PIM RPF Selection)

Syntax

```
source source-address {
    next-hop next-hop-address;
}
```

Hierarchy Level

[edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses*]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the source address for the PIM group.

Options

source-address—Specific source address for the PIM group.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

source (Protocols IGMP)

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*],
[edit protocols igmp interface *interface-name* static group *multicast-group-address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address—IPv4 unicast address.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

source (Protocols MLD)

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address*],
[edit protocols mld interface *interface-name* static group *multicast-group-address*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

IP version 6 (IPv6) unicast source address for the multicast group being statically configured on an interface.

Options

ip-address — One or more IPv6 unicast addresses.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

source (Protocols MSDP)

Syntax

```
source ip-address</prefix-length> {
  active-source-limit {
    maximum number;
    threshold number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Limit the number of active source messages the routing device accepts from sources in this address range.

Default

If you do not include this statement, the routing device accepts any number of MSDP active source messages.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262

Copyright © 2016, Juniper Networks, Inc.

source (Routing Instances)

Syntax source [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify IPv4 or IPv6 source addresses for an SSM map.

Options ***addresses***—IPv4 or IPv6 source addresses.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 4192

source-address

Syntax source-address *ip-address*;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping proxy],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* proxy],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping proxy],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* proxy]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured.

Options ***ip-address***—IP address to use as the source for proxy-mode IGMP snooping reports.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

source-count (Protocols IGMP)

Syntax source-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of multicast source addresses that should be accepted for each static group created.

Options *number*—Number of source addresses.

Default: 1

Range: 1 through 1024

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

source-count (Protocols MLD)

Syntax source-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address* source],
[edit protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the number of multicast source addresses that should be accepted for each static group created.

Options *number*—Number of source addresses.

Default: 1

Range: 1 through 1024

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

source-increment (Protocols IGMP)

Syntax source-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source],
[edit protocols igmp interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.

Options *increment*—Number of times the source address should be incremented.

Default: 0.0.0.1

Range: 0.0.0.1 through 255.255.255.255

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

source-increment (Protocols MLD)

Syntax source-increment *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name* static group *multicast-group-address* source],
[edit protocols mld interface *interface-name* static group *multicast-group-address* source]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.

Options *increment*—Number of times the source address should be incremented.

Default: ::1

Range: ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

spt-threshold

Syntax

```
spt-threshold {
  infinity [ policy-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the PIM SPT Threshold Policy on page 4177

Copyright © 2016, Juniper Networks, Inc.

ssm-groups

Syntax ssm-groups [*ip-addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure source-specific multicast (SSM) groups.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.

Options ***ip-addresses***—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4187

Copyright © 2016, Juniper Networks, Inc.

ssm-map (Protocols IGMP AMT)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Apply a source-specific multicast (SSM) map to all Automatic Multicast Tunneling (AMT) interfaces.

Options ***ssm-map-name***—Name of the SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

ssm-map (Protocols MLD)

Syntax ssm-map *ssm-map-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Apply an SSM map to an MLD interface.

Options ***ssm-map-name***—Name of SSM map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Mapping on page 4192

Copyright © 2016, Juniper Networks, Inc.

ssm-map (Multicast Routing Options)

Syntax

```
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure SSM mapping.

Options

ssm-map-name—Name of the SSM map.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring SSM Mapping on page 4192

ssm-map (Protocols IGMP)

Syntax

```
ssm-map ssm-map-name;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Apply an SSM map to an IGMP interface.

Options

ssm-map-name—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring SSM Mapping on page 4192

Copyright © 2016, Juniper Networks, Inc.

ssm-map-policy (IGMP)

Syntax ssm-map-policy *ssm-map-policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply an SSM map policy to an IGMP interface.

Options ***ssm-map-policy-name***—Name of SSM map policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Maps for Different Groups to Different Sources on page 4195

ssm-map-policy (MLD)

Syntax ssm-map-policy *ssm-map-policy-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*],
[edit protocols mld interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Apply an SSM map policy to an MLD interface.

Options ***ssm-map-policy-name***—Name of SSM map policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring SSM Maps for Different Groups to Different Sources on page 4195

Copyright © 2016, Juniper Networks, Inc.

static (Bridge Domains)

Syntax

```
static {
  group multicast-group-address {
    source ip-address;
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Define static multicast groups on an interface.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on page 4335

Copyright © 2016, Juniper Networks, Inc.

static (Protocols MLD)

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*], [edit protocols mld interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Test multicast forwarding on an interface.

The **static** statement simulates MLD joins on a routing device statically on an interface without any MLD hosts. It is supported for both MLDv1 and MLDv2 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.

NOTE: To prevent joining too many groups accidentally, the static statement is not supported with the interface all statement.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Enabling MLD Static Group Membership on page 4075

Copyright © 2016, Juniper Networks, Inc.

static (Protocols IGMP)

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*], [edit protocols igmp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.

NOTE: To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership on page 4050

Copyright © 2016, Juniper Networks, Inc.

static (Protocols PIM)

Syntax

```
static {
  address address {
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    override;
    version version;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more **address** statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4137

Copyright © 2016, Juniper Networks, Inc.

subscriber-leave-timer

Syntax subscriber-leave-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options ***seconds***—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.

Range: 0 through 30

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast with Subscriber VLANs on page 4358

Copyright © 2016, Juniper Networks, Inc.

threshold (Bridge Domains)

Syntax threshold suppress *value* <reuse *value*>;

Hierarchy Level [edit bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* multicast-snooping-options forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache],
[edit routing-instances *routing-instance-name* multicast-snooping-options forwarding-cache],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* multicast-snooping-options forwarding-cache]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the suppression and reuse thresholds for multicast snooping forwarding cache limits.

Options

suppress *value*—Value to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number must be greater than the reuse value.

Range: 1 through 200,000

reuse *value*—(Optional) Value to begin creating new multicast forwarding cache entries. If configured, this number must be less than the suppress value.

Range: 1 through 200,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Multicast Snooping on page 4345

Copyright © 2016, Juniper Networks, Inc.

threshold (Protocols MSDP)

Syntax threshold *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols msdp active-source-limit],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp active-source-limit],
[edit protocols msdp active-source-limit],
[edit routing-instances *routing-instance-name* protocols msdp active-source-limit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.

Options *number*—RED threshold for active source messages.

Range: 1 through 1,000,000

Default: 24,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4262
- maximum (MSDP Active Source Messages) on page 4499

Copyright © 2016, Juniper Networks, Inc.

threshold (PIM BFD Detection Time)

Syntax threshold *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection detection-time],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection detection-time]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

NOTE: The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the minimum-interval or the minimum-receive-interval statement.

Options *milliseconds*—Value for the detection time adaptation threshold.

Range: 1 through 255,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418
- detection-time on page 4429
- minimum-interval on page 4502
- minimum-receive-interval on page 4504

Copyright © 2016, Juniper Networks, Inc.

threshold (PIM BFD Transmit Interval)

Syntax threshold *milliseconds*;

Hierarchy Level [edit protocols pim interface *interface-name* bfd-liveness-detection transmit-interval],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection transmit-interval]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options ***milliseconds***—Value for the transmit interval adaptation threshold.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

NOTE: The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418

Copyright © 2016, Juniper Networks, Inc.

timeout (Flow Maps)

Syntax timeout (never non-discard-entry-only | *minutes*);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the timeout value for multicast forwarding cache entries associated with the flow map.

Options *minutes*—Length of time that the forwarding cache entry remains active.

Range: 1 through 720

never non-discard-entry-only—Specify that the forwarding cache entry always remain active. If you omit the **non-discard-entry-only** option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the **non-discard-entry-only** option, entries with forwarding states are kept forever, and entries with pruned states time out.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

timeout (Multicast)

Syntax timeout *minutes*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the timeout value for multicast forwarding cache entries.

Options ***minutes***—Length of time that the forwarding cache limit remains active.

Range: 1 through 720

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the Multicast Forwarding Cache on page 4374

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Multicast Snooping Options)

Syntax

```
traceoptions {
  file filename<files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level [edit multicast-snooping-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Set multicast snooping tracing options.

Default Tracing operations are disabled.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place multicast snooping tracing output in the file **/var/log/multicast-snooping-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

The following are the tracing options:

- **all**—All tracing operations
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.

Copyright © 2016, Juniper Networks, Inc.

- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Multicast Snooping on page 4344
- Example: Configuring Multicast Snooping on page 4345
- Enabling Bulk Updates for Multicast Snooping on page 4350
- Example: Configuring Multicast Snooping on page 4345

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols AMT)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols amt],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt],
[edit protocols amt],
[edit routing-instances *routing-instance-name* protocols amt]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure Automatic Multicast Tunneling (AMT) tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

AMT Tracing Flags

- **errors**—All error conditions
- **packets**—All AMT packets
- **tunnels**—All AMT tunnel-related information

Global Tracing Flags

- **all**—All tracing operations

Copyright © 2016, Juniper Networks, Inc.

- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring the AMT Protocol on page 4277

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols DVMRP)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols dvmrp],
[edit protocols dvmrp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure DVMRP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default The default DVMRP trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **dvmrp-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

DVMRP Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **graft**—Graft messages
- **neighbor**—Neighbor probe messages
- **normal**—All normal operations

Copyright © 2016, Juniper Networks, Inc.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **packets**—All DVMRP packets
- **poison**—Poison-route-reverse packets
- **probe**—Probe packets
- **prune**—Prune messages
- **report**—DVMRP route report packets
- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Copyright © 2016, Juniper Networks, Inc.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing DVMRP Protocol Traffic on page 4295

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols IGMP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure IGMP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

To trace the paths of multicast packets, use the **mtrace** command.

Default

The default IGMP trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

IGMP Tracing Flags

- **leave**—Leave group messages (for IGMP version 2 only).
- **mtrace**—Mtrace packets. Use the **mtrace** command to troubleshoot the software.
- **packets**—All IGMP packets.

Copyright © 2016, Juniper Networks, Inc.

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Copyright © 2016, Juniper Networks, Inc.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing IGMP Protocol Traffic on page 4059

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols IGMP Snooping)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> ;
    flag flag (detail | disable | receive | send);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* bridge-domains *domain-name* protocols igmp-snooping],
[edit logical-systems *logical-system-name* routing-instances instance-name bridge-domains *domain-name* protocols igmp-snooping],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols igmp-snooping],
[edit bridge-domains *domain-name* protocols igmp-snooping],
[edit routing-instances *instance-name* bridge-domains *domain-name* protocols igmp-snooping],
[edit routing-instances *instance-name* protocols igmp-snooping]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Define tracing operations for IGMP snooping.

Default

The **traceoptions** feature is disabled by default.

Options

file *filename*—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x***k to specify KB, ***x***m to specify MB, or ***x***g to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag* —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **client-notification**—Trace notifications.
- **general**—Trace general IGMP snooping protocol events.
- **group**—Trace group operations.
- **host-notification**—Trace host notifications.
- **leave**—Trace leave group messages (IGMPv2 only).
- **normal**—Trace normal IGMP snooping protocol events.
- **packets**—Trace all IGMP packets.

Copyright © 2016, Juniper Networks, Inc.

- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.
- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.

no-world-readable—(Optional) Restrict file access to the user who created the file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes

Range: 10 KB through 1 gigabytes

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring IGMP Snooping Trace Operations on page 4341
- Configuring IGMP Snooping on page 4333

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols MLD)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld],
[edit protocols mld]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure MLD tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

To trace the paths of multicast packets, use the **mtrace** command.

Default The default MLD trace options are those inherited from the **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **mld-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MLD Tracing Flags

- **leave**—Leave group messages.
- **mtrace**—Mtrace packets. Use the **mtrace** command to troubleshoot the software.
- **packets**—All MLD packets.
- **query**—MLD membership query messages, including general and group-specific queries.

Copyright © 2016, Juniper Networks, Inc.

- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—Traces errors and significant events during normal packet processing

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Copyright © 2016, Juniper Networks, Inc.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- *Tracing MLD Protocol Traffic*

traceoptions (Protocols MSDP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols msdp],
[edit logical-systems *logical-system-name* protocols msdp group *group-name*],
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* protocols msdp peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols msdp peer *address*],
[edit protocols msdp],
[edit protocols msdp group *group-name*],
[edit protocols msdp group *group-name* peer *address*],
[edit protocols msdp peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name*],
[edit routing-instances *routing-instance-name* protocols msdp group *group-name* peer *address*],
[edit routing-instances *routing-instance-name* protocols msdp peer *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure MSDP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default MSDP trace options are those inherited from the routing protocol's **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **msdp-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Copyright © 2016, Juniper Networks, Inc.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

Copyright © 2016, Juniper Networks, Inc.

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*****.0**. When ***trace-file*** again reaches this size, ***trace-file*****.0** is renamed ***trace-file*****.1** and ***trace-file*** is renamed ***trace-file*****.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing MSDP Protocol Traffic on page 4268

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols PGM)

Syntax

```
traceoptions {
  flag flag <flag-modifier>;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols pgm],
[edit protocols pgm]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure PGM tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default The default PGM trace options are those inherited from the routing protocol **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

PGM Tracing Flags

- **all**—Trace all PGM packets.
- **init**—Trace all PGM initialization events.
- **packets**—Trace all PGM packet processing.
- **parser**—Trace all PGM parser processing.
- **route-socket**—Trace all PGM route-socket events.
- **show**—Trace all PGM **show** command servicing.
- **state**—Trace all PGM state transitions.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions

Copyright © 2016, Juniper Networks, Inc.

- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of the following modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- PGM Configuration Guidelines on page 4251

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Protocols PIM)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure PIM tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default PIM trace options are those inherited from the routing protocol's **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **pim-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

PIM Tracing Flags

- **assert**—Assert messages
- **bidirectional-df-election**—Bidirectional PIM designated-forwarder (DF) election events

Copyright © 2016, Juniper Networks, Inc.

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

Copyright © 2016, Juniper Networks, Inc.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 0 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Trace Options on page 4098
- Tracing DVMRP Protocol Traffic on page 4295
- Tracing MSDP Protocol Traffic on page 4268

Copyright © 2016, Juniper Networks, Inc.

transmit-interval (PIM BFD Liveness Detection)

Syntax

```
transmit-interval {
  minimum-interval milliseconds;
  threshold milliseconds;
}
```

Hierarchy Level

[edit protocols pim interface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for BFD authentication introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Specify the transmit interval for the **bfd-liveness-detection** statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221
- bfd-liveness-detection on page 4418
- threshold on page 4593
- minimum-interval on page 4503
- minimum-receive-interval on page 4504

Copyright © 2016, Juniper Networks, Inc.

tunnel-devices (Routing Instances)

Syntax tunnel-devices [mt-*fpc/pic/port*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols pim], [edit routing-instances *instance-name* protocols pim]

Release Information Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 10.2 for EX Series switches.

Description List one or more tunnel-capable PICs to be used for creating multicast tunnel (**mt**) interfaces. Creating a PIC list enables you to control the load-balancing implementation.

Tunnel-capable PICs include:

- Adaptive Services PIC
- Multiservices PIC or Multiservices DPC
- Tunnel Services PIC
- On MX Series routers, a PIC created with the **tunnel-services** statement at the **[edit chassis fpc *slot-number* pic *number*]** hierarchy level.

The physical position of the PIC in the routing device determines the multicast tunnel interface name. For example, if you have an Adaptive Services PIC installed in FPC slot 0 and PIC slot 0, the corresponding multicast tunnel interface name is **mt-0/0/0**. The same is true for Tunnel Services PICs, Multiservices PICs, and Multiservices DPCs.

Default Multicast tunnel interfaces are created on all available tunnel-capable PICs, based on a round-robin algorithm.

Options **mt-*fpc/pic/port***—Interface that is automatically generated when a tunnel-capable PIC is installed in the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

tunnel-limit (Protocols AMT)

Syntax tunnel-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols amt relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols amt relay],
[edit protocols amt relay],
[edit routing-instances *routing-instance-name* protocols amt relay]

Release Information Statement introduced in Junos OS Release 10.2.

Description Limit the number of Automatic Multicast Tunneling (AMT) data tunnels created. The system might reach a dynamic upper limit of tunnels of all types before the static AMT limit is reached.

Options *number*—Maximum number of data AMTs that can be created on the system.

Range: 0 through 4294967295

Default: 1 tunnel

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the AMT Protocol on page 4277

 Copyright © 2016, Juniper Networks, Inc.

upstream-interface

Syntax upstream-interface [*interface-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-mld-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit routing-options multicast pim-to-igmp-proxy],
[edit routing-options multicast pim-to-mld-proxy]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the **pim-to-igmp-proxy** statement), or into corresponding MLD report or leave messages (if you include the **pim-to-mld-proxy** statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.

Options
interface-names—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([ ]).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM-to-IGMP Message Translation on page 4241
- Configuring PIM-to-MLD Message Translation on page 4242

Copyright © 2016, Juniper Networks, Inc.

version (BFD)

Syntax version (0 | 1 | automatic);

Hierarchy Level [edit protocols piminterface *interface-name* bfd-liveness-detection],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.

Options Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for PIM on page 4221

Copyright © 2016, Juniper Networks, Inc.

version (PIM)

Syntax version *version*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* protocols pim rp static address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp static address *address*],
[edit protocols pim interface *interface-name*],
[edit protocols pim rp static address *address*],
[edit routing-instances *routing-instance-name* protocols pim interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols pim rp static address *address*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Specify the version of PIM.

Options *version*—PIM version number.

Range: 1 or 2

Default: PIMv1 for rendezvous point (RP) mode (at the **[edit protocols pim rp static address *address*]** hierarchy level). PIMv2 for interface mode (at the **[edit protocols pim interface *interface-name*]** hierarchy level).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling PIM Sparse Mode on page 4118
- Configuring PIM Dense Mode Properties on page 4109
- Configuring PIM Sparse-Dense Mode Properties on page 4111

Copyright © 2016, Juniper Networks, Inc.

version (Protocols IGMP)

Syntax version *version*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*],
[edit protocols igmp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Specify the version of IGMP.

Options *version*—IGMP version number.

Range: 1, 2, or 3

Default: IGMP version 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Changing the IGMP Version on page 4049

version (Protocols IGMP AMT)

Syntax version *version*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp amt relay defaults],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols igmp amt relay defaults],
[edit protocols igmp amt relay defaults],
[edit routing-instances *routing-instance-name* protocols igmp amt relay defaults]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the version of IGMP used through an Automatic Multicast Tunneling (AMT) interface.

Options *version*—IGMP version number.

Range: 1, 2, or 3

Default: IGMP version 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Default IGMP Parameters for AMT Interfaces on page 4279

Copyright © 2016, Juniper Networks, Inc.

version (Protocols MLD)

Syntax version *version*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mld interface *interface-name*], [edit protocols mld interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the MLD version explicitly. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).

Options ***version***—MLD version to run on the interface.

Range: 1 or 2

Default: 1 (MLDv1)

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Modifying the MLD Version on page 4068

Copyright © 2016, Juniper Networks, Inc.

vlan (Bridge Domains)

Syntax

```
vlan vlan-id {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    multicast-router-interface;
    static {
      group multicast-group-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols igmp-snooping],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure IGMP snooping parameters for a particular VLAN.

Default

By default, IGMP snooping options apply to all VLANs.

Options

vlan-id—Apply the parameters to this VLAN.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN-Specific IGMP Snooping Parameters on page 4334
- igmp-snooping on page 4470

Copyright © 2016, Juniper Networks, Inc.

vpn-group-address

Syntax vpn-group-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim],
[edit routing-instances *routing-instance-name* protocols pim]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the group address for the Layer 3 VPN in the service provider's network.

Options ***address***—Address for the Layer 3 VPN in the service provider's network.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Multicast Feature Guide for Security Devices*

wildcard-source (PIM RPF Selection)

Syntax
```
wildcard-source {
    next-hop next-hop-address;
}
```

Hierarchy Level [edit routing-instances *routing-instance-name* protocols pim rpf-selection group *group-address*],
[edit routing-instances *routing-instance-name* protocols pim rpf-selection prefix-list *prefix-list-addresses*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.

The remaining statements are explained separately.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM RPF Selection on page 4323

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 197

Operational Commands

- clear amt statistics
- clear amt tunnel
- clear igmp membership
- clear igmp snooping membership
- clear igmp snooping statistics
- clear igmp statistics
- clear mld membership
- clear mld statistics
- clear multicast snooping statistics
- clear pgm negative-acknowledgments
- clear pgm source-path-messages
- clear pgm statistics
- clear pim join
- clear pim join-distribution
- clear pim register
- clear pim statistics
- request pim multicast-tunnel rebalance
- show amt statistics
- show amt summary
- show amt tunnel
- show dvmrp interfaces
- show dvmrp neighbors
- show dvmrp prefix
- show dvmrp prunes
- show igmp group
- show igmp interface
- show igmp snooping interface
- show igmp snooping membership

Copyright © 2016, Juniper Networks, Inc.

- show igmp snooping statistics
- show igmp statistics
- show mld group
- show mld interface
- show mld statistics
- show msdp
- show msdp source
- show msdp source-active
- show msdp statistics
- show multicast backup-pe-groups
- show multicast flow-map
- show multicast interface
- show multicast pim-to-igmp-proxy
- show multicast pim-to-mld-proxy
- show multicast route
- show multicast rpf
- show multicast scope
- show multicast sessions
- show multicast snooping route
- show multicast snooping statistics
- show multicast usage
- show pgm negative-acknowledgments
- show pgm source-path-messages
- show pgm statistics
- show pim bidirectional df-election
- show pim bidirectional df-election interface
- show pim bootstrap
- show pim interfaces
- show pim join
- show pim neighbors
- show pim rps
- show pim source
- show pim statistics
- show policy
- show route
- show route table
- show sap listen

Copyright © 2016, Juniper Networks, Inc.

clear amt statistics

Syntax
clear amt statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in Junos OS Release 10.2.

Description
Clear Automatic Multicast Tunneling (AMT) statistics.

Options
none—Clear the multicast statistics for all AMT tunnel interfaces.

instance *instance-name*—(Optional) Clear AMT multicast statistics for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show amt statistics on page 4658

List of Sample Output
clear amt statistics on page 4633

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt statistics

```
user@host> clear amt statistics
```

clear amt tunnel

Syntax

clear amt tunnel
<gateway *gateway-ip-addr*> <port *port-number*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<statistics>
<tunnel-interface *interface-name*>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Clear the Automatic Multicast Tunneling (AMT) multicast state. Optionally, clear AMT protocol statistics.

Options

none—Clear multicast state for all AMT tunnel interfaces.

gateway *gateway-ip-addr* port *port-number*—(Optional) Clear the AMT multicast state for the specified gateway address. If no port is specified, clear the AMT multicast state for all AMT gateways with the given IP address.

instance *instance-name*—(Optional) Clear the AMT multicast state for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

statistics—(Optional) Clear multicast statistics for all AMT tunnels or for specified tunnels.

tunnel-interface *interface-name*—(Optional) Clear the AMT multicast state for the specified AMT tunnel interface.

Required Privilege Level

clear

Related Documentation

- show amt tunnel on page 4663

List of Sample Output

clear amt tunnel on page 4634
clear amt tunnel statistics gateway-address on page 4634

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt tunnel

```
user@host> clear amt tunnel
```

clear amt tunnel statistics gateway-address

```
user@host> clear amt tunnel statistics gateway-address 100.31.1.21 port 4000
```


Copyright © 2016, Juniper Networks, Inc.

clear igmp membership

List of Syntax
Syntax on page 4635
Syntax (EX Series Switch and the QFX Series) on page 4635

Syntax
clear igmp membership
<group *address-range*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
clear igmp membership
<group *address-range*>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Clear Internet Group Management Protocol (IGMP) group members.

Options
none—Clear all IGMP members on all interfaces and for all address ranges.

group *address-range*—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is **224.2/16**. If you omit the destination prefix length, the default is **/32**.

interface *interface-name*—(Optional) Clear all IGMP group members on an interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show igmp group on page 4674
- show igmp interface on page 4678

List of Sample Output
clear igmp membership on page 4635
clear igmp membership interface on page 4636
clear igmp membership group on page 4637

Output Fields
See show igmp group for an explanation of output fields.

Sample Output

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          186
so-0/0/0         224.2.127.254   10.1.128.1          186
so-0/0/0         239.255.255.255 10.1.128.1          187
so-0/0/0         224.1.127.255   10.1.128.1          188
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```


Copyright © 2016, Juniper Networks, Inc.

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.1.127.255   10.1.128.1          231
so-0/0/0         224.2.127.254   10.1.128.1          233
so-0/0/0         224.2.127.253   10.1.128.1          236
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

Copyright © 2016, Juniper Networks, Inc.

clear igmp snooping membership

Syntax

clear igmp snooping membership
<group | source *address*>
<instance *instance-name*>
<interface *interface-name*>
<learning-domain *learning-domain-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Clear IP IGMP snooping membership information.

Options

none—Clear IGMP snooping membership for all supported address families on all interfaces.

group | source *address*—(Optional) Clear IGMP snooping membership for the specified multicast group or source address.

instance *instance-name*—(Optional) Clear IGMP snooping membership for the specified instance.

interface *interface-name*—(Optional) Clear IGMP snooping membership on a specific interface.

learning-domain *learning-domain-name*—(Optional) Perform this operation on all learning domains or on a particular learning domain.

vlan-id *vlan-identifier*—(Optional) Perform this operation on a particular VLAN.

Required Privilege Level

clear

Related Documentation

- show igmp snooping membership on page 4685

List of Sample Output

clear igmp snooping membership on page 4638

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping membership

```
user@host> clear igmp snooping membership
```


Copyright © 2016, Juniper Networks, Inc.

clear igmp snooping statistics

Syntax
clear igmp snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<learning-domain (all | *learning-domain-name*)>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in Junos OS Release 8.5.

Description
Clear IP IGMP snooping statistics.

Options
none—Clear IGMP snooping statistics for all supported address families on all interfaces.

instance *instance-name*—(Optional) Clear IGMP snooping statistics for the specified instance.

interface *interface-name*—(Optional) Clear IGMP snooping statistics on a specific interface.

learning-domain (all | *learning-domain-name*)—(Optional) Perform this operation on all learning domains or on a particular learning domain.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show igmp snooping statistics on page 4689

List of Sample Output
clear igmp snooping statistics on page 4639

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping statistics

```
user@host> clear igmp snooping statistics
```

Copyright © 2016, Juniper Networks, Inc.

clear igmp statistics

List of Syntax Syntax on page 4640
Syntax (EX Series Switches) on page 4640

Syntax

clear igmp statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

clear igmp statistics
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Clear Internet Group Management Protocol (IGMP) statistics.

Options

none—Clear IGMP statistics on all interfaces.

interface *interface-name*—(Optional) Clear IGMP statistics for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

List of Sample Output clear igmp statistics on page 4640

Output Fields

See show igmp statistics for an explanation of output fields.

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
DVMRP                     19784      35476          0
PIM V1                    18310          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
```


Copyright © 2016, Juniper Networks, Inc.

```
IGMP v3 unsupported type                          0
IGMP v3 source required for SSM                   0
IGMP v3 mode not applicable for SSM               0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local               1227
```

```
user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                              0
IGMP v3 unsupported type                         0
IGMP v3 source required for SSM                  0
IGMP v3 mode not applicable for SSM              0
IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local                  0
```

clear mld membership

Syntax
clear mld membership
<group *group-name*> | <interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Clear Multicast Listener Discovery (MLD) group membership.

Options
none—Clear all MLD memberships.

group *group-name*—(Optional) Clear MLD membership for the specified group.

interface *interface-name*—(Optional) Clear MLD group membership for the specified interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show mld group on page 4695

List of Sample Output
clear mld membership on page 4642

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld membership

```
user@host> clear mld membership
```


Copyright © 2016, Juniper Networks, Inc.

clear mld statistics

Syntax
clear mld statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Clear Multicast Listener Discovery (MLD) statistics.

Options
none—(Same as **logical-system all**) Clear MLD statistics for all interfaces.

interface *interface-name*—(Optional) Clear MLD statistics for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show mld statistics on page 4702

List of Sample Output
clear mld statistics on page 4643

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld statistics

```
user@host> clear mld statistics
```

clear multicast snooping statistics

Syntax

clear multicast snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Clear IP multicast snooping statistics.

Options

none—Clear multicast snooping statistics for all supported address families on all interfaces.

instance *instance-name*—(Optional) Clear multicast snooping statistics for the specified instance.

interface *interface-name*—(Optional) Clear multicast snooping statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show multicast snooping statistics on page 4743

List of Sample Output

clear multicast snooping statistics on page 4644

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast snooping statistics

```
user@host> clear multicast snooping statistics
```


Copyright © 2016, Juniper Networks, Inc.

clear pgm negative-acknowledgments

Syntax clear pgm negative-acknowledgments

Release Information Command introduced before Junos OS Release 7.4.

Description Clear the Pragmatic General Multicast (PGM) negative acknowledgment (NAK) state received.

Options This command has no options.

Required Privilege Level clear

Related Documentation
- show pgm negative-acknowledgments on page 4749

List of Sample Output clear pgm negative-acknowledgments on page 4645

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm negative-acknowledgments

```
user@host> clear pgm negative-acknowledgments
```

clear pgm source-path-messages

Syntax clear pgm source-path-messages

Release Information Command introduced before Junos OS Release 7.4.

Description Clear Pragmatic General Multicast (PGM) source-path messages.

Options This command has no options.

Required Privilege Level clear

Related Documentation
- show pgm source-path-messages on page 4751

List of Sample Output clear pgm source-path-messages on page 4646

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm source-path-messages

```
user@host> clear pgm source-path-messages
```


Copyright © 2016, Juniper Networks, Inc.

clear pgm statistics

Syntax clear pgm statistics

Release Information Command introduced before Junos OS Release 7.4.

Description Clear Pragmatic General Multicast (PGM) statistics.

Options This command has no options.

Required Privilege Level clear

Related Documentation
- show pgm statistics on page 4752

List of Sample Output clear pgm statistics on page 4647

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm statistics

```
user@host> clear pgm statistics
```

clear pim join

List of Syntax
Syntax on page 4648
Syntax (EX Series Switch and the QFX Series) on page 4648

Syntax

clear pim join
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear pim join
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Clear the Protocol Independent Multicast (PIM) join and prune states.

Options

none—Clear the PIM join and prune states for all groups, family addresses, and instances.

group-address—(Optional) Clear the PIM join and prune states for a group address.

inet | inet6—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- show pim join on page 4766

List of Sample Output

clear pim join on page 4649

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

clear pim join

```
user@host> clear pim join
```

Copyright © 2016, Juniper Networks, Inc.

clear pim join-distribution

Syntax

clear pim join-distribution
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 10.0.

Description

Redistribute the Protocol Independent Multicast (PIM) join states.

You can find out if there are multiple paths available for a source (for example, an RP) with the output of the **show pim source** command.

When you include the **join-load-balance** statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The **clear pim join-distribution** command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the **clear pim join-distribution** command during a maintenance window.

Options

none—Redistribute the PIM join states for the default master instance.

instance *instance-name*—(Optional) Redistribute the join states for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim join-distribution** command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- show pim neighbors on page 4787
- show pim join on page 4766
- join-load-balance on page 4489

List of Sample Output

clear pim join-distribution on page 4651

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show pim join** command before and after distributing the join state to verify the operation.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

clear pim join-distribution

```
user@host> clear pim join-distribution
```

Copyright © 2016, Juniper Networks, Inc.

clear pim register

List of Syntax
Syntax on page 4652
Syntax (EX Series Switch and the QFX Series) on page 4652
Syntax (PTX Series) on page 4652

Syntax
clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Syntax (PTX Series)
clear pim register
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Clear Protocol Independent Multicast (PIM) register message counters.

Options
none—Clear PIM register message counters for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear register message counters for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM register message counters for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
The **clear pim register** command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level
clear

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- show pim statistics on page 4801

List of Sample Output clear pim register on page 4653

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim register

```
user@host> clear pim register
```

clear pim statistics

List of Syntax
Syntax on page 4654
Syntax (EX Series Switch and the QFX Series) on page 4654

Syntax

clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Clear Protocol Independent Multicast (PIM) statistics.

Options

none—Clear PIM statistics for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear statistics for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM statistics for a specific interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim statistics** command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- show pim statistics on page 4801

List of Sample Output

clear pim statistics on page 4655

Output Fields

See show pim statistics for an explanation of output fields.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                   2111       4222          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune             14200      13115          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0
PIM statistics summary for all interfaces:
Unknown type                              0
V1 Unknown type                           0
Unknown Version                           0
Neighbor unknown                          0
Bad Length                                0
Bad Checksum                              0
Bad Receive If                            0
Rx Intf disabled                       2007
Rx V1 Require V2                          0
Rx Register not RP                        0
RP Filtered Source                        0
Unknown Reg Stop                          0
Rx Join/Prune no state                 1040
Rx Graft/Graft Ack no state               0
...

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                      1          0          0
V1 Register                   0          0          0
...
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

request pim multicast-tunnel rebalance

List of Syntax
Syntax on page 4657
Syntax (EX Series Switches) on page 4657

Syntax
request pim multicast-tunnel rebalance
<instance *instance-name*>
<logical-system (all | *logical-system-name*>

Syntax (EX Series Switches)
request pim multicast-tunnel rebalance
<instance *instance-name*>

Release Information
Command introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 10.2 for EX Series switches.

Description
Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the **show pim interfaces instance** ***instance-name*** command.

Options
none—Re-create and rebalance all tunnel interfaces for all routing instances.

instance ***instance-name***—Re-create and rebalance all tunnel interfaces for a specific instance.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
maintenance

Related Documentation
- show pim interfaces on page 4763

Output Fields
This command produces no output. To verify the operation of the command, run the **show pim interface instance** ***instance-name*** before and after running the **request pim multicast-tunnel rebalance** command.

Copyright © 2016, Juniper Networks, Inc.

show amt statistics

Syntax
show amt statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in JUNOS Release 10.2.

Description
Display information about the Automatic Multicast Tunneling (AMT) protocol tunnel statistics.

Options
none—Display summary information about all AMT Protocol tunnels.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear amt statistics on page 4633
- show amt summary on page 4661
- show amt tunnel on page 4663

List of Sample Output
show amt statistics on page 4659

Output Fields
Table 59 describes the output fields for the **show amt statistics** command. Output fields are listed in the approximate order in which they appear.

Table 389: show amt statistics Output Fields

Field Name	Field Description
AMT receive message count	Summary of AMT statistics for messages received on all interfaces. • **AMT relay discovery**—Number of AMT relay discovery messages received. • **AMT membership request**—Number of AMT membership request messages received. • **AMT membership update**—Number of AMT membership update messages received.
AMT send message count	Summary of AMT statistics for messages sent on all interfaces. • **AMT relay advertisement**—Number of AMT relay advertisement messages sent. • **AMT membership query**—Number of AMT membership query messages sent.

Copyright © 2016, Juniper Networks, Inc.

Table 389: show amt statistics Output Fields *(continued)*

Field Name	Field Description
AMT error message count	Summary of AMT statistics for error messages received on all interfaces. • **AMT incomplete packet**—Number of messages received with length errors so severe that further classification could not occur. • **AMT invalid mac**—Number of messages received with an invalid message authentication code (MAC). • **AMT unexpected type**—Number of messages received with an unknown message type specified. • **AMT invalid relay discovery address**—Number of AMT relay discovery messages received with an address other than the configured anycast address. • **AMT invalid membership request address**—Number of AMT membership request messages received with an address other than the configured AMT local address. • **AMT invalid membership update address**—Number of AMT membership update messages received with an address other than the configured AMT local address. • **AMT incomplete relay discovery messages**—Number of AMT relay discovery messages received that are not fully formed. • **AMT incomplete membership request messages**—Number of AMT membership request messages received that are not fully formed. • **AMT incomplete membership update messages**—Number of AMT membership update messages received that are not fully formed. • **AMT no active gateway**—Number of AMT membership update messages received for a tunnel that does not exist for the gateway that sent the message. • **AMT invalid inner header checksum**—Number of AMT membership update messages received with an invalid IP checksum. • **AMT gateways timed out**—Number of gateways that timed out because of inactivity.

Sample Output

show amt statistics

```
user@host> show amt statistics

AMT receive message count
AMT relay advertisement                 :          2
AMT membership request                  :          5
AMT membership update                   :          5

AMT send message count
AMT relay advertisement                 :          2
AMT membership query                    :          5

AMT error message count
AMT incomplete packet                   :          0
AMT invalid mac                         :          0
AMT unexpected type                     :          0
AMT invalid relay discovery address     :          0
AMT invalid membership request address  :          0
AMT invalid membership update address   :          0
AMT incomplete relay discovery messages :          0
AMT incomplete membership request messages :       0
AMT incomplete membership update messages  :       0
AMT no active gateway                   :          0
```

Copyright © 2016, Juniper Networks, Inc.

```
AMT invalid inner header checksum           :           0
AMT gateways timed out                      :           0
```


Copyright © 2016, Juniper Networks, Inc.

show amt summary

Syntax
show amt summary
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in Junos OS Release 10.2.

Description
Display summary information about the Automatic Multicast Tunneling (AMT) protocol.

Options
none—Display summary information about all AMT protocol instances.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear amt tunnel on page 4634
- show amt statistics on page 4658
- show amt tunnel on page 4663

List of Sample Output
show amt summary on page 4662

Output Fields
Table 390 describes the output fields for the **show amt summary** command. Output fields are listed in the approximate order in which they appear.

Table 390: show amt summary Output Fields

Field Name	Field Description	Level of Output
AMT anycast prefix	Prefix advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways.	All levels
AMT anycast address	Anycast address configured from which the anycast prefix is derived.	All levels
AMT local address	Local unique AMT relay IP address configured. Used to send AMT relay advertisement messages, it is the IP source address of AMT control messages and the source address of the data tunnel encapsulation.	All levels
AMT tunnel limit	Maximum number of AMT tunnels that can be created.	All levels
active tunnels	Number of active AMT tunnel interfaces.	All levels

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show amt summary

```
user@host> show amt summary
 AMT anycast prefix : 20.0.0.4/32
 AMT anycast address : 20.0.0.4
 AMT local address : 20.0.0.4
 AMT tunnel limit : 1000, active tunnels : 2
```


Copyright © 2016, Juniper Networks, Inc.

show amt tunnel

Syntax

show amt tunnel
<brief | detail>
<gateway-address *gateway-ip-address*> <port *port-number*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<tunnel-interface *interface-name*>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels.

Options

none—Display summary information about all AMT protocol instances.

brief | detail—(Optional) Display the specified level of detail.

gateway-address *gateway-ip-address* port *port-number*—(Optional) Display information for the specified AMT gateway only. If no port is specified, display information for all AMT gateways with the given IP address.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

tunnel-interface *interface-name*—(Optional) Display information for the specified AMT tunnel interface only.

Required Privilege Level

view

Related Documentation

- clear amt tunnel on page 4634
- show amt statistics on page 4658
- show amt summary on page 4661

List of Sample Output

show amt tunnel on page 4664
show amt tunnel detail on page 4664
show amt tunnel tunnel-interface on page 4665
show amt tunnel gateway-address on page 4665
show amt tunnel gateway-address detail on page 4665

Output Fields

Table 391 describes the output fields for the **show amt tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 391: show amt tunnel Output Fields

Field Name	Field Description	Level of Output
AMT gateway address	Address of the AMT gateway that is being connected by the AMT tunnel.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 391: show amt tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
port	Client port used by the AMT tunnel.	All levels
AMT tunnel interface	Dynamically created AMT logical interfaces used by the AMT tunnel in the format **ud-FPC/PIC/Port.unit**.	All levels
AMT tunnel state	State of the AMT tunnel. The state is normally **Active**. • **Active**—The tunnel is active. • **Pending**—The tunnel creation is pending. This is a transient state. • **Down**—The tunnel is in the down state. • **Graceful restart pending**—Graceful restart is in progress. • **Reviving**—The routing protocol daemon or Routing Engine was restarted (not gracefully). The tunnel remains in the reviving state until the AMT gateway sends a control message. When the message is received the tunnel is moved to the **Active** state. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
AMT tunnel inactivity timeout	Number of seconds since the most recent control message was received from an AMT gateway. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
Number of groups	Number of multicast groups using the tunnel.	All levels
Group	Multicast group address or addresses using the tunnel.	**detail**
Include Source	Multicast source address for each IGMPv3 group using the tunnel.	**detail**
AMT message count	Statistics for AMT messages: • **AMT Request**—Number of AMT relay tunnel request messages received. • **AMT membership update**—Number of AMT membership update messages received.	All levels

Sample Output

show amt tunnel

```
user@host> show amt tunnel
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/1/10.1120256
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 15
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel detail

```
user@host> show amt tunnel detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
```


Copyright © 2016, Juniper Networks, Inc.

```
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 62
 Number of groups : 1
  Group: 226.2.3.2

AMT message count:
AMT Request    AMT membership update
2              2

 AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel tunnel-interface

```
user@host> show amt tunnel tunnel-interface ud-5/3/10.1120512
 AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 145
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address

```
user@host> show amt tunnel gateway-address 11.11.11.3 port 2268
AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address detail

```
user@host> show amt tunnel gateway-address 11.11.11.2 detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 234
 Number of groups : 1
  Group: 226.2.3.2

AMT message count:
AMT Request    AMT membership update
2              2
```

Copyright © 2016, Juniper Networks, Inc.

show dvmrp interfaces

Syntax
show dvmrp interfaces
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display information about Distance Vector Multicast Routing Protocol (DVMRP)–enabled interfaces.

Options
none—(Same as **logical-system all**) Display information about DVMRP-enabled interfaces.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show dvmrp interfaces on page 4667

Output Fields
Table 392 describes the output fields for the **show dvmrp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 392: show dvmrp interfaces Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **up** or **down**.
Leaf	Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors.
Metric	Interface metric: a value from **1** through **31**.
Announce	Number of routes the interface is announcing.
Mode	DVMRP mode: • **Forwarding**—DVMRP does both the routing and the multicast data forwarding. • **Unicast-routing**—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show dvmrp interfaces

```
user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0    Up    N    1      4 Forwarding
fxp1.0    Up    N    1      4 Forwarding
fxp2.0    Up    N    1      3 Forwarding
lo0.0     Up    Y    1      0 Unicast-routing
```

Copyright © 2016, Juniper Networks, Inc.

show dvmrp neighbors

Syntax show dvmrp neighbors
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display information about Distance Vector Multicast Routing Protocol (DVMRP) neighbors.

Options **none**—(Same as **logical-system all**) Display information about DVMRP neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output show dvmrp neighbors on page 4669

Output Fields Table 393 describes the output fields for the **show dvmrp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 393: show dvmrp neighbors Output Fields

Field Name	Field Description
Neighbor	Address of the neighboring DVMRP router.
Interface	Interface through which the neighbor is reachable.
Version	Version of DVMRP that the neighbor is running, in the format ***majorminor***.
Flags	Information about the neighbor: • **1**—One way. The local router has seen the neighbor, but the neighbor has not seen the local router. • **G**—Neighbor supports generation ID. • **L**—Neighbor is a leaf router. • **M**—Neighbor supports mtrace. • **N**—Neighbor supports netmask in prune messages and graft messages. • **P**—Neighbor supports pruning. • **S**—Neighbor supports SNMP.
Routes	Number of routes learned from the neighbor.
Timeout	How long until the DVMRP neighbor information times out, in seconds.
Transitions	Number of generation ID changes that have occurred since the local router learned about the neighbor.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show dvmrp neighbors

```
user@host> show dvmrp neighbors
Neighbor        Interface        Version Flags   Routes Timeout Transitions
192.168.1.1    ipip.0            3.255  PGM         3      28           1
```

show dvmrp prefix

Syntax
show dvmrp prefix
<brief | detail>
<logical-system (all | *logical-system-name*)>
<*prefix*>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes.

Options
none—Display standard information about all DVMRP prefixes.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display information about specific prefixes.

Required Privilege Level
view

List of Sample Output
show dvmrp prefix on page 4671
show dvmrp prefix brief on page 4671
show dvmrp prefix detail on page 4671

Output Fields
Table 394 describes the output fields for the **show dvmrp prefix** command. Output fields are listed in the approximate order in which they appear.

Table 394: show dvmrp prefix Output Fields

Field Name	Field Description	Level of Output
Prefix	DVMRP route.	All levels
Next hop	Next hop from which the route was learned.	All levels
Age	Last time that the route was refreshed.	All levels
multicast-group	Multicast group address.	**detail**
Prunes sent	Number of prune messages sent to the multicast group.	**detail**
Grafts sent	Number of grafts sent to the multicast group.	**detail**
Cache lifetime	Lifetime of the group in the multicast cache, in seconds.	**detail**
Prune lifetime	Lifetime remaining and total lifetime of prune messages, in seconds.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show dvmrp prefix

```
user@host> show dvmrp prefix
Prefix              Next hop         Age
10.38.0.0      /30 10.38.0.1        00:06:17
10.38.0.4      /30 10.38.0.5        00:06:13
10.38.0.8      /30 10.38.0.2        00:00:04
10.38.0.12     /30 10.38.0.6        00:00:04
10.255.14.114  /32 10.255.14.114    00:06:17
10.255.14.142  /32 10.38.0.2        00:00:04
10.255.14.144  /32 10.38.0.2        00:00:04
10.255.70.15   /32 10.38.0.6        00:00:04
192.168.14.0   /24 192.168.14.114   00:06:17
192.168.195.40 /30 192.168.195.41   00:06:17
192.168.195.92 /30 10.38.0.2        00:00:04
```

show dvmrp prefix brief

The output for the **show dvmrp prefix brief** command is identical to that for the **show dvmrp prefix** command.

show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix              Next hop         Age
10.38.0.0      /30 10.38.0.1        00:06:28
10.38.0.4      /30 10.38.0.5        00:06:24
10.38.0.8      /30 10.38.0.2        00:00:15
10.38.0.12     /30 10.38.0.6        00:00:15
10.255.14.114  /32 10.255.14.114    00:06:28
10.255.14.142  /32 10.38.0.2        00:00:15
10.255.14.144  /32 10.38.0.2        00:00:15
10.255.70.15   /32 10.38.0.6        00:00:15
192.168.14.0   /24 192.168.14.114   00:06:28
192.168.195.40 /30 192.168.195.41   00:06:28
192.168.195.92 /30 10.38.0.2        00:00:15
```

Copyright © 2016, Juniper Networks, Inc.

show dvmrp prunes

Syntax
show dvmrp prunes
<all | rx | tx>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prune messages.

Options
none—Display received and transmitted DVMRP prune information.

all—(Optional) Display information about all received and transmitted prune messages.

rx—(Optional) Display information about received prune messages.

tx—(Optional) Display information about transmitted prune messages.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show dvmrp prunes on page 4672

Output Fields
Table 395 describes the output fields for the **show dvmrp prunes** command. Output fields are listed in the approximate order in which they appear.

Table 395: show dvmrp prunes Output Fields

Field Name	Field Description
Group	Group address.
Source prefix	Prefix for the prune.
Timeout	How long until the prune message expires, in seconds.
Neighbor	Neighbor to which the prune was sent or from which the prune was received.

Sample Output

show dvmrp prunes

```
user@host> show dvmrp prunes
Group           Source prefix      Timeout Neighbor
224.0.1.1       128.112.0.0    /12    7077 192.168.1.1
224.0.1.32      160.0.0.0      /3     7087 192.168.1.1
224.2.123.4     136.0.0.0      /5     6955 192.168.1.1
224.2.127.1     129.0.0.0      /8     7046 192.168.1.1
```


Copyright © 2016, Juniper Networks, Inc.

```
224.2.135.86    128.102.128.0  /17    7071 192.168.1.1
224.2.135.86    129.0.0.0      /8     7074 192.168.1.1
224.2.135.86    130.0.0.0      /7     7071 192.168.1.1
...
```

Copyright © 2016, Juniper Networks, Inc.

show igmp group

List of Syntax
Syntax on page 4674
Syntax (EX Series Switch and the QFX Series) on page 4674

Syntax
show igmp group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show igmp group
<brief | detail>
<*group-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display Internet Group Management Protocol (IGMP) group membership information.

Options
none—Display standard information about membership for all IGMP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display group membership for the specified IP address only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show igmp group (Include Mode) on page 4675
show igmp group (Exclude Mode) on page 4676
show igmp group brief on page 4676
show igmp group detail on page 4676

Output Fields
Table 396 describes the output fields for the **show igmp group** command. Output fields are listed in the approximate order in which they appear.

Table 396: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 396: show igmp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Last reported by: Local
    Group timeout:       0 Type: Dynamic
Group: 224.0.0.22
    Group mode: Exclude
    Source: 0.0.0.0
    Source timeout: 0
    Last reported by: Local
    Group timeout:       0 Type: Dynamic
```

Copyright © 2016, Juniper Networks, Inc.

show igmp interface

List of Syntax
Syntax on page 4678
Syntax (EX Series Switch and the QFX Series) on page 4678

Syntax
show igmp interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show igmp interface
<brief | detail>
<*interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options
none—Display standard information about all IGMP-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified IGMP-enabled interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear igmp membership on page 4635

List of Sample Output
show igmp interface on page 4680
show igmp interface brief on page 4680
show igmp interface detail on page 4681

Output Fields
Table 397 describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Table 397: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 397: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: **1** , **2** , or **3**.	All levels
Groups	Number of groups on the interface.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.	All levels
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic.	All levels
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 397: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Configured Parameters	Information configured by the user: • **IGMP Query Interval**—Interval (in seconds) at which this router sends membership queries when it is the querier. • **IGMP Query Response Interval**—Time (in seconds) that the router waits for a report in response to a general query. • **IGMP Last Member Query Interval**—Time (in seconds) that the router waits for a report in response to a group-specific query. • **IGMP Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information: • **IGMP Membership Timeout**—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • **IGMP Other Querier Present Timeout**—Time (in seconds) that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp interface

```
user@host> show igmp interface
Interface: at-0/3/1.0
    Querier: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
    Querier: 10.111.10.1
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
    Querier: 10.111.20.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 4680.

Copyright © 2016, Juniper Networks, Inc.

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 4680.

Copyright © 2016, Juniper Networks, Inc.

show igmp snooping interface

Syntax
show igmp snooping interface *interface-name*
<brief | detail>
<bridge-domain *bridge-domain-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information
Command introduced in Junos OS Release 8.5.

Description
Display IGMP snooping interface information.

Options
none—Display detailed information.

brief | detail—(Optional) Display the specified level of output.

bridge-domain *bridge-domain-name*—(Optional) Display information about a particular bridge domain.

virtual-switch *virtual-switch-name*—(Optional) Display information about a particular virtual switch.

vlan-id *vlan-identifier*—(Optional) Display information about a particular VLAN.

Required Privilege Level
view

Related Documentation
- show igmp snooping membership on page 4685
- show igmp snooping statistics on page 4689

List of Sample Output
show igmp snooping interface on page 4683
show igmp snooping interface (Group Limit Configured) on page 4684

Output Fields
Table 398 lists the output fields for the **show igmp snooping interface** command. Output fields are listed in the approximate order in which they appear.

Table 398: show igmp snooping interface Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels
IGMP Query Interval	Frequency (in seconds) with which this router sends membership queries when it is the querier.	**detail**
IGMP Query Response Interval	Time (in seconds) that the router waits for a response to a general query.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 398: show igmp snooping interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
IGMP Last Member Query Interval	Time (in seconds) that the router waits for a report in response to a group-specific query.	**detail**
IGMP Robustness Count	Number of times the router retries a query.	**detail**
immediate-leave	State of immediate leave: **On** or **Off**.	All levels
router-interface	Router interfaces that are part of this learning domain.	All levels
Group limit	Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown.	All levels
interface	Interfaces that are being snooped in this learning domain.	All levels
Groups	Number of groups on the interface.	none
State	State of the interface: **Up** or **Down**.	none
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
IGMP Membeship Timeout	Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.	none
IGMP Other Querier Present Timeout	Time that the router waits for the IGMP querier to send a query.	none

Sample Output

show igmp snooping interface

```
user@host> show igmp snooping interface
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/0.200
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ge-0/1/2.200
    State:         Up Groups:      2
    Immediate leave: On
    Router interface: no
Interface: ge-0/1/3.200
    State:         Up Groups:      1
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 130.0
IGMP Query Response Interval: 15.0
```

Copyright © 2016, Juniper Networks, Inc.

```
IGMP Last Member Query Interval: 2.0
IGMP Robustness Count: 3

Derived Parameters:
IGMP Membership Timeout: 405.0
IGMP Other Querier Present Timeout: 397.500
```

Sample Output

show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ge-1/3/8.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
    Group limit:   1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```


Copyright © 2016, Juniper Networks, Inc.

show igmp snooping membership

Syntax

show igmp snooping membership
<brief | detail>
<bridge-domain *bridge-domain-name*>
<group *group-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display IGMP snooping membership information.

Options

none—Display detailed information.

brief | detail—(Optional) Display the specified level of output.

bridge-domain *bridge-domain-name*—(Optional) Display information about a particular bridge domain.

group *group-name* —(Optional) Display information about this group address.

virtual-switch *virtual-switch-name*—(Optional) Display information about a particular virtual switch.

vlan-id *vlan-identifier*—(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Related Documentation

- show igmp snooping interface on page 4682
- show igmp snooping statistics on page 4689
- clear igmp snooping membership on page 4638

List of Sample Output

show igmp snooping membership on page 4686
show igmp snooping membership (Exclude Mode) on page 4687
show igmp snooping membership interface ge-0/1/2.200 on page 4687
show igmp snooping membership vlan-id 1 on page 4687

Output Fields

Table 399 lists the output fields for the **show igmp snooping membership** command. Output fields are listed in the approximate order in which they appear.

Table 399: show igmp snooping membership Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 399: show igmp snooping membership Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface	Interface on which this router is a proxy.	**detail**
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
Group	Multicast group address in the membership database.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address used on queries.	**detail**
Last reported by	Address of source last replying to the query.	**detail**
Group Timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	All levels
Timeout	Length of time (in seconds) left until the entry is purged.	**detail**
Type	Way that the group membership information was learned: • **Dynamic**—Group membership was learned by the IGMP protocol. • **Static**—Group membership was learned by configuration.	**detail**
Include receiver	Source address of receiver included in membership with timeout (in seconds).	**detail**

Sample Output

show igmp snooping membership

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 225.10.10.1
        Group mode: Exclude
        Source: 0.0.0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
        Last reported by: 100.6.85.2
        Group timeout:     173 Type: Dynamic
```

show igmp snooping membership (Exclude Mode)

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 225.10.10.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 100.6.85.2
        Group timeout:     173 Type: Dynamic
```

show igmp snooping membership interface ge-0/1/2.200

```
user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/2.200
    Group: 225.1.1.1
        Source: 0.0.0.0
        Timeout:     391 Type: Static
    Group: 232.1.1.1
        Source: 192.168.1.1
        Timeout:       0 Type: Static
```

show igmp snooping membership vlan-id 1

```
user@host> show igmp snooping membership vlan–id 1
Instance: vpls2

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 225.10.10.1
        Group mode: Exclude
        Source: 0.0.0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
Last reported by: 100.6.85.2
Group timeout:     209 Type: Dynamic
```


Copyright © 2016, Juniper Networks, Inc.

show igmp snooping statistics

Syntax
show igmp snooping statistics
<brief | detail>
<bridge-domain *bridge-domain-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information
Command introduced in Junos OS Release 8.5.

Description
Display IGMP snooping statistics.

Options
none—(Optional) Display detailed information.

brief | detail—(Optional) Display the specified level of output.

bridge-domain *bridge-domain-name*—(Optional) Display information about a particular bridge domain.

virtual-switch *virtual-switch-name*—(Optional) Display information about a particular virtual switch.

vlan-id *vlan-identifier*—(Optional) Display information about a particular VLAN.

Required Privilege Level
view

Related Documentation
- show igmp snooping interface on page 4682
- show igmp snooping membership on page 4685
- clear igmp snooping statistics on page 4639

List of Sample Output
show igmp snooping statistics on page 4690

Output Fields
Table 400 lists the output fields for the **show igmp snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 400: show igmp snooping statistics Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
IGMP packet statistics	Heading for IGMP snooping statistics for all interfaces or for the specified interface.	All levels
learning-domain	Appears at end of "IGMP packets statistics" line.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 400: show igmp snooping statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).	All levels
Received	Number of messages received.	All levels
Sent	Number of messages sent.	All levels
Rx errors	Number of received packets that contained errors.	All levels
IGMP Global Statistics	Summary of IGMP snooping statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Rx non-local**—Number of messages received from senders that are not local.	All levels

Sample Output

show igmp snooping statistics

```
user@host> show igmp snooping statistics
Routing-instance foo

 IGMP packet statistics for all interfaces in learning-domain vlan-100

 IGMP Message type      Received       Sent  Rx errors
 Membership Query             89         51          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
```


Copyright © 2016, Juniper Networks, Inc.

```
  PIM V1                          0          0          0
  Cisco Trace                     0          0          0
  V2 Membership Report          139          0          0
  Group Leave                     0          0          0
  Domain Wide Report              0          0          0
  V3 Membership Report          136          0          0
  Other Unknown types                                    0
  IGMP v3 unsupported type                               0
  IGMP v3 source required for SSM                       23
  IGMP v3 mode not applicable for SSM                    0

  IGMP Global Statistics
  Bad Length                      0
  Bad Checksum                    0
  Rx non-local                    0

Routing-instance bar

  IGMP packet statistics for all interfaces in learning-domain vlan-100

  IGMP Message type      Received       Sent  Rx errors
  Membership Query             89         51          0
  V1 Membership Report          0          0          0
  DVMRP                         0          0          0
  PIM V1                        0          0          0
  Cisco Trace                   0          0          0
  V2 Membership Report        139          0          0
  Group Leave                   0          0          0
  Domain Wide Report            0          0          0
  V3 Membership Report        136          0          0
  Other Unknown types                                  0
  IGMP v3 unsupported type                             0
  IGMP v3 source required for SSM                     23
  IGMP v3 mode not applicable for SSM                  0

  IGMP Global Statistics
  Bad Length                    0
  Bad Checksum                  0
  Rx non-local                  0
```

Copyright © 2016, Juniper Networks, Inc.

show igmp statistics

List of Syntax
Syntax on page 4692
Syntax (EX Series Switch and the QFX Series) on page 4692

Syntax
show igmp statistics
<brief | detail>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show igmp statistics
<brief | detail>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display Internet Group Management Protocol (IGMP) statistics.

Options
none—Display IGMP statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display IGMP statistics about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear igmp statistics on page 4640

List of Sample Output
show igmp statistics on page 4693
show igmp statistics interface on page 4694

Output Fields
Table 401 describes the output fields for the **show igmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 401: show igmp statistics Output Fields

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Copyright © 2016, Juniper Networks, Inc.

Table 401: show igmp statistics Output Fields *(continued)*

Field Name	Field Description
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Mtrace Response**—Number of Mtrace response messages sent or received. • **Mtrace Request**—Number of Mtrace request messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
IGMP Global Statistics	Summary of IGMP statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for IGMP. • **Rx non-local**—Number of messages received from senders that are not local. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the IGMP group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
```

Copyright © 2016, Juniper Networks, Inc.

```
DVMRP                          0          0          0
PIM V1                         0          0          0
Cisco Trace                    0          0          0
V2 Membership Report           0          0          0
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                     0
Bad Checksum                   0
Bad Receive If                 0
Rx non-local                1227
Timed out                      0
Rejected Report                0
Total Interfaces               2
```

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0        230          0
V1 Membership Report          0          0          0
```


Copyright © 2016, Juniper Networks, Inc.

show mld group

Syntax
show mld group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display information about Multicast Listener Discovery (MLD) group membership.

Options
none—Display standard information about all MLD groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display MLD information about the specified group.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear mld membership on page 4642

List of Sample Output
show mld group (Include Mode) on page 4696
show mld group (Exclude Mode) on page 4697
show mld group brief on page 4697
show mld group detail (Include Mode) on page 4697
show mld group detail (Exclude Mode) on page 4698

Output Fields
Table 402 describes the output fields for the **show mld group** command. Output fields are listed in the approximate order in which they appear.

Table 402: show mld group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the MLD membership report; **local** means that the local router joined the group itself.	All levels
Group	Group address.	All levels
Source	Source address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Last reported by	Address of the host that last reported membership in this group.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 402: show mld group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show mld group (Include Mode)

```
user@host> show mld group
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     245 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     241 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```


Copyright © 2016, Juniper Networks, Inc.

show mld group (Exclude Mode)

```
user@host> show mld group
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:     245 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:      28 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see show mld group (Include Mode) on page 4696
show mld group (Exclude Mode) on page 4697.

show mld group detail (Include Mode)

```
user@host> show mld group detail
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     224 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     220 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
Interface: so-1/0/1.0
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::280:42ff:fe15:f445
        Timeout:     258 Type: Dynamic
Interface: local
```

Copyright © 2016, Juniper Networks, Inc.

```
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group detail (Exclude Mode)

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Group timeout:     226 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Group timeout:     246 Type: Dynamic
Interface: local
    Group: ff02::2
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
```


Copyright © 2016, Juniper Networks, Inc.

show mld interface

Syntax
show mld interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display information about Multicast Listener Discovery (MLD)-enabled interfaces.

Options
none—Display standard information about all MLD-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear mld membership on page 4642

List of Sample Output
show mld interface on page 4701
show mld interface brief on page 4701
show mld interface detail on page 4701

Output Fields
Table 403 describes the output fields for the **show mld interface** command. Output fields are listed in the approximate order in which they appear.

Table 403: show mld interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the router that has been elected to send membership queries.	All levels
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the interface.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy at the MLD interface.	All levels
Timeout	How long until the MLD querier is declared to be unreachable, in seconds.	All levels
Version	MLD version being used on the interface: **1** or **2**.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 403: show mld interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Groups	Number of groups on the interface.	All levels
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map used on the interface, if configured.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • **Off**—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Configured Parameters	Information configured by the user. • **MLD Query Interval (.1 secs)**—Interval at which this router sends membership queries when it is the querier. • **MLD Query Response Interval (.1 secs)**—Time that the router waits for a report in response to a general query. • **MLD Last Member Query Interval (.1 secs)**—Time that the router waits for a report in response to a group-specific query. • **MLD Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information. • **MLD Membership Timeout (.1 secs)**—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed. • **MLD Other Querier Present Timeout (.1 secs)**—Time that the router waits for the IGMP querier to send a query.	All levels

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show mld interface

```
user@host> show mld interface
Interface: fe-0/0/0
    Querier: None
    State: Up         Timeout:       0    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
    Querier:  8038::c0a8:c345
    State: Up         Timeout:    None    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
    Querier:  ::192.168.195.73
    State: Up         Timeout:    None    Version:  1    Groups:      3
    SSM Map Policy: ssm-policy-C
    SSM map: ipv6map1
Immediate Leave: On

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550
```

show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see show mld interface on page 4701.

show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see show mld interface on page 4701.

Copyright © 2016, Juniper Networks, Inc.

show mld statistics

Syntax

show mld statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about Multicast Listener Discovery (MLD) statistics.

Options

none—Display MLD statistics for all interfaces.

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear mld statistics on page 4643

List of Sample Output

show mld statistics on page 4703
show mld statistics interface on page 4704

Output Fields

Table 404 describes the output fields for the **show mld statistics** command. Output fields are listed in the approximate order in which they appear.

Table 404: show mld statistics Output Fields

Field Name	Field Description
Received	Number of received packets.
Sent	Number of transmitted packets.
Rx errors	Number of received packets that contained errors.

Copyright © 2016, Juniper Networks, Inc.

Table 404: show mld statistics Output Fields *(continued)*

Field Name	Field Description
MLD Message type	Summary of MLD statistics. • **Listener Query (v1/v2)**—Number of membership queries sent and received. • **Listener Report (v1)**—Number of version 1 membership reports sent and received. • **Listener Done (v1/v2)**—Number of Listener Done messages sent and received. • **Listener Report (v2)**—Number of version 2 membership reports sent and received. • **Other Unknown types**—Number of unknown message types received. • **MLD v2 source required for SSM**—Number of MLD version 2 messages received that contained no source. • **MLD v2 mode not applicable for SSM**—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).
MLD Global Statistics	Summary of MLD statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with an invalid IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for MLD. • **Rx non-local**—Number of messages received from nonlocal senders. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the MLD group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show mld statistics

```
user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          2          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               0          0          0
Other Unknown types                                      0
MLD v2 source required for SSM     2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                         0
Bad Checksum                       0
Bad Receive If                     0
Rx non-local                       0
Timed out                          0
```

Copyright © 2016, Juniper Networks, Inc.

```
Rejected Report                   0
Total Interfaces                  2
```

show mld statistics interface

```
user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type             Received       Sent  Rx errors
Listener Query (v1/v2)              0          2          0
Listener Report (v1)                0          0          0
Listener Done (v1/v2)               0          0          0
Listener Report (v2)                0          0          0
Other Unknown types                                        0
MLD v2 source required for SSM      2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                          0
Bad Checksum                        0
Bad Receive If                      0
Rx non-local                        0
Timed out                           0
Rejected Report                     0
Total Interfaces                    2
```


Copyright © 2016, Juniper Networks, Inc.

show msdp

Syntax

show msdp
<brief | detail>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display Multicast Source Discovery Protocol (MSDP) information.

Options

none—Display standard MSDP information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display information about the specified peer only,

Required Privilege Level

view

Related Documentation

- show msdp source on page 4707
- show msdp source-active on page 4709
- show msdp statistics on page 4711

List of Sample Output

show msdp on page 4706
show msdp brief on page 4706
show msdp detail on page 4706

Output Fields

Table 405 describes the output fields for the **show msdp** command. Output fields are listed in the approximate order in which they appear.

Table 405: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: **Listen**, **Established**, or **Inactive**.	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 405: show msdp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format *number-accepted/number-received*.	All levels
Peer Connect Retries	Number of peer connection retries.	**detail**
State timer expires	Number of seconds before another message is sent to a peer.	**detail**
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	**detail**
SA accepted	Number of entries in the source-active cache accepted from the peer.	**detail**
SA received	Number of entries in the source-active cache received by the peer.	**detail**

Sample Output

show msdp

```
user@host> show msdp
Peer address     Local address   State       Last up/down Peer-Group SA Count
198.32.8.193     198.32.8.195    Established  5d 19:25:44 North23    120/150
198.32.8.194     198.32.8.195    Established  3d 19:27:27 North23    300/345
198.32.8.196     198.32.8.195    Established  5d 19:39:36 North23    10/13
198.32.8.197     198.32.8.195    Established  5d 19:32:27 North23    5/6
198.32.8.198     198.32.8.195    Established  3d 19:33:04 North23    2305/3000
```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see show msdp on page 4706.

show msdp detail

```
user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0
```


Copyright © 2016, Juniper Networks, Inc.

show msdp source

Syntax

show msdp source
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).

Options

none—Display standard MSDP source information for all routing instances.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-address—(Optional) IP address and optional prefix length. Display information for the specified source address only.

Required Privilege Level

view

Related Documentation

- show msdp on page 4705
- show msdp source-active on page 4709
- show msdp statistics on page 4711

List of Sample Output

show msdp source on page 4708

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 406 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 406: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: • **Configured**—Source-active limit explicitly configured for this source. • **Dynamic**—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```
user@host> show msdp source
Source address /Len   Type       Maximum   Threshold    Exceeded
0.0.0.0        /0     Configured        5       none           0
10.1.0.0       /16    Configured      500       none           0
10.1.1.1       /32    Configured    10000       none           0
10.1.1.2       /32    Dynamic        6936       none           0
10.1.5.5       /32    Dynamic         500       none         123
10.2.1.1       /32    Dynamic           2       none           0
```


Copyright © 2016, Juniper Networks, Inc.

show msdp source-active

Syntax

show msdp source-active
<brief | detail>
<group *group*>
<instance *instance-name*>
<local>
<logical-system (all | *logical-system-name*)>
<originator *originator*>
<peer *peer-address*>
<source *source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

none—Display standard MSDP source-active cache information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

group *group*—(Optional) Display source-active cache information for the specified group.

instance *instance-name*—(Optional) Display information for the specified instance.

local—(Optional) Display all source-active caches originated by this router.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

originator *originator*—(Optional) Display information about the peer that originated the source-active cache entries.

peer *peer-address*—(Optional) Display the source-active cache of the specified peer.

source *source-address*—(Optional) Display the source-active cache of the specified source.

Required Privilege Level

view

Related Documentation

- show msdp on page 4705
- show msdp source on page 4707
- show msdp statistics on page 4711

List of Sample Output

show msdp source-active on page 4710
show msdp source-active brief on page 4710
show msdp source-active detail on page 4710

Output Fields

Table 407 describes the output fields for the **show msdp source-active** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 407: show msdp source-active Output Fields

Field Name	Field Description
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Address of the rendezvous point (RP) that originated the message.
Flags	Flags: **Accept**, **Reject**, or **Filtered**.

Sample Output

show msdp source-active

```
user@host> show msdp source-active
Group address   Source address  Peer address    Originator      Flags
230.0.0.0       192.168.195.46  local           10.255.14.30    Accept
230.0.0.1       192.168.195.46  local           10.255.14.30    Accept
230.0.0.2       192.168.195.46  local           10.255.14.30    Accept
230.0.0.3       192.168.195.46  local           10.255.14.30    Accept
230.0.0.4       192.168.195.46  local           10.255.14.30    Accept
```

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 4710.

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 4710.

Copyright © 2016, Juniper Networks, Inc.

show msdp statistics

Syntax

show msdp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display statistics about Multicast Source Discovery Protocol (MSDP) peers.

Options

none—Display statistics about all MSDP peers for all routing instances.

instance *instance-name*—(Optional) Display statistics about a specific MSDP instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display statistics about a particular MSDP peer.

Required Privilege Level

view

Related Documentation

- *clear msdp statistics*

List of Sample Output

show msdp statistics on page 4713
show msdp statistics peer on page 4713

Output Fields

Table 408 describes the output fields for the **show msdp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 408: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.

Copyright © 2016, Juniper Networks, Inc.

Table 408: show msdp statistics Output Fields *(continued)*

Field Name	Field Description
Peer	Address of peer.
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
SA response messages received	Entry Count is a field within SA message that defines how many source/group tuples are present in the SA message. The counter is incremented each time an SA with an Entry Count of zero is received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.
Keepalive messages sent	Number of keepalive messages sent.

Copyright © 2016, Juniper Networks, Inc.

Table 408: show msdp statistics Output Fields *(continued)*

Field Name	Field Description
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```
user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA messages with zero Entry Count received: 0
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

show msdp statistics peer

```
user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
    Last State Change: 8:19:23 (00:01:08)
    Last message received from peer: 8:20:05 (00:00:26)
    RPF Failures: 0
    Remote Closes: 0
    Peer Timeouts: 0
    SA messages sent: 17
    SA messages received: 16
    SA request messages sent: 0
    SA request messages received: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 20
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval: 120
Keepalive messages sent: 0
Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0
```


Copyright © 2016, Juniper Networks, Inc.

show multicast backup-pe-groups

Syntax

show multicast backup-pe-groups
<address *pe-address*>
<group *group-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 9.0.

Description

Display backup PE router group information when ingress PE redundancy is configured. Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution.

Options

none—Display standard information about all backup PE groups.

address *pe-address*—(Optional) Display the groups that a PE address is associated with.

group *group*—(Optional) Display the backup PE group information for a particular group.

instance *instance-name*—(Optional) Display backup PE group information for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast backup-pe-groups on page 4716

Output Fields

Table 409 describes the output fields for the **show multicast backup-pe-groups** command. Output fields are listed in the approximate order in which they appear.

Table 409: show multicast backup-pe-groups Output Fields

Field Name	Field Description
Backup PE Group	Group name.
Designated PE	Primary PE router. Address of the PE router that is currently forwarding traffic on the static route.
Transitions	Number of times that the designated PE router has transitioned from the most eligible PE router to a backup PE router and back again to the most eligible PE router.
Last Transition	Time of the most recent transition.
Local Address	Address of the local PE router.
Backup PE List	List of PE routers that are configured to be backups for the group.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show multicast backup-pe-groups

```
user@host> show multicast backup-pe-groups
Instance: master

Backup PE group: b1
    Designated PE:   10.255.165.7
    Transitions:     1
    Last Transition: 03:15:01
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.8

Backup PE group: b2
    Designated PE:   10.255.165.7
    Transitions:     2
    Last Transition: 02:58:20
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.9
                     10.255.165.8
```


Copyright © 2016, Juniper Networks, Inc.

show multicast flow-map

List of Syntax
Syntax on page 4717
Syntax (EX Series Switch and the QFX Series) on page 4717

Syntax
show multicast flow-map
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast flow-map
<brief | detail>

Release Information
Command introduced in Junos OS Release 8.2.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display configuration information about IP multicast flow maps.

Options
none—Display configuration information about IP multicast flow maps on all systems.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast flow-map on page 4718
show multicast flow-map detail on page 4718

Output Fields
Table 410 describes the output fields for the **show multicast flow-map** command. Output fields are listed in the approximate order in which they appear.

Table 410: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated with the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none
Flow-map	Name of the flow map.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 410: show multicast flow-map Output Fields *(continued)*

Field Name	Field Description	Levels of Output
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	**detail**
Redundant Sources	Redundant sources defined for the same destination group.	**detail**

Sample Output

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name            Policy                  Cache timeout        Bandwidth Adaptive
map2            policy2                         never          2000000 no
map1            policy1                    60 seconds          2000000 no
```

Sample Output

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
    Policy:             policy1
    Cache Timeout:      600 seconds
    Bandwidth:          2000000
    Adaptive Bandwidth: yes
    Redundant Sources:  11.11.11.11
    Redundant Sources:  11.11.11.12
    Redundant Sources:  11.11.11.13
```


Copyright © 2016, Juniper Networks, Inc.

show multicast interface

List of Syntax
Syntax on page 4719
Syntax (EX Series Switch and the QFX Series) on page 4719

Syntax
show multicast interface
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast interface

Release Information
Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display bandwidth information about IP multicast interfaces.

Options
none—Display all interfaces that have multicast configured.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast interface on page 4720

Output Fields
Table 411 describes the output fields for the **show multicast interface** command. Output fields are listed in the approximate order in which they appear.

Table 411: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.

Copyright © 2016, Juniper Networks, Inc.

Table 411: show multicast interface Output Fields *(continued)*

Field Name	Field Description
Local bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping	State of the reverse OIF mapping feature (**on** or **off**). NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping no QoS adjustment	State of the no QoS adjustment feature (**on** or **off**) for interfaces that are using reverse OIF mapping. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Leave timer	Amount of time a mapped interface remains active after the last mapping ends. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
No QoS adjustment	State (**on**) of the no QoS adjustment feature when this feature is enabled. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.

Sample Output

show multicast interface

```
user@host> show multicast interface
Interface               Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3                              10000000                          0
fe-0/0/3.210                          10000000                   -2000000
fe-0/0/3.220                         100000000                  100000000
fe-0/0/3.230                          20000000                   18000000
fe-0/0/2.200                         100000000                  100000000
```


Copyright © 2016, Juniper Networks, Inc.

show multicast pim-to-igmp-proxy

List of Syntax
Syntax on page 4721
Syntax (EX Series Switch and the QFX Series) on page 4721

Syntax

show multicast pim-to-igmp-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast pim-to-igmp-proxy
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.0.
instance option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.

Options

none—Display configuration information about PIM-to-IGMP message translation for all routing instances.

instance *instance-name*—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output
show multicast pim-to-igmp-proxy on page 4722
show multicast pim-to-igmp-proxy instance on page 4722

Output Fields

Table 412 describes the output fields for the **show multicast pim-to-igmp-proxy** command. Output fields are listed in the order in which they appear.

Table 412: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is **master** (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```


Copyright © 2016, Juniper Networks, Inc.

show multicast pim-to-mld-proxy

List of Syntax

Syntax on page 4723
Syntax (EX Series Switch and the QFX Series) on page 4723

Syntax

show multicast pim-to-mld-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast pim-to-mld-proxy
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.0.
instance option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.

Options

none—Display configuration information about PIM-to-MLD message translation for all routing instances.

instance ***instance-name***—(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast pim-to-mld-proxy on page 4724
show multicast pim-to-mld-proxy instance on page 4724

Output Fields

Table 413 describes the output fields for the **show multicast pim-to-mld-proxy** command. Output fields are listed in the order in which they appear.

Table 413: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```


Copyright © 2016, Juniper Networks, Inc.

show multicast route

List of Syntax
Syntax on page 4725
Syntax (EX Series Switch and the QFX Series) on page 4725

Syntax

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>
<source-prefix *source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<*regular-expression*>
<source-prefix *source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

Options

none—Display standard information about all entries in the multicast forwarding table for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

group *group*—(Optional) Display the cache entries for a particular group.

inet | inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Copyright © 2016, Juniper Networks, Inc.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

source-prefix *source-prefix*—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

Related Documentation
- Example: Configuring Bidirectional PIM on page 4199

List of Sample Output
show multicast route on page 4727
show multicast route (Bidirectional PIM) on page 4728
show multicast route brief on page 4728
show multicast route detail on page 4729
show multicast route extensive (Bidirectional PIM) on page 4729
show multicast route instance <instance-name> extensive on page 4730
show multicast route summary on page 4730

Output Fields Table 414 describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 414: show multicast route Output Fields

Field Name	Field Description	Level of Output
family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Upstream rpf interface list	When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded.	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	**extensive**
Session description	Name of the multicast session.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 414: show multicast route Output Fields *(continued)*

Field Name	Field Description	Level of Output
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays **Forwarding statistics are not available**. NOTE: On QFX Series switches and OCX Series switches, this field does not report valid statistics.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	**detail extensive**
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source. When the **show multicast route extensive** command is used with the **display-origin-protocol** option, the field name is only **Protocol** and not **Upstream Protocol**. However, this field also displays the protocol that installed the active route.	**detail extensive**
Route type	Type of multicast route. Values can be (S,G) or (*,G).	**summary**
Route state	Whether the group is **Active** or **Inactive**.	**summary extensive**
Route count	Number of multicast routes.	**summary**
Forwarding state	Whether the prefix is pruned or forwarding.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry. A value of **forever** indicates routes that do not have keepalive times.	**extensive**
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	**extensive**
Uptime	Time since the creation of a multicast route.	**extensive**

Sample Output

show multicast route

```
user@host> show multicast route
Family: INET

Group: 228.0.0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344

Family: INET6
```

show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 224.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0

Group: 225.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 225.1.3.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
Family: INET6
```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see show multicast route on page 4727 or show multicast route (Bidirectional PIM) on page 4728.

 Copyright © 2016, Juniper Networks, Inc.

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
    Statistics: 8 kBps, 100 pps, 45272 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344
    Session description: Administratively Scoped
    Statistics: 46 kBps, 1000 pps, 921077 packets

    Next-hop ID: 262143
    Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 585
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
```

```
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 589
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

show multicast route instance <instance-name> extensive

```
user@host> show multicast route instance mvpn extensive
Family: INET

Group: 239.10.10.10
    Source: 2.0.0.2/32
    Upstream interface: xe-0/0/0.102
    Downstream interface list:
        xe-10/3/0.0 xe-0/3/0.0 xe-0/0/0.106 xe-0/0/0.105
        xe-0/0/0.103 xe-0/0/0.104 xe-0/0/0.107 xe-0/0/0.108
    Session description: Administratively Scoped
    Statistics: 256 kBps, 3998 pps, 670150 packets
    Next-hop ID: 1048579
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 58
    Uptime: 00:00:04
```

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET

Route type   Route state    Route count
(S,G)             Active             2
(S,G)             Inactive           3

Instance: master Family: INET6
```


Copyright © 2016, Juniper Networks, Inc.

show multicast rpf

List of Syntax
Syntax on page 4731
Syntax (EX Series Switch and the QFX Series) on page 4731

Syntax

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<summary>

Syntax (EX Series Switch and the QFX Series)

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<*prefix*>
<summary>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display information about multicast reverse-path-forwarding (RPF) calculations.

Options

none—Display RPF calculation information for all supported address families.

inet | inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about multicast RPF calculations for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display the RPF calculation information for the specified prefix.

summary—(Optional) Display a summary of all multicast RPF information.

Required Privilege Level

view

List of Sample Output

show multicast rpf on page 4732
show multicast rpf inet6 on page 4733
show multicast rpf prefix on page 4734
show multicast rpf summary on page 4734

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 415 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 415: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.
Interface	Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the **show pim join extensive** command when bidirectional PIM is configured.
Neighbor	Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the **show pim join extensive** command when bidirectional PIM is configured.

Sample Output

show multicast rpf

```
user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
    Protocol: Static

10.255.14.132/32
    Protocol: Direct
    Interface: lo0.0

10.255.245.91/32
    Protocol: IS-IS
    Interface: so-1/1/1.0
    Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21
```

show multicast rpf inet6

```
user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
    Protocol: Direct
    Interface: lo0.0

::10.255.245.91/128
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0

::192.168.195.22/128
Protocol: Local

::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0

::192.168.195.77/128
Protocol: Local
```

Copyright © 2016, Juniper Networks, Inc.

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```


Copyright © 2016, Juniper Networks, Inc.

show multicast scope

List of Syntax
Syntax on page 4735
Syntax (EX Series Switch and the QFX Series) on page 4735

Syntax
show multicast scope
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast scope
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display administratively scoped IP multicast information.

Options
none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.

inet | inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display administratively scoped information for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast scope on page 4736
show multicast scope inet on page 4736
show multicast scope inet6 on page 4736

Output Fields
Table 416 describes the output fields for the **show multicast scope** command. Output fields are listed in the approximate order in which they appear.

Table 416: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.

Copyright © 2016, Juniper Networks, Inc.

Table 416: show multicast scope Output Fields *(continued)*

Field Name	Field Description
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
232-net              232.232.0.0/16     fe-0/0/0.1              0
local                239.255.0.0/16     fe-0/0/0.1              0
local                ff05::/16          fe-0/0/0.1              0
larry                ff05::1234/128     fe-0/0/0.1              0
```

show multicast scope inet

```
user@host> show multicast scope inet
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
232-net              232.232.0.0/16     fe-0/0/0.1              0
local                239.255.0.0/16     fe-0/0/0.1              0
```

show multicast scope inet6

```
user@host> show multicast scope inet6
                                                          Resolve
Scope name           Group Prefix       Interface         Rejects
local                ff05::/16          fe-0/0/0.1              0
larry                ff05::1234/128     fe-0/0/0.1              0
```


Copyright © 2016, Juniper Networks, Inc.

show multicast sessions

List of Syntax
Syntax on page 4737
Syntax (EX Series Switch and the QFX Series) on page 4737

Syntax
show multicast sessions
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (EX Series Switch and the QFX Series)
show multicast sessions
<brief | detail | extensive>
<*regular-expression*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display information about announced IP multicast sessions.

Options
none—Display standard information about all multicast sessions for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Required Privilege Level
view

List of Sample Output
show multicast sessions on page 4738
show multicast sessions regular-expression detail on page 4738

Output Fields
Table 417 describes the output fields for the **show multicast sessions** command. Output fields are listed in the approximate order in which they appear.

Table 417: show multicast sessions Output Fields

Field Name	Field Description
session-name	Name of the known announced multicast sessions.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
  ...
  Monterey Bay - DockCam
  Monterey Bay - JettyCam
  Monterey Bay - StandCam
  Monterey DockCam
  Monterey DockCam / ROV cam
  ...
  NASA TV (MPEG-1)
  ...
  UO Broadcast - NASA Videos - 25 Years of Progress
  UO Broadcast - NASA Videos - Journey through the Solar System
  UO Broadcast - NASA Videos - Life in the Universe
  UO Broadcast - NASA Videos - Nasa and the Airplane
  UO Broadcasts OPB's Oregon Story
  UO DOD News Clips
  UO Medical Management of Biological Casualties (1)
  UO Medical Management of Biological Casualties (2)
  UO Medical Management of Biological Casualties (3)
...
  376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
 Session: NASA TV (MPEG-1)
 Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
 Email: Your Name Here <multicast@lists.private.edu>
 Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
 Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
```


Copyright © 2016, Juniper Networks, Inc.

```
Attribute: rtpmap:104 L16/22050

  1 matching sessions.
```

Copyright © 2016, Juniper Networks, Inc.

show multicast snooping route

Syntax

show multicast snooping route
<brief | detail | extensive>
<active | all | inactive>
<bridge-domain *bridge-domain-name*>
<group *group*>
<instance *instance-name*>
<mesh-group *mesh-group-name*>
<*regular-expression*>
<source-prefix *source-prefix*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display the entries in the IP multicast snooping forwarding table. You can display some of this information with the **show route table inet.1** command.

Options

none—Display standard information about all entries in the multicast snooping table for all virtual switches and all bridge domains.

brief | detail | extensive—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast snooping table.

bridge-domain *bridge-domain*—(Optional) Display the entries for a particular bridge domain.

group *group*—(Optional) Display the entries for a particular group.

instance *instance-name*—(Optional) Display the entries for a multicast instance.

mesh-group *mesh-group-name*—(Optional) Display the entries for a particular mesh group.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.

source-prefix *source-prefix*—(Optional) Display the entries for a particular source prefix.

Required Privilege Level

view

List of Sample Output

show multicast snooping route bridge-domain on page 4741
show multicast snooping route instance vs on page 4741

Output Fields

Table 418 describes the output fields for the **show multicast snooping route** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 418: show multicast snooping route Output Fields

Field Name	Field Description	Level of Output
Nexthop Bulking	Displays whether next-hop bulk updating is **ON** or **OFF** (only for routing-instance type of **virtual switch** or **vpls**).	All levels
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Routing-instance	Name of the routing instance to which this routing information applies. (Displayed when multicast is configured within a routing instance.)	All levels
Learning Domain	Name of the learning domain to which this routing information applies.	**detail extensive**
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Route state	Whether the group is **Active** or **Inactive**.	**extensive**
Forwarding state	Whether the prefix is **Pruned** or **Forwarding**.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry.	**extensive**

Sample Output

show multicast snooping route bridge-domain

```
user@host> show multicast snooping route bridge-domain br-dom-1 extensive
Family: INET

Group: 232.1.1.1
    Source: 192.168.3.100/32
    Downstream interface list:
        ge-0/1/0.200
    Statistics: 0 kBps, 0 pps, 1 packets
    Next-hop ID: 1048577
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 240 seconds
```

show multicast snooping route instance vs

```
user@host> show multicast snooping route instance vs
Nexthop Bulking: ON

Family: INET
```

Copyright © 2016, Juniper Networks, Inc.

```
Group: 224.0.0.0
    Bridge-domain: vsid500

Group: 225.1.0.1
    Bridge-domain: vsid500
    Downstream interface list: vsid500
       ge-0/3/8.500 ge-1/1/9.500 ge1/2/5.500
```


Copyright © 2016, Juniper Networks, Inc.

show multicast snooping statistics

Syntax
show multicast snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced in Junos OS Release 8.5.

Description
Display IP multicast snooping statistics.

Options
none—Display multicast snooping statistics for all supported address families for all routing instances.

instance *instance-name*—(Optional) Display statistics for a specific routing instance.

interface *interface-name*—(Optional) Display statistics for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
The input and output interface multicast snooping statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds.

Required Privilege Level
view

Related Documentation
- clear multicast snooping statistics on page 4644

List of Sample Output
show multicast snooping statistics on page 4745

Output Fields
Table 419 describes the output fields for the **show multicast snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 419: show multicast snooping statistics Output Fields

Field Name	Field Description
Routing-instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Family	Protocol family for which multicast statistics are displayed: **INET** or **INET6**.
Interface	Name of the interface for which statistics are being reported.
Routing Protocol	Primary multicast protocol on the interface: **PIM**, **DVMRP for INET**, or **PIM for INET6**.
Mismatch	Number of multicast packets that did not arrive on the correct upstream interface.

Copyright © 2016, Juniper Networks, Inc.

Table 419: show multicast snooping statistics Output Fields *(continued)*

Field Name	Field Description
Kernel Resolve	Number of resolve requests processed by the primary multicast protocol on the interface.
Resolve No Route	Number of resolve requests that were ignored because there was no route to the source.
In Kbytes	Total accumulated incoming packets (in KB) since the last time the **clear multicast snooping statistics** command was issued.
Out Kbytes	Total accumulated outgoing packets (in KB) since the last time the **clear multicast snooping statistics** command was issued.
Mismatch error	Number of mismatches that were ignored because of internal errors.
Mismatch No Route	Number of mismatches that were ignored because there was no route to the source.
Routing Notify	Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol.
Resolve Error	Number of resolve requests that were ignored because of internal errors.
In packets	Total number of incoming packets since the last time the **clear multicast snooping statistics** command was issued.
Out packets	Total number of outgoing packets since the last time the **clear multicast snooping statistics** command was issued.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show multicast snooping statistics

```
user@host> show multicast snooping statistics
Routing-instance: foo
 Family: INET
 Interface: fe-0/0/2.200
     Routing protocol:          PIM   Mismatch error:              0
     Mismatch:                    0   Mismatch no route:           0
     Kernel resolve:             22   Routing notify:              0
     Resolve no route:            0   Resolve error:               0
     Resolve filtered:            0   Notify filtered:             0
     In kbytes:                   0   In packets:                  0
     Out kbytes:                  0   Out packets:                 0

Routing-instance: bar
 Family: INET
 Interface: fe-0/1/2.200
     Routing protocol:          PIM   Mismatch error:              0
     Mismatch:                    0   Mismatch no route:           0
     Kernel resolve:             22   Routing notify:              0
     Resolve no route:            0   Resolve error:               0
     Resolve filtered:            0   Notify filtered:             0
     In kbytes:                   0   In packets:                  0
     Out kbytes:                  0   Out packets:                 0
```

Copyright © 2016, Juniper Networks, Inc.

show multicast usage

List of Syntax
Syntax on page 4746
Syntax (EX Series Switch and the QFX Series) on page 4746

Syntax
show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.

Options
none—Display multicast usage information for all supported address families for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast usage on page 4747
show multicast usage brief on page 4747
show multicast usage instance on page 4747
show multicast usage detail on page 4748

Output Fields
Table 420 describes the output fields for the **show multicast usage** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 420: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Group	Group address.
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays **unavailable**.
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays **unavailable**.
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.

Sample Output

show multicast usage

```
user@host> show multicast usage
Group           Sources Packets             Bytes
228.0.0.0       1       52847               4439148
239.1.1.1       2       13450               1125530


Prefix          /len Groups Packets             Bytes
10.255.14.144   /32  2      66254               5561304
10.255.70.15    /32  1      43                  3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see show multicast usage on page 4747.

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group           Sources Packets             Bytes
224.2.127.254   1       5538                509496
224.0.1.39      1       13                  624
224.0.1.40      1       13                  624

Prefix          /len Groups Packets             Bytes
192.168.195.34  /32  1      5538                509496
10.255.14.30    /32  1      13                  624
```

Copyright © 2016, Juniper Networks, Inc.

```
10.255.245.91   /32  1      13                  624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group           Sources Packets              Bytes
228.0.0.0       1       53159                4465356
    Source: 10.255.14.144   /32 Packets: 53159 Bytes: 4465356
239.1.1.1       2       13450                1125530
    Source: 10.255.14.144   /32 Packets: 13407 Bytes: 1122156
    Source: 10.255.70.15    /32 Packets: 43 Bytes: 3374


Prefix          /len Groups Packets              Bytes
10.255.14.144   /32  2      66566               5587512
    Group: 228.0.0.0       Packets: 53159 Bytes: 4465356
    Group: 239.1.1.1       Packets: 13407 Bytes: 1122156
10.255.70.15    /32  1      43                  3374
    Group: 239.1.1.1       Packets: 43 Bytes: 3374
```


Copyright © 2016, Juniper Networks, Inc.

show pgm negative-acknowledgments

Syntax show pgm negative-acknowledgments

Release Information Command introduced before Junos OS Release 7.4.

Description Display the sent or received Pragmatic General Multicast (PGM) negative acknowledgments (NAKs), the source-path message (SPM) sequence number being negatively acknowledged, and the current state of repair.

Options This command has no options.

Required Privilege Level view

List of Sample Output show pgm negative-acknowledgments on page 4750

Output Fields Table 421 describes the output fields for the **show pgm negative-acknowledgments** command. Output fields are listed in the approximate order in which they appear.

Table 421: show pgm negative-acknowledgments Output Fields

Field Name	Field Description
Global source id	Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI).
Network layer address	Network layer address of the local system.
Source port	Source port number, which is combined with the GSI to determine the TSI.
SPM sequence number	Numeric sequence identifier of the source-path message.
Window (trailing/leading sequence)	Range of sequence numbers used by the source for sequentially numbering and transmitting the most recent packets. The trailing (or left) edge of the transmit window is the sequence number of the oldest data packet available for repair from a source. The leading (or right) edge of the transmit window is defined as the sequence number of the most recent data packet a source has transmitted.
Outstanding NAKS	Total number of outstanding negative acknowledgments sent or received by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing. • **Sequence number**—Numeric sequence identifier of the source-path message. • **Group**—Group address. • **Source**—Multicast source. • **Interface**—Interface name. • **Receiver**—IP address receiving the multicast.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show pgm negative-acknowledgments

```
user@host> show pgm negative-acknowledgments
Global source ID: 010203040506 Source port: 1111
      Network layer address: 10.38.0.1
      SPM sequence number: 1
      Window (trailing/leading sequence): 0/1
      Outstanding NAKs:
              Sequence number: 1
              Group: 225.1.1.1
              Source: 192.168.195.121
              Interface: t3-0/2/0:0 Receiver: 10.38.0.10
```


Copyright © 2016, Juniper Networks, Inc.

show pgm source-path-messages

Syntax show pgm source-path-messages

Release Information Command introduced before Junos OS Release 7.4.

Description Display the Pragmatic General Multicast (PGM) source-path messages received.

Options This command has no options.

Required Privilege Level view

List of Sample Output show pgm source-path-messages on page 4751

Output Fields Table 422 describes the output fields for the **show pgm source-path-messages** command. Output fields are listed in the approximate order in which they appear.

Table 422: show pgm source-path-messages Output Fields

Field Name	Field Description
Global source ID	Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI).
Port	Source port number, which combines with the GSI to determine the TSI.
SPM number	Numeric sequence identifier of the source-path message.
Trail number	Sequence number of the oldest data packet available for repair from a source.
Lead number	Sequence number of the most recent data packet a source has transmitted.
Network layer address	Network layer address of the local system.

Sample Output

show pgm source-path-messages

```
user@host> show pgm source-path-messages
Global source ID  Port SPM number Trail number Lead number Network layer address
010203040506      1111          1            0           1 10.38.0.1
```

Copyright © 2016, Juniper Networks, Inc.

show pgm statistics

Syntax show pgm statistics

Release Information Command introduced before Junos OS Release 7.4.

Description Display Pragmatic General Multicast (PGM) packet statistics, including general loss and repair statistics.

Options This command has no options.

Required Privilege Level view

List of Sample Output show pgm statistics on page 4754

Output Fields Table 423 describes the output fields for the **show pgm statistics** command. Output fields are listed in the approximate order in which they appear.

Table 423: show pgm statistics Output Fields

Field Name	Field Description
PGM type, # received, # sent	Number of packets received and sent for the following PGM packet types: • **SPM**—Number of total source path messages received and sent by the local system. Source path messages (SPMs) are sent by a source to establish the source path state in network elements and to provide the transmit-window state to receivers. • **POLL**—Total number of poll requests received and sent by the local system. • **POLR**—Total number of poll responses received and sent by the local system. • **ODATA**—Total number of original data packets received and sent by the local system. • **RDATA**—Total number of repair data packets received and sent by the local system. RDATA packets are generated in response to negative acknowledgments (NAKs), which indicate a missing packet from the original data sequence. • **NAK**—Total number of negative acknowledgments received and sent by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing. • **NULLNAK**—Total number of null negative acknowledgments received and sent by the local system. NULLNAKs are transmitted by a designated local repairer that receives NAKs redirected to it by either receivers or network elements to provide flow-control feedback to a source. • **NCF**—Total number of NAK confirmations received and sent by the local system. NAK confirmations are generated in response to NAK packets that are received. • **SPMR**—Total number of source path message requests (SPMRs) received and sent by the local system. SPMRs are used to solicit a source path message from a source in a nonimplosive way. The typical application is for late-joining receivers to solicit source path messages directly from a source in order to be able to send NAKs for missing packets, without having to wait for a regularly scheduled source path message from that source. • **OTHER**—Total number of other PGM packets received and sent by the local system.
packets shorter than minimum PGM header length	Total number of packets received with headers that are shorter than the minimum required PGM header length.

Copyright © 2016, Juniper Networks, Inc.

Table 423: show pgm statistics Output Fields *(continued)*

Field Name	Field Description
packets received with incorrect check sum	Total number of packets received with an incorrect checksum. The checksum field is the 1's complement of the 1's complement sum of the entire PGM packet, including the header.
packets received with zero check sum	Total number of packets received with a zero checksum. If the computed checksum is zero, it is transmitted as all ones. A value of zero in this field means that the transmitter generated no checksum.
packets received with TSDU length incorrect	Total number of packets received with an incorrect Transport Service Data Unit (TSDU) length (16 bits).
packets received with SPM length incorrect	Total number of packets received with an incorrect source path message length.
packets received with unknown SPM address family	Total number of packets received with an unknown source path message address family indicator (AFI).
packets received with NAK length incorrect	Total number of packets received with an incorrect NAK length.
packets received with unknown NAK address family	Total number of packets received with an unknown NAK address family indicator (AFI).
packets received with NAK for unknown TSI	Total number of NAK packets received with an unknown transport session identifier (TSI).
packets received when NAK throttled	Total number of packets received when NAK is throttled.
packets received with NCF length incorrect	Total number of packets received with an incorrect NAK confirmation length.
packets received with unknown NCF address family	Total number of packets received with an unknown NAK confirmation address family indicator (AFI).
packets received with NCF for unknown TSI	Total number of NAK confirmation packets received with an unknown transport session identifier (TSI).
packets received with RDATA length incorrect	Total number of packets received with an incorrect RDATA length.
packets received with RDATA for unknown TSI	Total number of RDATA packets received with an unknown transport session identifier (TSI).

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show pgm statistics

```
user@host> show pgm statistics
PGM type     # received     # sent
SPM                   0          0
POLL                  0          0
POLR                  0          0
ODATA                 0          0
RDATA                 0          0
NAK                   0          0
NULLNAK               0          0
NCF                   0          0
SPMR                  0          0
OTHER                 0          0

packets shorter than minimum PGM header length  :          0
packets received with incorrect check sum       :          0
packets received with zero check sum            :          0
packets received with TSDU length incorrect     :          0
packets received with SPM length incorrect      :          0
packets received with unknown SPM address family:          0
packets received with NAK length incorrect      :          0
packets received with unknown NAK address family:          0
packets received with NAK for unknown TSI       :          0
packets received when NAK throttled             :          0
packets received with NCF length incorrect      :          0
packets received with unknown NCF address family:          0
packets received with NCF for unknown TSI       :          0
packets received with RDATA length incorrect    :          0
packets received with RDATA for unknown TSI     :          0
```


Copyright © 2016, Juniper Networks, Inc.

show pim bidirectional df-election

Syntax

show pim bidirectional df-election
<brief | detail >
<inet | inet6>
<instance *instance name*>
<logical-system (all | *logical-system-name*)>
<rpa *address*>

Release Information

Command introduced in Junos OS Release 12.1.

Description

For bidirectional PIM, display the designated forwarder (DF) election results for each interface grouped by the rendezvous point addresses (RPAs).

Options

none—Display standard information about all interfaces.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display DF election results for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display DF election results for a specific routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rpa *address*—(Optional) Display the DF election results for an RP address.

Required Privilege Level

view

List of Sample Output

show pim bidirectional df-election on page 4756
show pim bidirectional df-election brief on page 4756

Output Fields

Table 424 describes the output fields for the **show pim bidirectional df-election** command. Output fields are listed in the approximate order in which they appear.

Table 424: show pim bidirectional df-election Output Fields

Field Name	Field Description	Level of Output
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Instance	Name of the routing instance.	All levels
RPA	RP address.	All levels
Group ranges	Address ranges of the multicast groups mapped to this RP address.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 424: show pim bidirectional df-election Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interfaces	Bidirectional PIM interfaces on this routing device. An interface can win the DF election (**Win**), lose the DF election (**Lose**), or be the RP link (**RPL**). The RP link is the interface directly connected to a subnet that contains a phantom RP address. A phantom RP address is an RP address that is not assigned to a routing device interface.	All levels **brief** displays the DF election winner only.
DF	IP address of the designated forwarder.	All levels

Sample Output

show pim bidirectional df-election

```
user@host> show pim bidirectional df-election
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
    ge-0/0/1.0     (RPL)      DF: none
    lo0.0          (Win)      DF: 10.255.179.246
    xe-4/1/0.0     (Win)      DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
    ge-0/0/1.0     (Lose)     DF: 10.10.1.2
    lo0.0          (Win)      DF: 10.255.179.246
    xe-4/1/0.0     (Lose)     DF: 10.10.2.2

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
    ge-0/0/1.0     (Lose)     DF: fe80::b2c6:9aff:fe95:86fa
    lo0.0          (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0     (Win)      DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
    ge-0/0/1.0     (Lose)     DF: fe80::b2c6:9aff:fe95:86fa
    lo0.0          (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0     (Win)      DF: fe80::226:88ff:fec5:3c37
```

show pim bidirectional df-election brief

```
user@host> show pim bidirectional df-election brief
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
    lo0.0          (Win)      DF: 10.255.179.246
    xe-4/1/0.0     (Win)      DF: 10.10.2.1
```


Copyright © 2016, Juniper Networks, Inc.

```
RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
    lo0.0         (Win)      DF: 10.255.179.246

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
    lo0.0         (Win)      DF: fe80::2a0:a50f:fc64:e661
    xe-4/1/0.0    (Win)      DF: fe80::226:88ff:fec5:3c37
```

Copyright © 2016, Juniper Networks, Inc.

show pim bidirectional df-election interface

Syntax

show pim bidirectional df-election interface
<inet | inet6>
<instance *instance name*>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 12.1.

Description

For bidirectional PIM, display the default and the configured designated forwarder (DF) election parameters for each interface.

Options

none—Display standard information about all interfaces.

inet | inet6—(Optional) Display DF election parameters for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display DF election parameters for a specific routing instance.

interface-name—(Optional) Display DF election parameters for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim bidirectional df-election interface on page 4759

Output Fields

Table 425 describes the output fields for the **show pim bidirectional df-election interface** command. Output fields are listed in the approximate order in which they appear.

Table 425: show pim bidirectional df-election interface Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).
Interface	Name of the bidirectional PIM interface.
Robustnes Count	Minimum number of DF election messages that must fail to be received for DF election to fail.
Offer Period	Interval between repeated DF election messages.
Backoff Period	Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.

Copyright © 2016, Juniper Networks, Inc.

Table 425: show pim bidirectional df-election interface Output Fields *(continued)*

Field Name	Field Description
RPA	RP address.
State	For each RP address, state of each interface with respect to the DF election: **Offer** (when the election is in progress), **Win**, or **Lose**.
DF	IP address of the designated forwarder.

Sample Output

show pim bidirectional df-election interface

```
user@host> show pim bidirectional df-election interface
Instance: PIM.master Family: INET

Interface: ge-0/0/1.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                           State  DF
  10.10.1.3                     Offer  none
  10.10.13.2                    Lose   10.10.1.2

Interface: lo0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                           State  DF
  10.10.1.3                     Win    10.255.179.246
  10.10.13.2                    Win    10.255.179.246

Interface: xe-4/1/0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                           State  DF
  10.10.1.3                     Win    10.10.2.1
  10.10.13.2                    Lose   10.10.2.2

Instance: PIM.master Family: INET6

Interface: ge-0/0/1.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                           State  DF
  fec0::10:10:1:3               Lose   fe80::b2c6:9aff:fe95:86fa
  fec0::10:10:13:2              Lose   fe80::b2c6:9aff:fe95:86fa

Interface: lo0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                             State  DF
  fec0::10:10:1:3                 Win    fe80::2a0:a50f:fc64:e661
  fec0::10:10:13:2                Win    fe80::2a0:a50f:fc64:e661

Interface: xe-4/1/0.0
  Robustness Count: 3
  Offer Period: 100 ms
  Backoff Period: 1000 ms

  RPA                             State  DF
  fec0::10:10:1:3                 Win    fe80::226:88ff:fec5:3c37
  fec0::10:10:13:2                Win    fe80::226:88ff:fec5:3c37
```


Copyright © 2016, Juniper Networks, Inc.

show pim bootstrap

List of Syntax
Syntax on page 4761
Syntax (EX Series Switch and the QFX Series) on page 4761

Syntax
show pim bootstrap
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim bootstrap
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
instance option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description
For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.

Options
none—Display PIM bootstrap router information for all routing instances.

instance *instance-name*—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim bootstrap on page 4762
show pim bootstrap instance on page 4762

Output Fields
Table 426 describes the output fields for the **show pim bootstrap** command. Output fields are listed in the approximate order in which they appear.

Table 426: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.

Copyright © 2016, Juniper Networks, Inc.

Table 426: show pim bootstrap Output Fields *(continued)*

Field Name	Field Description
State	Local routing device election state: **Candidate**, **Elected**, or **Ineligible**.
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master

BSR                    Pri Local address           Pri State      Timeout
None                     0 10.255.71.46              0 InEligible       0
feco:1:1:1:1:0:aff:785c 34 feco:1:1:1:1:0:aff:7c12   0 InEligible       0
```

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR             Pri Local address   Pri State      Timeout
None              0 192.168.196.105   0 InEligible       0
```


Copyright © 2016, Juniper Networks, Inc.

show pim interfaces

List of Syntax
Syntax on page 4763
Syntax (EX Series Switch and the QFX Series) on page 4763

Syntax
show pim interfaces
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim interfaces
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Commmand introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description
Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

Options
none—Display interface information for all family addresses for all routing instances.

inet | inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about interfaces for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim interfaces on page 4764

Output Fields
Table 427 describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 427: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the **show interfaces** command.

Copyright © 2016, Juniper Networks, Inc.

Table 427: show pim interfaces Output Fields *(continued)*

Field Name	Field Description
Mode	PIM mode running on the interface: • **B**—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • **S**—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • **Dense**—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • **Sparse-Dense**—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as **dense** is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as **sparse** is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.) When sparse-dense mode is configured, the output includes both **S** and **D**. When bidirectional-sparse mode is configured, the output includes **S** and **B**. When bidirectional-sparse-dense mode is configured, the output includes **B**, **S**, and **D**.
IP	Version number of the address family on the interface: **4** (IPv4) or **6** (IPv6).
V	PIM version running on the interface: **1** or **2**.
State	State of PIM on the interface: • **Active**—Bidirectional mode is enabled on the interface and on all PIM neighbors. • **DR**—Designated router. • **NotCap**—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • **NotDR**—Not the designated router. • **P2P**—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*,g) join messages that have been seen on the interface.
DR address	Address of the designated router.

Sample Output

show pim interfaces

```
user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
```


Copyright © 2016, Juniper Networks, Inc.

```
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/3/0.0         Up   S     4 2 NotDR,NotCap      1 0/0            40.0.0.3
ge-0/3/3.50        Up   S     4 2 DR,NotCap         1 9901/100       50.0.0.2
ge-0/3/3.51        Up   S     4 2 DR,NotCap         1 0/0            51.0.0.2
pe-1/2/0.32769     Up   S     4 2 P2P,NotCap        0 0/0
```

Copyright © 2016, Juniper Networks, Inc.

show pim join

List of Syntax

Syntax on page 4766
Syntax (EX Series Switch and the QFX Series) on page 4766

Syntax

show pim join
<brief | detail | extensive | summary>
<bidirectional | dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Syntax (EX Series Switch and the QFX Series)

show pim join
<brief | detail | extensive | summary>
<dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
summary option introduced in Junos OS Release 9.6.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Multiple new filter options introduced in Junos OS Release 13.2.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (*,G-range) for each active bidirectional RP group range, in addition to each of the joined (*,G) routes.

Options

none—Display the standard information about PIM groups for all supported family addresses for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

bidirectional | dense | sparse—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Display information about only the group that exactly matches the specified group address.

Copyright © 2016, Juniper Networks, Inc.

inet | inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as ***prefix/prefix-length***.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Display information about PIM (S,G) or (*,G) entries.

Required Privilege Level

view

Related Documentation

- clear pim join on page 4648
- Example: Configuring Bidirectional PIM on page 4205

List of Sample Output

show pim join summary on page 4771
show pim join (PIM Sparse Mode) on page 4771
show pim join (Bidirectional PIM) on page 4772
show pim join inet6 on page 4772
show pim join inet6 star-g on page 4773
show pim join instance <instance-name> on page 4773
show pim join detail on page 4773
show pim join extensive (PIM Sparse Mode) on page 4774
show pim join extensive (Bidirectional PIM) on page 4775
show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP) on page 4776
show pim join instance <instance-name> extensive on page 4776
show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4777
show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4778
show pim join summary on page 4779
show pim join (PIM Sparse Mode) on page 4780
show pim join (Bidirectional PIM) on page 4780
show pim join inet6 on page 4781
show pim join inet6 star-g on page 4781
show pim join instance <instance-name> on page 4781
show pim join detail on page 4782
show pim join extensive (PIM Sparse Mode) on page 4782
show pim join extensive (Bidirectional PIM) on page 4783

Copyright © 2016, Juniper Networks, Inc.

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP) on page 4784
show pim join instance <instance-name> extensive on page 4784
show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4785
show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute) on page 4786

Output Fields Table 428 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 428: show pim join Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	**brief detail extensive summary** none
Family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	**brief detail extensive summary** none
Route type	Type of multicast route: (S,G) or (*,G).	**summary**
Route count	Number of (S,G) routes and number of (*,G) routes.	**summary**
R	Rendezvous Point Tree.	**brief detail extensive** none
S	Sparse.	**brief detail extensive** none
W	Wildcard.	**brief detail extensive** none
Group	Group address.	**brief detail extensive** none
Bidirectional group prefix length	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
Source	Multicast source: • ***** (wildcard value) • ***ipv4-address*** • ***ipv6-address***	**brief detail extensive** none
RP	Rendezvous point for the PIM group.	**brief detail extensive** none
Flags	PIM flags: • **bidirectional**—Bidirectional mode entry. • **dense**—Dense mode entry. • **rptree**—Entry is on the rendezvous point tree. • **sparse**—Sparse mode entry. • **spt**—Entry is on the shortest-path tree for the source. • **wildcard**—Entry is on the shared tree.	**brief detail extensive** none

Copyright © 2016, Juniper Networks, Inc.

Table 428: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Upstream interface	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, **RP Link** means that the interface is directly connected to a subnet that contains a phantom RP address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**brief detail extensive** none
Upstream neighbor	Information about the upstream neighbor: **Direct**, **Local**, **Unknown**, or a specific IP address. For bidirectional PIM, **Direct** means that the interface is directly connected to a subnet that contains a phantom RP address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**extensive**
Upstream state	When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.	**extensive**
Active upstream neighbor	On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.	**extensive**
MoFRR Backup upstream interface	The MoFRR upstream interface that is used when the primary path fails. When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 428: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Upstream state	Information about the upstream interface: • **Join to RP**—Sending a join to the rendezvous point. • **Join to Source**—Sending a join to the source. • **Local RP**—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • **Local Source**—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • **No Prune to RP**—Automatically sent to RP when SPT and RPT are on the same path. • **Prune to RP**—Sending a prune to the rendezvous point. • **Prune to Source**—Sending a prune to the source. NOTE: RP group range entries have **None** in the **Upstream state** field because RP group ranges do not trigger actual PIM join messages between routing devices.	**extensive**
Downstream neighbors	Information about downstream interfaces: • **Interface**—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • **Interface address**—Address of the downstream neighbor. • **State**—Information about the downstream neighbor: **join** or **prune**. • **Flags**—PIM join flags: **R (RPtree)**, **S (Sparse)**, **W (Wildcard)**, or **zero**. • **Uptime**—Time since the downstream interface joined the group. • **Time since last Join**—Time since the last join message was received from the downstream interface. • **Time since last Prune**—Time since the last prune message was received from the downstream interface.	**extensive**
Assert Timeout		
Assert Timeout	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 428: show pim join Output Fields *(continued)*

Field Name	Field Description	Level of Output
Keepalive timeout	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, **Keepalive timeout** is **Infinity**.	**extensive**
Uptime	Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.	**extensive**
Bidirectional accepting interfaces	Interfaces on the router that forward bidirectional PIM traffic. The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (**DF Winner**), or the interface is the reverse path forwarding (RPF) interface toward the RP (**RPF**).	**extensive**

Sample Output

show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type            Route count
(s,g)                 2
(*,g)                 1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local

 Group: ff04::e000:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: ff04::e800:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```


Copyright © 2016, Juniper Networks, Inc.

```
 Group: ff04::e800:101
     Source: ::1.1.1.2
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
```

Copyright © 2016, Juniper Networks, Inc.

```
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 2

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source, Local RP
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: S Timeout: 174
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 2

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0
    Upstream neighbor: 10.111.10.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: Pseudo-GMP
            fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
        Interface: so-1/0/0.0 (pruned)
            10.111.10.2 State: Prune Flags: SR Timeout: 174
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
```


Copyright © 2016, Juniper Networks, Inc.

```
                10.10.47.100 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:03:49 Time since last Prune: 00:01:49
        Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
    Number of downstream interfaces: 0

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
     Downstream neighbors:
         Interface: lt-1/0/10.24
             10.0.24.4 State: Join   RW  Timeout: 185
         Interface: lt-1/0/10.23
             10.0.23.3 State: Join   RW  Timeout: 184
     Number of downstream interfaces: 2

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 1

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52
```


Copyright © 2016, Juniper Networks, Inc.

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:55
    Downstream neighbors:
        Interface: Pseudo-MLDP
        Interface: lt-1/2/0.25
            1.2.5.2 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.1.1.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.2.2.2
    Source: 1.2.7.7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:25
    Downstream neighbors:
        Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
    Source: abcd::1:2:7:7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
```

Copyright © 2016, Juniper Networks, Inc.

```
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:26
    Downstream neighbors:
        Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
    Source: *
    RP: 1.1.1.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 11:31:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: lt-1/2/0.14
            1.1.4.4 State: Join Flags: S Timeout: 177
            Uptime: 11:30:33 Time since last Join: 00:00:33
    Downstream neighbors:
```


Copyright © 2016, Juniper Networks, Inc.

```
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2
    Source: 1.2.7.7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:30
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
    Source: abcd::1:2:7:7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <1.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32
```

Sample Output

show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type           Route count
(s,g)                2
(*,g)                1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```


Copyright © 2016, Juniper Networks, Inc.

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local

 Group: ff04::e000:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: ff04::e800:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: ff04::e800:101
     Source: ::1.1.1.2
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: ff04::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
```

Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 2

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source, Local RP
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
                10.111.10.2 State: Join Flags: S Timeout: 174
                Uptime: 00:03:49 Time since last Prune: 00:01:49
            Interface: mt-1/1/0.32768
                10.10.47.100 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:03:49 Time since last Prune: 00:01:49
        Number of downstream interfaces: 2

    Group: 239.1.1.1
        Source: 10.255.70.15
        Flags: sparse,spt
        Upstream interface: so-1/0/0.0
        Upstream neighbor: 10.111.10.2
        Upstream state: Local RP, Join to Source
        Keepalive timeout: 344
        Uptime: 00:03:49
        Downstream neighbors:
            Interface: Pseudo-GMP
                fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
            Interface: so-1/0/0.0 (pruned)
                10.111.10.2 State: Prune Flags: SR Timeout: 174
                Uptime: 00:03:49 Time since last Prune: 00:01:49
            Interface: mt-1/1/0.32768
                10.10.47.100 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:03:49 Time since last Prune: 00:01:49
        Number of downstream interfaces: 3

    Instance: PIM.master Family: INET6
    R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
    Number of downstream interfaces: 0

Group: 225.1.1.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
```

Copyright © 2016, Juniper Networks, Inc.

```
        Downstream neighbors:
            Interface: lt-1/0/10.24
                10.0.24.4 State: Join   RW  Timeout: 185
            Interface: lt-1/0/10.23
                10.0.23.3 State: Join   RW  Timeout: 184
        Number of downstream interfaces: 2

    Group: 225.1.3.0
        Bidirectional group prefix length: 24
        Source: *
        RP: 10.10.1.3
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0 (RP Link)
        Upstream neighbor: Direct
        Upstream state: Local RP
        Uptime: 00:03:49
        Bidirectional accepting interfaces:
            Interface: ge-0/0/1.0    (RPF)
            Interface: lo0.0         (DF Winner)
            Interface: xe-4/1/0.0    (DF Winner)
        Number of downstream interfaces: 0

    Instance: PIM.master Family: INET6
    R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
```


Copyright © 2016, Juniper Networks, Inc.

```
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 1

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:55
    Downstream neighbors:
        Interface: Pseudo-MLDP
        Interface: lt-1/2/0.25
            1.2.5.2 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.1.1.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
```

Copyright © 2016, Juniper Networks, Inc.

```
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 232.2.2.2
    Source: 1.2.7.7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:25
    Downstream neighbors:
        Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
    Source: abcd::1:2:7:7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:26
    Downstream neighbors:
        Interface: Pseudo-MLDP
```

show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
    Source: 10.0.0.1
    Flags: sparse,spt
    Active upstream interface: fe-1/2/13.0
    Active upstream neighbor: 10.0.0.9
    MoFRR Backup upstream interface: fe-1/2/14.0
    MoFRR Backup upstream neighbor: 10.0.0.21
    Upstream state: Join to Source, No Prune to RP
    Keepalive timeout: 354
    Uptime: 00:00:06
    Downstream neighbors:
        Interface: fe-1/2/15.0
            10.0.0.13 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:00:06 Time since last Join: 00:00:06
    Number of downstream interfaces: 1
```


Copyright © 2016, Juniper Networks, Inc.

show pim neighbors

List of Syntax
Syntax on page 4787
Syntax (EX Series Switch and the QFX Series) on page 4787

Syntax
show pim neighbors
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)
show pim neighbors
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Support for the **instance all** option added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display information about Protocol Independent Multicast (PIM) neighbors.

Options
none—(Same as **brief**) Display standard information about PIM neighbors for all supported family addresses for the main instance.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.

instance (*instance-name* | all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim neighbors on page 4789
show pim neighbors brief on page 4789
show pim neighbors instance on page 4789
show pim neighbors detail on page 4789
show pim neighbors detail (With BFD) on page 4790

Output Fields
Table 429 describes the output fields for the **show pim neighbors** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 429: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
Neighbor addr	Address of the neighboring PIM routing device.	All levels
IP	IP version: **4** or **6**.	All levels
V	PIM version running on the neighbor: **1** or **2**.	All levels
Mode	PIM mode of the neighbor: **Sparse**, **Dense**, **SparseDense**, or **Unknown**. When the neighbor is running PIM version 2, this mode is always **Unknown**.	All levels
Option	Can be one or more of the following: • **B**—Bidirectional Capable. • **G**—Generation Identifier. • **H**—Hello Option Holdtime. • **L**—Hello Option LAN Prune Delay. • **P**—Hello Option DR Priority. • **T**—Tracking bit.	**brief** none
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format ***dd*:*hh*:*mm*:*ss* ago** for less than a week and ***nwnd*:*hh*:*mm*:*ss* ago** for more than a week.	All levels
Address	Address of the neighboring PIM routing device.	**detail**
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: **Enabled**, **Operational state is up**, or **Disabled**.	**detail**
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is **0** through **65,535.**	**detail**
Hello Default Holdtime	Default holdtime and the time remaining if the **holdtime** option is not in the received hello message.	**detail**
Hello Option DR Priority	Designated router election priority. The range of values is **0** through **255**.	**detail**
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	**detail**
Hello Option Bi-Directional PIM supported	Neighbor can process bidirectional PIM messages.	**detail**
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format **delay *nnn* ms override *nnnn* ms**.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 429: show pim neighbors Output Fields *(continued)*

Field Name	Field Description	Level of Output
Join Suppression supported	Neighbor is capable of join suppression.	**detail**
Rx Join	Information about joins received from the neighbor. • **Group**—Group addresses in the join message. • **Source**—Address of the source in the join message. • **Timeout**—Time for which the join is valid.	**detail**

Sample Output

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface           IP V Mode        Option      Uptime Neighbor addr
so-1/0/0.0           4 2             HPLG      00:07:10 10.111.10.2
```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see show pim neighbors on page 4789.

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface           IP V Mode        Option      Uptime Neighbor addr
at-0/3/1.0           4 2             HPLG      00:07:54 10.111.30.2
mt-1/1/0.32768       4 2             HPLG      00:07:22 10.10.47.101
so-1/0/1.0           4 2             HPLG      00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

    Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsg
Join Count: 2
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 2053759302
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

Copyright © 2016, Juniper Networks, Inc.

```
    Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 93 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1734018161
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported

Interface: lo0.0

    Address: 10.255.179.246,    IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1997462267
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
    Address: 192.168.11.1,        IPv4, PIM v2, Mode: Sparse
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 836607909
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.11.2,        IPv4, PIM v2
        BFD: Enabled, Operational state is up
        Hello Default Holdtime: 105 seconds 104 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1907549685
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
    Address: 192.168.12.1,        IPv4, PIM v2
        BFD: Disabled
        Hello Default Holdtime: 105 seconds 80 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1971554705
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```


Copyright © 2016, Juniper Networks, Inc.

show pim rps

List of Syntax
Syntax on page 4791
Syntax (EX Series Switch and the QFX Series) on page 4791

Syntax

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).

Options

none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about RPs for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- Example: Configuring Bidirectional PIM on page 4199

List of Sample Output

show pim rps on page 4794
show pim rps brief on page 4794

Copyright © 2016, Juniper Networks, Inc.

show pim rps <group-address> on page 4794
show pim rps <group-address> (Bidirectional PIM) on page 4794
show pim rps <group-address> (PIM Dense Mode) on page 4795
show pim rps <group-address> (SSM Range Without asm-override-ssm Configured) on page 4795
show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP) on page 4795
show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP) on page 4795
show pim rps instance on page 4795
show pim rps extensive (PIM Sparse Mode) on page 4795
show pim rps extensive (Bidirectional PIM) on page 4796
show pim rps extensive (PIM Anycast RP in Use) on page 4796

Output Fields Table 430 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 430: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Family or Address family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: • **auto-rp**—Address of the RP known through the Auto-RP protocol. • **bootstrap**—Address of the RP known through the bootstrap router protocol (BSR). • **embedded**—Address of the RP known through an embedded RP (IPv6). • **static**—Address of RP known through static configuration.	**brief** none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	**brief** none
Learned via	Address and method by which the RP was learned.	**detail extensive**
Mode	The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 430: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output
Time Active	How long the RP has been active, in the format ***hh:mm:ss***.	**detail extensive**
Device Index	Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the **Device Index** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Subunit	Logical unit number of the interface. For bidirectional RPs, the **Subunit** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the **Interface** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Group Ranges	Addresses of groups that this RP spans.	**detail extensive** *group-address*
Active groups using RP	Number of groups currently using this RP.	**detail extensive**
total	Total number of active groups for this RP.	**detail extensive**
Register State for RP	Current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • **First Hop**—PIM-designated routing device that sent the Register message (the source address in the IP header). • **RP Address**—RP to which the Register message was sent (the destination address in the IP header). • **State**: On the designated router: ◦ **Send**—Sending Register messages. ◦ **Probe**—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. ◦ **Suppress**—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to **Probe** state. • On the RP: ◦ **Receive**—Receiving Register messages.	**extensive**
Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 430: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	**extensive**

Anycast-PIM Register State	If anycast RP is configured, the current register state for each group:
	• **Group**—Multicast group address.
	• **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.
	• **Origin**—How the information was obtained:
	  • **DIRECT**—From a local attachment
	  • **MSDP**—From the Multicast Source Discovery Protocol (MSDP)
	  • **DR**—From the designated router
		extensive

RP selected	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	*group-address*

Sample Output

show pim rps

```
user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type        Mode   Holdtime Timeout Groups Group prefixes
10.10.1.3       static      bidir       150    None      2 224.1.3.0/24
                                                           225.1.3.0/24
10.10.13.2      static      bidir       150    None      2 224.1.1.0/24
                                                           225.1.1.0/24
```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see show pim rps on page 4794.

show pim rps <group-address>

```
user@host> show pim rps 235.100.100.0
Instance: PIM.master
Instance: PIM.master

RP selected: 100.100.100.100
```

show pim rps <group-address> (Bidirectional PIM)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
        11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75
```


Copyright © 2016, Juniper Networks, Inc.

show pim rps <group-address> (PIM Dense Mode)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
        11.4.12.75

RP selected: 11.4.12.75
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

```
user@host> show pim rps 224.1.1.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
        11.4.12.75 (Bidirectional)

RP selected: (null)
```

show pim rps instance

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address              Type        Holdtime Timeout Groups Group prefixes
10.10.47.100            static             0    None      1 224.0.0.0/4

Address family INET6
```

show pim rps extensive (PIM Sparse Mode)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
```

Copyright © 2016, Juniper Networks, Inc.

```
    224.0.0.0/4, 36s remaining
Active groups using RP:
    225.1.1.1

    total 1 groups active

Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
225.1.1.1       192.168.195.78  10.255.14.132   10.255.245.91   Receive        0
```

show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        224.1.3.0/24
        225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        224.1.1.0/24
        225.1.1.0/24
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
           10.100.111.34
           10.100.111.17
           10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:
```


Copyright © 2016, Juniper Networks, Inc.

```
Group                    Source                           Origin
224.1.1.1                10.10.95.2                       DIRECT
224.1.1.2                10.10.95.2                       DIRECT
224.10.10.10             10.10.70.1                       MSDP
224.10.10.11             10.10.70.1                       MSDP
224.20.20.1              10.10.71.1                       DR

Address family INET6

Anycast-PIM rpset:
           ab::1
           ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:
Group                   Source                          Origin
::224.1.1.1             ::10.10.95.2                    DIRECT
::224.1.1.2             ::10.10.95.2                    DIRECT
::224.20.20.1           ::10.10.71.1                    DR
```

show pim source

List of Syntax
Syntax on page 4798
Syntax (EX Series Switch and the QFX Series) on page 4798

Syntax

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<*source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.

Options

none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-prefix—(Optional) Display the state for source RPF states in the given range.

Required Privilege Level

view

List of Sample Output

show pim source on page 4799
show pim source brief on page 4799
show pim source detail on page 4799

Output Fields

Table 431 describes the output fields for the **show pim source** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 431: show pim source Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Source	Address of the source or reverse path.
Prefix/length	Prefix and prefix length for the route used to reach the RPF address.
Upstream interface	RPF interface toward the source address.
Upstream Neighbor	Address of the RPF neighbor used to reach the source address.

Sample Output

show pim source

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6
```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see show pim source on page 4799.

show pim source detail

```
user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local
    Active groups:228.0.0.0
        239.1.1.1
        239.1.1.1

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2
    Active groups:239.1.1.1
```

Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.master Family: INET6
```


Copyright © 2016, Juniper Networks, Inc.

show pim statistics

List of Syntax
Syntax on page 4801
Syntax (EX Series Switch and the QFX Series) on page 4801

Syntax

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.

Description

Display Protocol Independent Multicast (PIM) statistics.

Options

none—Display PIM statistics.

inet | inet6—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.

instance *instance-name*—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear pim statistics on page 4654

List of Sample Output

show pim statistics on page 4807
show pim statistics inet interface <interface-name> on page 4808
show pim statistics inet6 interface <interface-name> on page 4809
show pim statistics interface <interface-name> on page 4809

Output Fields

Table 432 describes the output fields for the **show pim statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance. This field only appears if you specify an interface, for example: • **inet interface *interface-name*** • **inet6 interface *interface-name*** • **interface *interface-name***
Family	Output is for IPv4 or IPv6 PIM statistics. **INET** indicates IPv4 statistics, and **INET6** indicates IPv6 statistics. This field only appears if you specify an interface, for example: • **inet interface *interface-name*** • **inet6 interface *interface-name*** • **interface *interface-name***
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields *(continued)*

Field Name	Field Description
V2 State Refresh	PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS.
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the router is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the router has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the router has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream router, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream router, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the router has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the router has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encap if	Number of PIM register packets received when the first-hop router does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the router has an RP mismatch.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields *(continued)*

Field Name	Field Description
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the router has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the **maximum-rps** statement is exceeded. The **maximum-rps** statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	Number of packets in which the embedded RP for IPv6 is added. The following receive events trigger extraction of an IPv6 embedded RP address on the router: • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.

Copyright © 2016, Juniper Networks, Inc.

Table 432: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.

Sample Output

show pim statistics

```
user@host> show pim statistics
PIM Message type        Received       Sent  Rx errors
V2 Hello                      15         32          0
V2 Register                    0        362          0
V2 Register Stop             483          0          0
V2 Join Prune                 18        518          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
Rx Unknown Reg Stop                           0
Rx Join/Prune no state                        0
```

Copyright © 2016, Juniper Networks, Inc.

```
Rx Join/Prune on upstream if                0
Rx Join/Prune for invalid group             5
Rx Join/Prune messages dropped              0
Rx sparse join for dense group              0
Rx Graft/Graft Ack no state                 0
Rx Graft on upstream if                     0
Rx CRP not BSR                              0
Rx BSR when BSR                             0
Rx BSR not RPF if                           0
Rx unknown hello opt                        0
Rx data no state                            0
Rx RP no state                              0
Rx aggregate                                0
Rx malformed packet                         0
Rx illegal TTL                              0
Rx illegal destination address              0
No RP                                       0
No register encap if                        0
No route upstream                           0
Nexthop Unusable                            0
RP mismatch                                 0
RP mode mismatch                            0
RPF neighbor unknown                        0
Rx Joins/Prunes filtered                    0
Tx Joins/Prunes filtered                    0
Embedded-RP invalid addr                    0
Embedded-RP limit exceed                    0
Embedded-RP added                           0
Embedded-RP removed                         0
Rx Register msgs filtering drop             0
Tx Register msgs filtering drop             0
Rx Bidir Join/Prune on non-Bidir if         0
Rx Bidir Join/Prune on non-DF if            0
```

Sample Output

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
```


Copyright © 2016, Juniper Networks, Inc.

```
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

Sample Output

show pim statistics inet6 interface <interface-name>

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

Sample Output

show pim statistics interface <interface-name>

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

Copyright © 2016, Juniper Networks, Inc.

```
Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```


Copyright © 2016, Juniper Networks, Inc.

show policy

List of Syntax
Syntax on page 4811
Syntax (EX Series Switches) on page 4811

Syntax
show policy
<logical-system (all | *logical-system-name*)>
<*policy-name*>

Syntax (EX Series Switches)
show policy
<*policy-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display information about configured routing policies.

Options
none—List the names of all configured routing policies.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

policy-name—(Optional) Show the contents of the specified policy.

Required Privilege Level
view

List of Sample Output
show policy on page 4811
show policy policy-name on page 4812
show policy (Multicast Scoping) on page 4812

Output Fields
Table 433 lists the output fields for the **show policy** command. Output fields are listed in the approximate order in which they appear.

Table 433: show policy Output Fields

Field Name	Field Description
policy-name	Name of the policy listed.
term	Policy term listed.
from	Match condition for the policy.
then	Action for the policy.

Sample Output

show policy

```
user@host> show policy
```

```
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
all_routes
```

show policy policy-name

```
user@host> show policy test-statics
Policy test-statics:
    from
         3.0.0.0/8  accept
         3.1.0.0/16  accept
    then reject
```

show policy (Multicast Scoping)

```
user@host> show policy test-statics
Policy test-statics:
    from
         multicast-scoping == 8
```


Copyright © 2016, Juniper Networks, Inc.

show route

List of Syntax
Syntax on page 4813
Syntax (EX Series Switches) on page 4813

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1 on MX Series routers for enhanced subscriber management.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

Required Privilege Level

view

Related Documentation

- *Example: Configuring RIP*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*

List of Sample Output

show route on page 4816
show route destination-prefix on page 4817
show route destination-prefix detail on page 4817

Copyright © 2016, Juniper Networks, Inc.

show route extensive on page 4818
show route (Enhanced Subscriber Management) on page 4818

Output Fields Table 434 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 434: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).
destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as: • ***MPLS-label*** (for example, 80001). • ***interface-name*** (for example, ge-1/0/2). • ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): ◦ ***neighbor-address***—Address of the neighbor. ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. ◦ ***vc-id***—Virtual circuit identifier. ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

Copyright © 2016, Juniper Networks, Inc.

Table 434: show route Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Copyright © 2016, Juniper Networks, Inc.

Table 434: show route Output Fields *(continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.

Sample Output

show route

```
user@host> show route
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
```


Copyright © 2016, Juniper Networks, Inc.

```
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0

green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
224.0.0.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.10.10.10/32     *[Direct/0] 01:08:46
                    > via lo0.1
10.255.245.212/32  *[BGP/170] 00:01:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.213/32  *[BGP/170] 00:40:47, localpref 100
                      AS path: 100 I
                    > to 100.1.1.1 via so-0/0/1.0
```

show route destination-prefix

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 5.5.5.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
5.5.5.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

Copyright © 2016, Juniper Networks, Inc.

show route extensive

```
user@host> show route extensive
inet.0: 335844 destinations, 335845 routes (335395 active, 0 holddown, 450 hidden)
1.9.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006553
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:58:10        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
                Indirect next hops: 1
                        Protocol next hop: 192.168.69.71
                        Indirect next hop: 8e166c0 342
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 192.168.167.254 via fxp0.0
                        192.168.0.0/16 Originating RIB: inet.0
                          Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 192.168.167.254 via fxp0.0
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

100.20.0.111/32    *[Access-internal/12] 00:00:08
                  > to #0 10.0.0.1.93.65 via demux0.1073741824
100.20.0.112/32    *[Access-internal/12] 00:00:08
                   Private unicast
.
.
.
```


Copyright © 2016, Juniper Networks, Inc.

show route table

List of Syntax
Syntax on page 4819
Syntax (EX Series Switches) on page 4819

Syntax
show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route entries in a particular routing table.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level
view

Related Documentation
- *show route summary*

List of Sample Output
show route table bgp.l2.vpn on page 4820
show route table bgp.l3vpn.0 on page 4820
show route table bgp.l3vpn.0 detail on page 4820
show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 4822
show route table bgp.evpn.0 on page 4822
show route table inet.0 on page 4822
show route table inet.3 on page 4823
show route table inet6.0 on page 4823
show route table inet6.3 on page 4823
show route table inetflow detail on page 4824
show route table l2circuit.0 on page 4824
show route table mpls on page 4824
show route table mpls extensive on page 4825
show route table mpls.0 on page 4825
show route table mpls.0 detail (PTX Series) on page 4825
show route table mpls.0 extensive (PTX Series) on page 4826
show route table mpls.0 (RSVP Route—Transit LSP) on page 4827

Copyright © 2016, Juniper Networks, Inc.

show route table vpls_1 detail on page 4827
show route table vpn-a on page 4827
show route table vpn-a.mdt.0 on page 4828
show route table VPN-A detail on page 4828
show route table VPN-AB.inet.0 on page 4829
show route table VPN_blue.mvpn-inet6.0 on page 4829
show route table VPN-A detail on page 4829
show route table inetflow detail on page 4830

Output Fields For information about output fields, see the output field tables for the show route command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182449
         Protocol next hop: 10.255.245.12
         Push 182449
         Indirect next hop: 863a630 297
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
```

 Copyright © 2016, Juniper Networks, Inc.

```
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

                Communities: 2914:420 target:11111:1 origin:56:78
                VPN Label: 182449
                Localpref: 100
                 Router ID: 10.255.245.12

        10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
          *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 863a8f0 305
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
        AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                       Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                       VPN Label: 182465
                       Localpref: 100
                       Router ID: 10.255.245.12

        10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
          *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

        10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
          *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
```

Copyright © 2016, Juniper Networks, Inc.

```
State: <Active Int Ext>
Local AS:    35 Peer AS:    35
Age: 12:19      Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.o
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table bgp.evpn.0

```
user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```


Copyright © 2016, Juniper Networks, Inc.

```
0.0.0.0/0          *[Static/5] 00:51:57
                    > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                      Local
12.12.12.21/32     *[Local/0] 00:51:57
                      Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                      Local
13.13.13.21/32     *[Local/0] 00:51:58
                      Local
13.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                      Local
```

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 1.2.94.2 via lt-1/2/0.49
                    > to 1.2.3.2 via lt-1/2/0.23
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

Copyright © 2016, Juniper Networks, Inc.

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 65002 Peer AS: 65000
                Age: 4
                Task: BGP_65000.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 65000 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 65002
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
```


Copyright © 2016, Juniper Networks, Inc.

```
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:45:09, metric 1
                      Receive
1                  *[MPLS/0] 00:45:09, metric 1
                      Receive
2                  *[MPLS/0] 00:45:09, metric 1
                      Receive
100000             *[L2VPN/7] 00:43:04
                    > via so-0/1/0.1, Pop
100001             *[L2VPN/7] 00:43:03
                    > via so-0/1/0.2, Pop       Offset: 4
100002             *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100002(S=0)        *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100003             *[LDP/9] 00:43:22, metric 1
                    > via so-0/1/2.0, Swap 100002
                      via so-0/1/3.0, Swap 100002
100004             *[LDP/9] 00:43:16, metric 1
                      via so-0/1/2.0, Swap 100049
                    > via so-0/1/3.0, Swap 100049
so-0/1/0.1         *[L2VPN/7] 00:43:04
                    > via so-0/1/2.0, Push 100001, Push 100049(top)
                      via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2         *[L2VPN/7] 00:43:03
                      via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                    > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4
```

show route table mpls.0 detail (PTX Series)

```
user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
```

Copyright © 2016, Juniper Networks, Inc.

```
                        Address: 0x9438f34
                        Next-hop reference count: 2
                        Next hop type: Router, Next hop index: 567
                        Next hop: 3.0.0.1 via ge-0/0/1.0, selected
                        Label operation: Push 299808
                        Label TTL action: prop-ttl
                        Load balance label: Label 299808:None;
                        Session Id: 0x1
                        Protocol next hop: 10.255.255.1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        State: <Active Int>
                        Age: 21         Metric2: 1
                        Validation State: unverified
                        Task: Common L2 VC
                        Announcement bits (2): 0-KRT 2-Common L2 VC
                        AS path: I
```

show route table mpls.0 extensive (PTX Series)

```
user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 3.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 47         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
                Composite next hops: 1
                        Protocol next hop: 10.255.255.1 Metric: 1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 3.0.0.1 via ge-0/0/1.0
                                Session Id: 0x1
                        10.255.255.1/32 Originating RIB: inet.3
```


Copyright © 2016, Juniper Networks, Inc.

```
                                   Metric: 1                       Node path count: 1
                                   Forwarding nexthops: 1
                                         Nexthop: 3.0.0.1 via ge-0/0/1.0
```

show route table mpls.0 (RSVP Route—Transit LSP)

In the sample output, the 1 in **[RSVP/7/1]** indicates the secondary preference value. The secondary preference value becomes significant when multiple RSVP LSPs of different types are signaled to the destination. The possible values of RSVP secondary preferences are:

1—Normal Point-to-Point RSVP-TE LSP

2—Point-to-Multipoint (P2MP) RSVP-TE LSP

3—Dynamic RSVP-TE LSP

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
13                 *[MPLS/0] 00:37:31, metric 1
                      Receive
300352             *[RSVP/7/1] 00:08:00, metric 1
                    > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)        *[RSVP/7/1] 00:08:00, metric 1
                    > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384             *[RSVP/7/2] 00:05:20, metric 1
                    > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)        *[RSVP/7/2] 00:05:20, metric 1
                    > to 8.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
```

Copyright © 2016, Juniper Networks, Inc.

```
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                   *[VPN/7] 05:48:27
                     Discard
192.168.24.1:1:2:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                   *[MVPN/70] 01:23:05, metric2 1
                      Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                      AS path: I
                    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
                      AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```


Copyright © 2016, Juniper Networks, Inc.

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                             AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65535:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
6:10.255.2.203:65535:65000:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
```

Copyright © 2016, Juniper Networks, Inc.

```
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: **Active Ext>
                Local AS: 65002 Peer AS: 65000
                Age: 4
                Task: BGP_65000.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 65000 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: **Active>
                Local AS: 65002
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
```


Copyright © 2016, Juniper Networks, Inc.

```
1.1.1.4:100:1.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
1.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
22.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
                Source: 2.2.0.0
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Protocol next hop: 2.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     2 Peer AS:     2
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: BGP_2.2.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 2.2.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 2.2.0.0 Metric: 35
                        Push 16
                        Composite next hop: 0x25805988 - INH Session ID: 0x193c
```

Copyright © 2016, Juniper Networks, Inc.

```
                                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                                Indirect path forwarding next hops: 1
                                        Next hop type: Router
                                        Next hop: 10.1.1.1 via ge-1/1/9.0
                                        Session Id: 0x17d8
                                2.2.0.0/32 Originating RIB: inet.3
                                  Metric: 35                    Node path count: 1
                                  Forwarding nexthops: 1
                                        Nexthop: 10.1.1.1 via ge-1/1/9.0
                        BGP     Preference: 170/-1
                                Route Distinguisher: 2:1
                                Next hop type: Indirect
                                Address: 0x9347028
                                Next-hop reference count: 3
                                Source: 2.3.0.0
                                Next hop type: Router, Next hop index: 702
                                Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                                Label operation: Push 634278
                                Label TTL action: prop-ttl
                                Session Id: 0x17d9
                                Protocol next hop: 2.3.0.0
                                Push 16
                                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                                Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                                State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                                Inactive reason: Not Best in its group - IGP metric
                                Local AS:     2 Peer AS:     2
                                Age: 3:34       Metric2: 70
                                Validation State: unverified
                                Task: BGP_2.2.3.0.0+32805
                                Announcement bits (2): 0-KRT 1-BGP_RT_Background
                                AS path: I
                                Communities: target:2:1
                                Import Accepted
                                VPN Label: 16
                                Localpref: 0
                                Router ID: 2.3.0.0
                                Primary Routing Table bgp.l3vpn.0
                                Composite next hops: 1
                                        Protocol next hop: 2.3.0.0 Metric: 70
                                        Push 16
                                        Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                                        Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                                        Indirect path forwarding next hops: 1
                                                Next hop type: Router
                                                Next hop: 10.1.4.2 via ge-1/0/0.0
                                                Session Id: 0x17d9
                                        2.3.0.0/32 Originating RIB: inet.3
                                          Metric: 70                    Node path count: 1
                                          Forwarding nexthops: 1
                                                Nexthop: 10.1.4.2 via ge-1/0/0.0
                                #Multipath Preference: 255
                                Next hop type: Indirect
                                Address: 0x24afca30
                                Next-hop reference count: 1
                                Next hop type: Router
                                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                                Label operation: Push 707633
```


Copyright © 2016, Juniper Networks, Inc.

```
                    Label TTL action: prop-ttl
                    Session Id: 0x17d8
                    Next hop type: Router, Next hop index: 702
                    Next hop: 10.1.4.2 via ge-1/0/0.0
                    Label operation: Push 634278
                    Label TTL action: prop-ttl
                    Session Id: 0x17d9
                    Protocol next hop: 2.2.0.0
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                   Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                    Protocol next hop: 2.3.0.0
                    Push 16
                    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                  Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
0x4000
                    State: <ForwardingOnly Int Ext>
                    Inactive reason: Forwarding use only
                    Age: 23         Metric2: 35
                    Validation State: unverified
                    Task: RT
                    AS path: I
                    Communities: target:2:1
```

show sap listen

Syntax

show sap listen
<brief | detail>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the addresses that the router is listening to in order to receive multicast Session Announcement Protocol (SAP) session announcements.

Options

none—Display standard information about the addresses that the router is listening to in order to receive multicast SAP session announcements.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show sap listen on page 4834
show sap listen brief on page 4834
show sap listen detail on page 4835

Output Fields

Table 435 describes the output fields for the **show sap listen** command. Output fields are listed in the approximate order in which they appear.

Table 435: show sap listen Output Fields

Field Name	Field Description
Group address	Address of the group that the local router is listening to for SAP messages.
Port	UDP port number used for SAP.

Sample Output

show sap listen

```
user@host> show sap listen
Group address   Port
224.2.127.254   9875
239.255.255.255 9875
```

show sap listen brief

The output for the **show sap listen brief** command is identical to that for the **show sap listen** command. For sample output, see show sap listen on page 4834.

Copyright © 2016, Juniper Networks, Inc.

show sap listen detail

The output for the **show sap listen detail** command is identical to that for the **show sap listen** command. For sample output, see show sap listen on page 4834.

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Network Address Translation Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 198

Overview

- Introduction to NAT on page 4839
- Understanding NAT Rule Sets and Rules on page 4840

Introduction to NAT

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

NAT is described in RFC 1631 to solve IP (version 4) address depletion problems. Since then, NAT has been found to be a useful tool for firewalls, traffic redirect, load sharing, network migrations, and so on.

The following types of NAT are supported on Juniper Networks devices:

- Static NAT
- Destination NAT
- Source NAT

NOTE: **SRX Series devices perform both policy lookup and service lookup based on the translated destination port.**

You can use the NAT Wizard to perform basic NAT configuration. To perform more advanced configuration, use the J-Web interface or the CLI.

Related Documentation

- Understanding NAT Rule Sets and Rules on page 4840
- Understanding Static NAT on page 4947
- Understanding Destination NAT on page 4923
- Understanding Source NAT on page 4863
- Understanding Central Point Architecture Enhancements for NAT on page 4880

Copyright © 2016, Juniper Networks, Inc.

Understanding NAT Rule Sets and Rules

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. For example, a rule set can select traffic from a particular interface or to a specific zone. A rule set can contain multiple rules. Once a rule set is found that matches specific traffic, each rule in the rule set is evaluated for a match. Each rule in the rule set further specifies the traffic to be matched and the action to be taken when traffic matches the rule.

This topic includes the following sections:

- NAT Rule Sets on page 4840
- NAT Rules on page 4841
- Rule Processing on page 4841
- NAT Rule Capacity on page 4842

NAT Rule Sets

A rule set specifies a general set of matching conditions for traffic. For static NAT and destination NAT, a rule set specifies one of the following:

- Source interface
- Source zone
- Source routing instance

For source NAT rule sets, you configure both source and destination conditions:

- Source interface, zone, or routing instance
- Destination interface, zone, or routing instance

It is possible for a packet to match more than one rule set; in this case, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match. If a packet matches both a destination NAT rule set that specifies a source zone and a destination NAT rule set that specifies a source interface, the rule set that specifies the source interface is the more specific match.

Source NAT rule set matching is more complex because you specify both source and destination conditions in a source NAT rule set. In the case where a packet matches more than one source NAT rule set, the rule set chosen is based on the following source/destination conditions (in order of priority):

1. Source interface/destination interface
2. Source zone/destination interface
3. Source routing instance/destination interface
4. Source interface/destination zone

Copyright © 2016, Juniper Networks, Inc.

5. Source zone/destination zone
6. Source routing instance/destination zone
7. Source interface/destination routing instance
8. Source zone/destination routing instance
9. Source routing instance/destination routing instance

For example, you can configure rule set A, which specifies a source interface and a destination zone, and rule set B, which specifies a source zone and a destination interface. If a packet matches both rule sets, rule set B is the more specific match.

NOTE: **You cannot specify the same source and destination conditions for source NAT rule sets.**

NAT Rules

Once a rule set that matches the traffic has been found, each rule in the rule set is evaluated in order for a match. NAT rules can match on the following packet information:

- Source and destination address
- Source port (for source and static NAT only)
- Destination port

The first rule in the rule set that matches the traffic is used. If a packet matches a rule in a rule set during session establishment, traffic is processed according to the action specified by that rule.

You can use the **show security nat source rule** and **show security nat destination rule** and the **show security nat static rule** commands to view the number of sessions for a specific rule.

Rule Processing

The NAT type determines the order in which NAT rules are processed. During the first packet processing for a flow, NAT rules are applied in the following order:

1. Static NAT rules
2. Destination NAT rules
3. Route lookup
4. Security policy lookup
5. Reverse mapping of static NAT rules
6. Source NAT rules

Figure 202 illustrates the order for NAT rule processing.

Copyright © 2016, Juniper Networks, Inc.

Figure 202: NAT Rule Processing

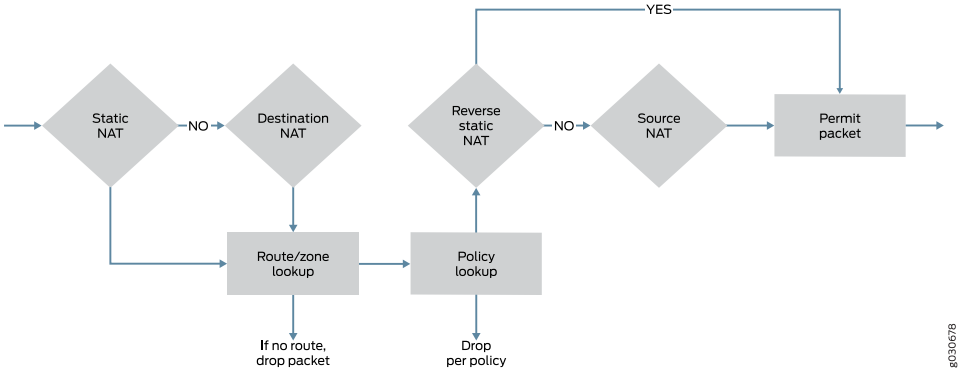

Static NAT and destination NAT rules are processed before route and security policy lookup. Static NAT rules take precedence over destination NAT rules. Reverse mapping of static NAT rules takes place after route and security policy lookup and takes precedence over source NAT rules. Source NAT rules are processed after route and security policy lookup and after reverse mapping of static NAT rules.

The configuration of rules and rule sets is basically the same for each type of NAT—source, destination, or static. But because both destination and static NAT are processed before route lookup, you cannot specify the destination zone, interface or routing instance in the rule set.

NAT Rule Capacity

Table 436 provides the NAT rule capacity requirements per device.

Table 436: Number of Rules on SRX Series Devices

NAT Rule Type	SRX300	SRX320	SRX340	SRX345	SRX1500	SRX5400 SRX5600 SRX5800
Source NAT rule	1024	1024	1024	1024	8192	30720
Destination NAT rule	1024	1024	1024	1024	8192	30720
Static NAT rule	6144	6144	6144	6144	8192	30720

The restriction on the number of rules per rule set is a device-wide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

For memory consumption, there is no guarantee to support these numbers (maximum source rule or rule set + maximum destination rule or rule set + maximum static rule or rule-set) at the same time for SRX5400, SRX5600, and SRX5800 devices.

Copyright © 2016, Juniper Networks, Inc.

Table 437 provides the suggested total number of rules and rule sets for SRX5400, SRX5600, and SRX5800 devices.

Table 437: Number of Rules and Rule Sets

Objects	SRX5400 SRX5600 SRX5800
Total NAT rule sets per system	30,720
Total NAT rules per rule set	30,720

Related Documentation

- Introduction to NAT on page 4839
- Static NAT Configuration Overview on page 4949
- Destination NAT Configuration Overview on page 4926
- Source NAT Configuration Overview on page 4864

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 199

Configuring General NAT Options

- Configuring NAT Using the NAT Wizard on page 4845
- Example: Configuring NAT for Multiple ISPs on page 4845
- Configuring Proxy ARP (CLI Procedure) on page 4857
- Verifying NAT Configuration on page 4858
- Monitoring Incoming Table Information on page 4859
- Monitoring Interface NAT Port Information on page 4860

Configuring NAT Using the NAT Wizard

You can use the NAT Wizard to perform basic NAT configuration. To perform more advanced configuration, use the J-Web interface or the CLI.

To configure NAT using the NAT Wizard:

1. Select **Configure>Tasks>Configure NAT** in the J-Web interface.
2. Click the Launch NAT Wizard button.
3. Follow the wizard prompts.

The upper left area of the wizard page shows where you are in the configuration process. The lower left area of the page shows field-sensitive help. When you click a link under the Resources heading, the document opens in your browser. If the document opens in a new tab, be sure to close only the tab (not the browser window) when you close the document.

Related Documentation

- Introduction to NAT on page 4839

Example: Configuring NAT for Multiple ISPs

This example shows how to configure a Juniper Networks device for address translation of multiple ISPs.

- Requirements on page 4846
- Overview on page 4846

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 4846
- Verification on page 4857

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

In this example, you can configure an SRX Series Services Gateway by connecting the LAN to the Internet by using NAT feature through two ISP connections. In this configuration, trust is the security zone for the private address space and the two untrust security zones for the public address space are used to connect from LAN to the two ISPs and vice versa. The example is a combination of source NAT rules to connect to Internet from the LAN, and destination and static NAT rules to connect to the LAN from Internet.

Configuration

Configuring NAT for Multiple ISPs

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances isp1 instance-type virtual-router
set routing-instances isp1 interface ge-0/0/2.0
set routing-instances isp1 routing-options static route 10.0.0.0/8 next-table inet.0
set routing-instances isp1 routing-options static route 0.0.0.0/0 next-hop 192.0.2.20
set routing-instances isp2 instance-type virtual-router
set routing-instances isp2 interface ge-0/0/3.0
set routing-instances isp2 routing-options static route 10.0.0.0/8 next-table inet.0
set routing-instances isp2 routing-options static route 0.0.0.0/0 next-hop 198.51.100.251
set routing-options interface-routes rib-group inet isp
set routing-options static route 10.0.0.0/8 next-hop 10.0.21.254
set routing-options rib-groups isp import-rib inet.0
set routing-options rib-groups isp import-rib isp1.inet.0
set routing-options rib-groups isp import-rib isp2.inet.0
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match
    source-address any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match
    destination-address any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match application
    any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol then permit
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match
    source-address any
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match
    destination-address any
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match application
    any
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol then permit
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
    source-address any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
    destination-address any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
    application any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol then reject
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
    source-address any
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
    destination-address any
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
    application any
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol then reject
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match
    source-address any
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match
    destination-address ftp-ser
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match
    destination-address telnet-ser
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match application
    junos-ftp
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match application
    junos-telnet
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol then permit
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match
    source-address any
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match
    destination-address 10.171.9.23/32
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match
    destination-address http-ser
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match
    destination-address 10.103.12.0/24
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
    junos-http
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
    junos-icmp-all
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
    junos-dhcp-server
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol then permit
set security nat source pool pool_1 address 192.0.2.40/32 to 192.0.2.190/32
set security nat source pool pool_2 address 192.0.2.250/32
set security nat source pool pool_3 address 198.51.100.20/32 to 198.51.100.30/32
set security nat source address-persistent
set security nat source pool-utilization-alarm raise-threshold 90
set security nat source pool-utilization-alarm clear-threshold 80
set security nat source rule-set SR_SET_1 from zone trust
set security nat source rule-set SR_SET_1 to zone untrust1
set security nat source rule-set SR_SET_1 rule rule1 match source-address 10.11.0.0/16
set security nat source rule-set SR_SET_1 rule rule1 match source-address 10.147.0.0/16
set security nat source rule-set SR_SET_1 rule rule1 match destination-address 0.0.0.0/0
```

Copyright © 2016, Juniper Networks, Inc.

```
set security nat source rule-set SR_SET_1 rule rule1 then source-nat pool pool_1
set security nat source rule-set SR_SET_1 rule rule2 match source-address 10.148.1.0/27
set security nat source rule-set SR_SET_1 rule rule2 match destination-address 0.0.0.0/0
set security nat source rule-set SR_SET_1 rule rule2 then source-nat interface
set security nat source rule-set SR_SET_2 from zone trust
set security nat source rule-set SR_SET_2 to zone untrust2
set security nat source rule-set SR_SET_2 rule rule3 match source-address 10.140.21.0/27
set security nat source rule-set SR_SET_2 rule rule3 then source-nat pool pool_3
set security nat source rule-set SR_SET_2 rule rule4 match source-address 10.150.45.0/24
set security nat source rule-set SR_SET_2 rule rule4 then source-nat off
set security nat destination pool dppol_1 address 10.101.1.10/32
set security nat destination pool dppol_1 address port 21
set security nat destination pool dppol_2 address 10.101.1.11/32
set security nat destination pool dppol_2 address port 2101
set security nat destination pool dppol_3 address 10.103.12.251/32
set security nat destination pool dppol_3 address port 23
set security nat destination pool dppol_4 address 10.103.12.241/32
set security nat destination pool dppol_4 address port 23
set security nat destination pool dppol_5 address 10.103.1.11/32
set security nat destination pool dppol_5 address port 22
set security nat destination rule-set DR_SET1 from routing-instance isp1
set security nat destination rule-set DR_SET1 rule rule1 match destination-address
  192.168.0.10/32
set security nat destination rule-set DR_SET1 rule rule1 match destination-port 7230
set security nat destination rule-set DR_SET1 rule rule1 then destination-nat pool dppol_1
set security nat destination rule-set DR_SET1 rule rule2 match destination-address
  192.169.1.0/24
set security nat destination rule-set DR_SET1 rule rule2 then destination-nat pool dppol_2
set security nat destination rule-set DR_SET2 from routing-instance isp2
set security nat destination rule-set DR_SET2 rule rule3 match destination-address
  192.168.2.2/32
set security nat destination rule-set DR_SET2 rule rule3 match destination-port 7351
set security nat destination rule-set DR_SET2 rule rule3 then destination-nat pool dppol_3
set security nat destination rule-set DR_SET2 rule rule4 match destination-address
  192.168.4.171/32
set security nat destination rule-set DR_SET2 rule rule4 match destination-port 3451
set security nat destination rule-set DR_SET2 rule rule4 then destination-nat pool dppol_4
set security nat static rule-set ST_SET1 from zone trust
set security nat static rule-set ST_SET1 rule rule1 match destination-address 10.0.10.0/24
set security nat static rule-set ST_SET1 rule rule1 then static-nat prefix 192.168.5.0/24
set security nat static rule-set ST_SET2 from routing-instance isp1
set security nat static rule-set ST_SET2 rule rule2 match destination-address
  192.168.6.0/24
set security nat static rule-set ST_SET2 rule rule2 then static-nat prefix 10.107.30.0/24
set security nat static rule-set ST_SET2 rule rule3 match destination-address
  192.168.7.2/32
set security nat static rule-set ST_SET2 rule rule3 then static-nat prefix 10.171.9.23/32
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure routing instances.

```
[edit ]
user@host# set routing-instances isp1 instance-type virtual-router
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set routing-instances isp1 interface ge-0/0/2.0
user@host# set routing-instances isp1 routing-options static route 10.0.0.0/8
  next-table inet.0
user@host# set routing-instances isp1 routing-options static route 0.0.0.0/0
  next-hop 192.0.2.20
user@host# set routing-instances isp2 instance-type virtual-router
user@host# set routing-instances isp2 interface ge-0/0/3.0
user@host# set routing-instances isp2 routing-options static route 10.0.0.0/8
  next-table inet.0
user@host# set routing-instances isp2 routing-options static route 0.0.0.0/0
  next-hop 198.51.100.251
```

2. Configure rib groups and routing options.

```
[edit ]
user@host# set routing-options interface-routes rib-group inet isp
user@host# set routing-options static route 10.0.0.0/8 next-hop 10.0.21.254
user@host# set routing-options rib-groups isp import-rib inet.0
user@host# set routing-options rib-groups isp import-rib isp1.inet.0
user@host# set routing-options rib-groups isp import-rib isp2.inet.0
```

3. Configure security policies.

```
[edit security policies]
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match
  source-address any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match
  destination-address any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match
  application any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol then permit
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match
  source-address any
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match
  destination-address any
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match
  application any
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol then permit
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
  source-address any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
  destination-address anyfrom-zone untrust1 to-zone untrust2 policy untr1-untr2-pol
  match destination-address any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match
  application any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol then
  reject
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
  source-address any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
  destination-address any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
  application any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol then
  reject
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match
  source-address any
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match
  destination-address ftp-ser
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match
  destination-address telnet-ser
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match
  application junos-ftp
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match
  application junos-telnet
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol then permit
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  source-address any
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  destination-address 10.171.9.23/32
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  destination-address http-ser
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  destination-address 10.103.12.0/24
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  application junos-http
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  application junos-icmp-all
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match
  application junos-dhcp-server
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol then permit
```

4. Configure source NAT pools and rules.

```
[edit security nat]
user@host# set source pool pool_1 address 192.0.2.40/32 to 192.0.2.190/32
user@host# set source pool pool_2 address 192.0.2.250/32
user@host# set source pool pool_3 address198.51.100.20/32 to 198.51.100.30/32
user@host# set source address-persistent
user@host# set source pool-utilization-alarm raise-threshold 90
user@host# set source pool-utilization-alarm clear-threshold 80
user@host# set source rule-set SR_SET_1 from zone trust
user@host# set source rule-set SR_SET_1 to zone untrust1
user@host# set source rule-set SR_SET_1 rule rule1 match source-address 10.11.0.0/16
user@host# set source rule-set SR_SET_1 rule rule1 match source-address
  10.147.0.0/16
user@host# set source rule-set SR_SET_1 rule rule1 match destination-address
  0.0.0.0/0
user@host# set source rule-set SR_SET_1 rule rule1 then source-nat pool pool_1
user@host# set source rule-set SR_SET_1 rule rule2 match source-address
  10.148.1.0/27
user@host# set source rule-set SR_SET_1 rule rule2 match destination-address
  0.0.0.0/0
user@host# set source rule-set SR_SET_1 rule rule2 then source-nat interface
user@host# set source rule-set SR_SET_2 from zone trust
user@host# set source rule-set SR_SET_2 to zone untrust2
user@host# set source rule-set SR_SET_2 rule rule3 match source-address
  10.140.21.0/27
user@host# set source rule-set SR_SET_2 rule rule3 then source-nat pool pool_3
user@host# set source rule-set SR_SET_2 rule rule4 match source-address
  10.150.45.0/24
user@host# set source rule-set SR_SET_2 rule rule4 then source-nat off
```

5. Configure destination NAT pools and rules.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat]
user@host#set destination pool dppol_1 address 10.101.1.10/32
user@host#set destination pool dppol_1 address port 21
user@host#set destination pool dppol_2 address 10.101.1.11/32
user@host#set destination pool dppol_2 address port 2101
user@host#set destination pool dppol_3 address 10.103.12.251/32
user@host#set destination pool dppol_3 address port 23
user@host#set destination pool dppol_4 address 10.103.12.241/32
user@host#set destination pool dppol_4 address port 23
user@host#set destination pool dppol_5 address 10.103.1.11/32
user@host#set destination pool dppol_5 address port 22
user@host#set destination rule-set DR_SET1 from routing-instance isp1
user@host#set destination rule-set DR_SET1 rule rule1 match destination-address
  192.168.0.10/32
user@host#set destination rule-set DR_SET1 rule rule1 match destination-port 7230
user@host#set destination rule-set DR_SET1 rule rule1 then destination-nat pool
  dppol_1
user@host#set destination rule-set DR_SET1 rule rule2 match destination-address
  192.169.1.0/24
user@host#set destination rule-set DR_SET1 rule rule2 then destination-nat pool
  dppol_2
user@host#set destination rule-set DR_SET2 from routing-instance isp2
user@host#set destination rule-set DR_SET2 rule rule3 match destination-address
  192.168.2.2/32
user@host#set destination rule-set DR_SET2 rule rule3 match destination-port 7351
user@host#set destination rule-set DR_SET2 rule rule3 then destination-nat pool
  dppol_3
user@host#set destination rule-set DR_SET2 rule rule4 match destination-address
  192.168.4.171/32
user@host#set destination rule-set DR_SET2 rule rule4 match destination-port
  3451
user@host#set destination rule-set DR_SET2 rule rule4 then destination-nat pool
  dppol_4
```

6. Configure static NAT rules.

```
[edit security nat]
user@host#set static rule-set ST_SET1 from zone trust
user@host#set static rule-set ST_SET1 rule rule1 match destination-address
  10.0.10.0/24
user@host#set static rule-set ST_SET1 rule rule1 then static-nat prefix 192.168.5.0/24
user@host#set static rule-set ST_SET2 from routing-instance isp1
user@host#set static rule-set ST_SET2 rule rule2 match destination-address
  192.168.6.0/24
user@host#set static rule-set ST_SET2 rule rule2 then static-nat prefix
  10.107.30.0/24
user@host#set static rule-set ST_SET2 rule rule3 match destination-address
  192.168.7.2/32
user@host#set static rule-set ST_SET2 rule rule3 then static-nat prefix 10.171.9.23/32
```

Results From configuration mode, confirm your configuration by entering **show configuration** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show configuration routing-intances
  routing-instances {
```

Copyright © 2016, Juniper Networks, Inc.

```
    isp1 {
      instance-type virtual-router;
      interface ge-0/0/2.0;
      routing-options {
        static {
          route 10.0.0.0/8 next-table inet.0;
          route 0.0.0.0/0 next-hop 192.0.2.20;
        }
      }
    }
    isp2 {
      instance-type virtual-router;
      interface ge-0/0/3.0;
      routing-options {
        static {
          route 10.0.0.0/8 next-table inet.0;
          route 0.0.0.0/0 next-hop 198.51.100.251;
        }
      }
    }
  }
```

```
user@host# show configuration routing-options
routing-options {
  interface-routes {
    rib-group inet isp;
  }
  static {
    route 10.0.0.0/8 next-hop 10.0.21.254;
  }
  rib-groups {
    isp {
      import-rib [ isp1.inet.0 isp2.inet.0 ];
    }
  }
}
```

```
user@host# show configuration policies
policies {
  from-zone trust to-zone untrust1 {
    policy tr-untr1-pol {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone trust to-zone untrust2 {
    policy tr-untr2-pol {
      match {
        source-address any;
        destination-address any;
        application any;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          permit;
        }
      }
    }
    from-zone untrust1 to-zone untrust2 {
      policy untr1-untr2-pol {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          reject;
        }
      }
    }
    from-zone untrust2 to-zone untrust1 {
      policy untr2-untr1-pol {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          reject;
        }
      }
    }
    from-zone untrust1 to-zone trust {
      policy untr1-tr-pol {
        match {
          source-address any;
          destination-address [ ftp-ser telnet-ser ];
          application [ junos-ftp junos-telnet ];
        }
        then {
          permit;
        }
      }
    }
    from-zone untrust2 to-zone trust {
      policy untr2-tr-pol {
        match {
          source-address any;
          destination-address [ 10.171.9.23/32 http-ser 10.103.12.0/24 ];
          application [ junos-http junos-icmp-all junos-dhcp-server ];
        }
        then {
          permit;
        }
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show configuration security nat
security {
    nat {
        source {
            pool pool_1 {
                address {
                    192.0.2.40/32 to 192.0.2.190/32;
                }
            }
            pool pool_2 {
                address {
                    192.0.2.250/32;
                }
            }
            pool pool_3 {
                address {
                    198.51.100.20/32 to 198.51.100.30/32;
                }
            }
            address-persistent;
            pool-utilization-alarm raise-threshold 90 clear-threshold 80;
            rule-set SR_SET_1 {
                from zone trust;
                to zone untrust1;
                rule rule1 {
                    match {
                        source-address [ 10.11.0.0/16 10.147.0.0/16 ];
                        destination-address 0.0.0.0/0;
                    }
                    then {
                        source-nat {
                            pool {
                                pool_1;
                            }
                        }
                    }
                }
                rule rule2 {
                    match {
                        source-address 10.148.1.0/27;
                        destination-address 0.0.0.0/0;
                    }
                    then {
                        source-nat {
                            interface;
                        }
                    }
                }
            }
            rule-set SR_SET_2 {
                from zone trust;
                to zone untrust2;
                rule rule3 {
                    match {
                        source-address 10.140.21.0/27;
                    }
```


Copyright © 2016, Juniper Networks, Inc.

```
            then {
              source-nat {
                pool {
                  pool_3;
                }
              }
            }
          }
          rule rule4 {
            match {
              source-address 10.150.45.0/24;
            }
            then {
              source-nat {
                off;
              }
            }
          }
        }
      }
```

```
user@host# show configuration security nat
destination {
  pool dppol_1 {
    address 10.101.1.10/32 port 21;
  }
  pool dppol_2 {
    address 10.101.1.11/32 port 2101;
  }
  pool dppol_3 {
    address 10.103.12.251/32 port 23;
  }
  pool dppol_4 {
    address 10.103.12.241/32 port 23;
  }
  pool dppol_5 {
    address 10.103.1.11/32 port 22;
  }
  rule-set DR_SET1 {
    from routing-instance isp1;
    rule rule1 {
      match {
        destination-address 192.168.0.10/32;
        destination-port 7230;
      }
      then {
        destination-nat pool dppol_1;
      }
    }
    rule rule2 {
      match {
        destination-address 192.169.1.0/24;
      }
      then {
        destination-nat pool dppol_2;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  rule-set DR_SET2 {
    from routing-instance isp2;
    rule rule3 {
      match {
        destination-address 192.168.2.2/32;
        destination-port 7351;
      }
      then {
        destination-nat pool dppol_3;
      }
    }
    rule rule4 {
      match {
        destination-address 192.168.4.171/32;
        destination-port 3451;
      }
      then {
        destination-nat pool dppol_4;
      }
    }
  }
}
```

```
user@host# show configuration static nat
static {
  rule-set ST_SET1 {
    from zone trust;
    rule rule1 {
      match {
        destination-address 10.0.10.0/24;
      }
      then {
        static-nat prefix 192.168.5.0/24;
      }
    }
  }
  rule-set ST_SET2 {
    from routing-instance isp1;
    rule rule2 {
      match {
        destination-address 192.168.6.0/24;
      }
      then {
        static-nat prefix 10.107.30.0/24;
      }
    }
    rule rule3 {
      match {
        destination-address 192.168.7.2/32;
      }
      then {
        static-nat prefix 10.171.9.23/32;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Interfaces

Purpose Verify that the interfaces are configured correctly.

Action From operational mode, enter the following commands:

- **show interfaces**
- **show zones**
- **show routing-instances**
- **show routing-options**
- **show policies**
- **show source nat**
- **show destination nat**
- **show static nat**

Related Documentation

- Introduction to NAT on page 4839

Configuring Proxy ARP (CLI Procedure)

You use NAT proxy ARP functionality to configure proxy ARP entries for IP addresses that require either source or destination NAT and that are in the same subnet as the ingress interface.

NOTE: On SRX Series devices, you must explicitly configure NAT proxy ARP.

When configuring NAT proxy ARP, you must specify the logical interface on which to configure proxy ARP. Then you enter an address or address range.

The device performs proxy ARP for the following conditions:

- When addresses defined in the static NAT and source NAT pool are in the same subnet as that of the ingress interface
- When addresses in the original destination address entry in the destination NAT rules are in the same subnet as that of the ingress interface

```
user@host# set security nat proxy-arp interface fe-0/0/0.0 address 10.1.1.10 to 10.1.1.20
```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Static NAT Configuration Overview on page 4949
- Destination NAT Configuration Overview on page 4926
- Source NAT Configuration Overview on page 4864

Verifying NAT Configuration

Purpose The NAT trace options hierarchy configures trace file and flags for verification purposes.

SRX Series devices have two main components: the Routing Engine (RE) and the Packet Forwarding Engine (PFE). The PFE is divided into the ukernel portion and the real-time portion.

When a NAT configuration is committed, the configuration is first checked and validated on the RE. After validation, the configuration is pushed to the PFE. The configuration is installed on the ukernel PFE, then action is taken on each packet that matches NAT rules on the real-time PFE.

For verification, you can turn on flags individually to debug NAT functionality on the RE, ukernel PFE, or real-time PFE:

- The **nat-re** flag records the trace of the NAT configuration validation on the RE and the configuration push to the PFE.
- The **nat-pfe** flag records the trace of the NAT configuration installation on the ukernel PFE.
- The **nat-rt** flag records the trace of the NAT rule match, and subsequent action on the real-time PFE.

The trace data is written to **/var/log/security-trace** by default, and can be viewed using the command **show log security-trace**.

NOTE: If session logging has been enabled in the policy configurations on the device, the session logs will include specific NAT details for each session. See "Monitoring Policy Statistics" on page 1162 for information on how to enable session logging and "Information Provided in Session Log Entries for SRX Series Services Gateways" on page 1787 for a description of information provided in session logs.

Action To verify that NAT configurations are correctly updated to the device upon commit, and that the NAT rule match and subsequent actions are correct, use the **security nat traceoptions** statement.

```
user@host# set security nat traceoptions flag all
user@host# set security nat traceoptions flag destination-nat-pfe
user@host# set security nat traceoptions flag destination-nat-re
user@host# set security nat traceoptions flag destination-nat-rti
user@host# set security nat traceoptions flag source-nat-pfe
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security nat traceoptions flag source-nat-re
user@host# set security nat traceoptions flag source-nat-rt
user@host# set security nat traceoptions flag static-nat-pfe
user@host# set security nat traceoptions flag static-nat-re
user@host# set security nat traceoptions flag static-nat-rt
```

To verify that NAT translations are being applied to the traffic, and to view individual traffic flow processing with NAT translations, use both the **security nat traceoptions** command and the **security flow traceoptions** command together. The commands are used together because the NAT trace, configured using the **security nat traceoptions** command, is not recorded unless the **flow traceoptions** command is also configured.

To filter a specific flow, you can define a packet filter and use it as a traceoption :

```
user@host# set security flow traceoptions packet-filter packet-filter
user@host# set security flow traceoptions packet-filter packet-filter apply-groups
user@host# set security flow traceoptions packet-filter packet-filter apply-groups-except
user@host# set security flow traceoptions packet-filter packet-filter destination-port
user@host# set security flow traceoptions packet-filter packet-filter destination-prefix
user@host# set security flow traceoptions packet-filter packet-filter interface
user@host# set security flow traceoptions packet-filter packet-filter protocol
user@host# set security flow traceoptions packet-filter packet-filter source-port
user@host# set security flow traceoptions packet-filter packet-filter source-prefix
```

To verify NAT traffic and to enable all traffic trace in data plane, use the traceoptions **set security flow traceoptions flag basic-datapath** command, as shown in the following example using a simple packet filter:

```
user@host# set security flow traceoptions file filename
user@host# set security flow traceoptions flag basic-datapath
user@host# set security flow traceoptions packet-filter client-traffic source-prefixprefix
user@host# set security flow traceoptions packet-filter client-traffic
  destination-prefixprefix
user@host# set security nat traceoptions flag all
```

Related Documentation

- traceoptions (Security NAT) on page 5117
- traceoptions (Security Flow) on page 1890
- Static NAT Configuration Overview on page 4949
- Destination NAT Configuration Overview on page 4926
- Source NAT Configuration Overview on page 4864

Monitoring Incoming Table Information

Purpose View NAT table information.

Action Select **Monitor>NAT>Incoming Table** in the J-Web user interface, or enter the following CLI command:

show security nat incoming-table

Table 438 summarizes key output fields in the incoming table display.

Copyright © 2016, Juniper Networks, Inc.

Table 438: Summary of Key Incoming Table Output Fields

Field	Values
Statistics	
In use	Number of entries in the NAT table.
Maximum	Maximum number of entries possible in the NAT table.
Entry allocation failed	Number of entries failed for allocation.
Incoming Table	
Clear	
Destination	Destination IP address and port number.
Host	Host IP address and port number that the destination IP address is mapped to.
References	Number of sessions referencing the entry.
Timeout	Timeout, in seconds, of the entry in the NAT table.
Source-pool	Name of source pool where translation is allocated.

Related Documentation

- Monitoring Source NAT Information on page 4889
- Monitoring Destination NAT Information on page 4944
- Monitoring Static NAT Information on page 4965
- Monitoring Interface NAT Port Information on page 4860

Monitoring Interface NAT Port Information

Purpose View port usage for an interface source pool information.

Action Select **Monitor>Firewall/NAT>Interface NAT** in the J-Web user interface, or enter the following CLI command:

- **show security nat interface-nat-ports**

Table 439 summarizes key output fields in the interface NAT display.

Table 439: Summary of Key Interface NAT Output Fields

Field	Values	Additional Information
Interface NAT Summary Table		
Pool Index	Port pool index.	–

Copyright © 2016, Juniper Networks, Inc.

Table 439: Summary of Key Interface NAT Output Fields *(continued)*

Field	Values	Additional Information
Total Ports	Total number of ports in a port pool.	–
Single Ports Allocated	Number of ports allocated one at a time that are in use.	–
Single Ports Available	Number of ports allocated one at a time that are free for use.	–
Twin Ports Allocated	Number of ports allocated two at a time that are in use.	–
Twin Ports Available	Number of ports allocated two at a time that are free for use.	–

Related Documentation

- Monitoring Source NAT Information on page 4889
- Monitoring Destination NAT Information on page 4944
- Monitoring Static NAT Information on page 4965
- Monitoring Incoming Table Information on page 4859

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 200

Configuring Source NAT

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Example: Configuring Source NAT for Egress Interface Translation on page 4865
- Example: Configuring Source NAT for Single Address Translation on page 4869
- Example: Configuring Source and Destination NAT Translations on page 4873
- Understanding Central Point Architecture Enhancements for NAT on page 4880
- Understanding Reverse NAT Enhancements for Central Point Architecture on page 4880
- Understanding Source NAT Rules on page 4881
- Example: Configuring Source NAT with Multiple Rules on page 4882
- Disabling Port Randomization for Source NAT (CLI Procedure) on page 4888
- Monitoring Source NAT Information on page 4889

Understanding Source NAT

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

Source NAT allows connections to be initiated only for outgoing network connections—for example, from a private network to the Internet. Source NAT is commonly used to perform the following translations:

- Translate a single IP address to another address (for example, to provide a single device in a private network with access to the Internet).
- Translate a contiguous block of addresses to another block of addresses of the same size.
- Translate a contiguous block of addresses to another block of addresses of smaller size.
- Translate a contiguous block of addresses to a single IP address or a smaller block of addresses using port translation.
- Translate a contiguous block of addresses to the address of the egress interface.

Copyright © 2016, Juniper Networks, Inc.

Translation to the address of the egress interface does not require an address pool; all other source NAT translations require configuration of an address pool. One-to-one and many-to-many translations for address blocks of the same size do not require port translation because there is an available address in the pool for every address that would be translated.

If the size of the address pool is smaller than the number of addresses that would be translated, either the total number of concurrent addresses that can be translated is limited by the size of the address pool or port translation must be used. For example, if a block of 253 addresses is translated to an address pool of 10 addresses, a maximum of 10 devices can be connected concurrently unless port translation is used.

The following types of source NAT are supported:

- Translation of the original source IP address to the egress interface's IP address (also called interface NAT). Port address translation is always performed.
- Translation of the original source IP address to an IP address from a user-defined address pool without port address translation. The association between the original source IP address to the translated source IP address is dynamic. However, once there is an association, the same association is used for the same original source IP address for new traffic that matches the same NAT rule.
- Translation of the original source IP address to an IP address from a user-defined address pool with port address translation. The association between the original source IP address to the translated source IP address is dynamic. Even if an association exists, the same original source IP address may be translated to a different address for new traffic that matches the same NAT rule.
- Translation of the original source IP address to an IP address from a user-defined address pool by shifting the IP addresses. This type of translation is one-to-one, static, and without port address translation. If the original source IP address range is larger than the IP address range in the user-defined pool, untranslated packets are dropped.

Related Documentation

- Understanding Source NAT Pools on page 4895
- Understanding Source NAT Rules on page 4881
- Source NAT Configuration Overview on page 4864
- Understanding Reverse NAT Enhancements for Central Point Architecture on page 4880

Source NAT Configuration Overview

The main configuration tasks for source NAT are as follows:

1. Configure an address pool or an interface NAT mapping of private addresses to the public address of an egress interface.

   For an address pool, also do the following:

   a. Specify the name of the pool, the addresses or address ranges, the routing instance, and whether to perform port address translation (PAT).

Copyright © 2016, Juniper Networks, Inc.

b. (Optional) Configure address pool options, such as overflow pool, IP address shifting, address sharing, address pooling, and pool utilization alarms.

c. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

2. (Optional) Configure the persistent address.
3. Configure source NAT rules that align with your network and security requirements.

Related Documentation

- Understanding Source NAT on page 4863
- Understanding Source NAT Pools on page 4895
- Understanding Source NAT Rules on page 4881
- **pool (Security Source NAT) on page 5076**
- **interface (Security Source NAT) on page 5066**
- **rule (Security Source NAT) on page 5093**
- **persistent-nat on page 5074**

Example: Configuring Source NAT for Egress Interface Translation

This example describes how to configure a source NAT mapping of private addresses to the public address of an egress interface.

- Requirements on page 4865
- Overview on page 4865
- Configuration on page 4867
- Verification on page 4868

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 203, devices with private addresses in the trust zone access a public network through the egress interface ge-0/0/0. For packets that enter the Juniper Networks security device from the trust zone with a destination address in the untrust zone, the source IP address is translated to the IP address of the egress interface.

Copyright © 2016, Juniper Networks, Inc.

NOTE: No source NAT pool is required for source NAT using an egress interface. Proxy ARP does not need to be configured for the egress interface.

Figure 203: Source NAT Egress Interface Translation

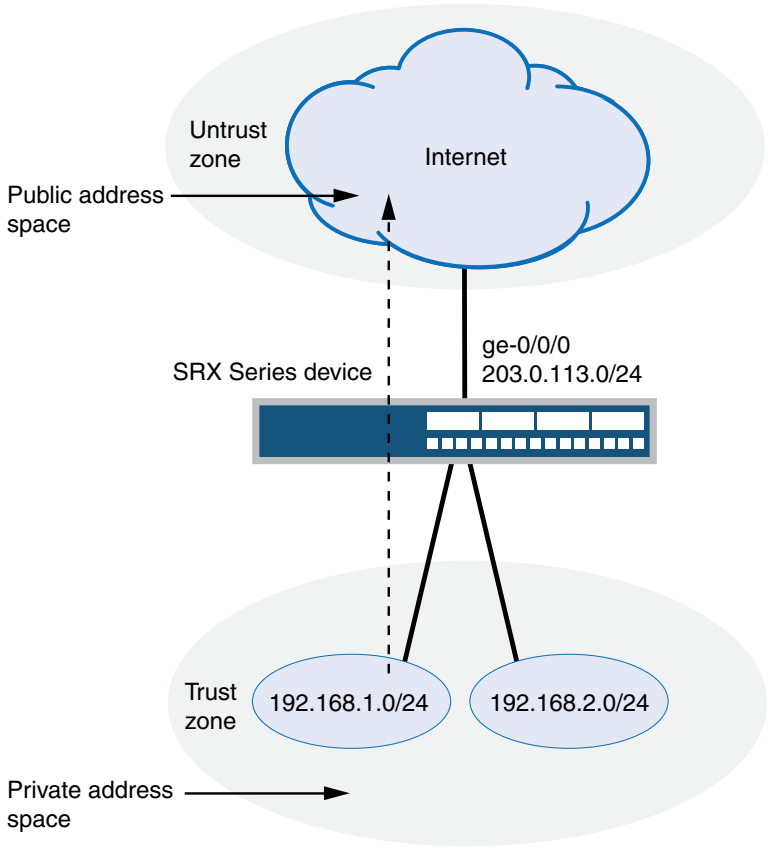

Original Source IP	Translated Source IP
0.0.0.0/0	203.0.113.63 (Interface IP)

g030668

This example describes the following configurations:

- Source NAT rule set **rs1** with a rule **r1** to match any packet from the trust zone to the untrust zone. For matching packets, the source address is translated to the IP address of the egress interface.
- Security policies to permit traffic from the trust zone to the untrust zone.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat interface
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT translation to an egress interface:

1. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

2. Configure a rule that matches packets and translates the source address to the address of the egress interface.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 then source-nat interface
   ```

3. Configure a security policy that allows traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy internet-access match source-address any
     destination-address any application any
   user@host# set policy internet-access then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  rule-set rs1 {
```

Copyright © 2016, Juniper Networks, Inc.

```
        from zone trust;
        to zone untrust;
        rule r1 {
          match {
            source-address 0.0.0.0/0;
            destination-address 0.0.0.0/0;
          }
          then {
            source-nat {
              interface;
            }
          }
        }
      }
    }
    user@host# show security policies
    from-zone trust to-zone untrust {
      policy internet-access {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Rule Usage on page 4868
- Verifying NAT Application to Traffic on page 4868

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864

Example: Configuring Source NAT for Single Address Translation

This example describes how to configure a source NAT mapping of a single private address to a public address.

- Requirements on page 4869
- Overview on page 4869
- Configuration on page 4871
- Verification on page 4873

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 204, a device with the private address 192.168.1.200 in the trust zone accesses a public network. For packets sent by the device to a destination address in the untrust zone, the Juniper Networks security device translates the source IP address to the public IP address 203.0.113.200/32.

Copyright © 2016, Juniper Networks, Inc.

Figure 204: Source NAT Single Address Translation

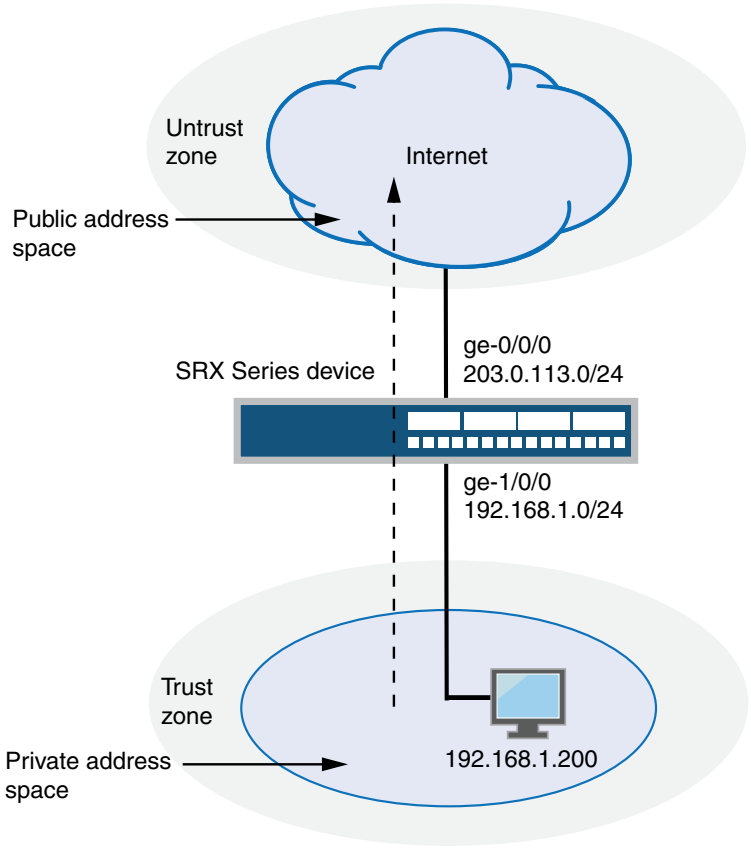

Original Source IP	Translated Source IP
192.168.1.200/32	203.0.113.200/32

g030669

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address 203.0.113.200/32.
- Source NAT rule set **rs1** with rule **r1** to match packets from the trust zone to the untrust zone with the source IP address 192.168.1.200/32. For matching packets, the source address is translated to the IP address in **src-nat-pool-1** pool.
- Proxy ARP for the address 203.0.113.200 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the trust zone to the untrust zone.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.200/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.200/32
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT translation for a single IP address:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.200/32
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 192.168.1.200/32
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```

4. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200
   ```

5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy internet-access match source-address any
     destination-address any application any
   user@host# set policy internet-access then permit
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
   pool src-nat-pool-1 {
      address {
         203.0.113.200/32;
      }
   }
   rule-set rs1 {
      from zone trust;
      to zone untrust;
      rule r1 {
         match {
            source-address 192.168.1.200/32;
         }
         then {
            source-nat {
               pool {
                  src-nat-pool-1;
               }
            }
         }
      }
   }
}
   proxy-arp {
      interface ge-0/0/0.0 {
         address {
            203.0.113.200/32;
         }
      }
   }
user@host# show security policies
from-zone trust to-zone untrust {
   policy internet-access {
      match {
         source-address any;
         destination-address any;
         application any;
      }
      then {
         permit;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4873
- Verifying Source NAT Rule Usage on page 4873
- Verifying NAT Application to Traffic on page 4873

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864

Example: Configuring Source and Destination NAT Translations

This example describes how to configure both source and destination NAT mappings.

- Requirements on page 4873
- Overview on page 4874
- Configuration on page 4875
- Verification on page 4878

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.

Copyright © 2016, Juniper Networks, Inc.

2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 205, the following translations are performed on the Juniper Networks security device:

- The source IP address in packets sent by the device with the private address 192.168.1.200 in the trust zone to any address in the untrust zone is translated to a public address in the range from 203.0.113.10 through 203.0.113.14.
- The destination IP address 203.0.113.100/32 in packets sent from the trust zone to the untrust zone is translated to the address 10.1.1.200/32.

Figure 205: Source and Destination NAT Translations

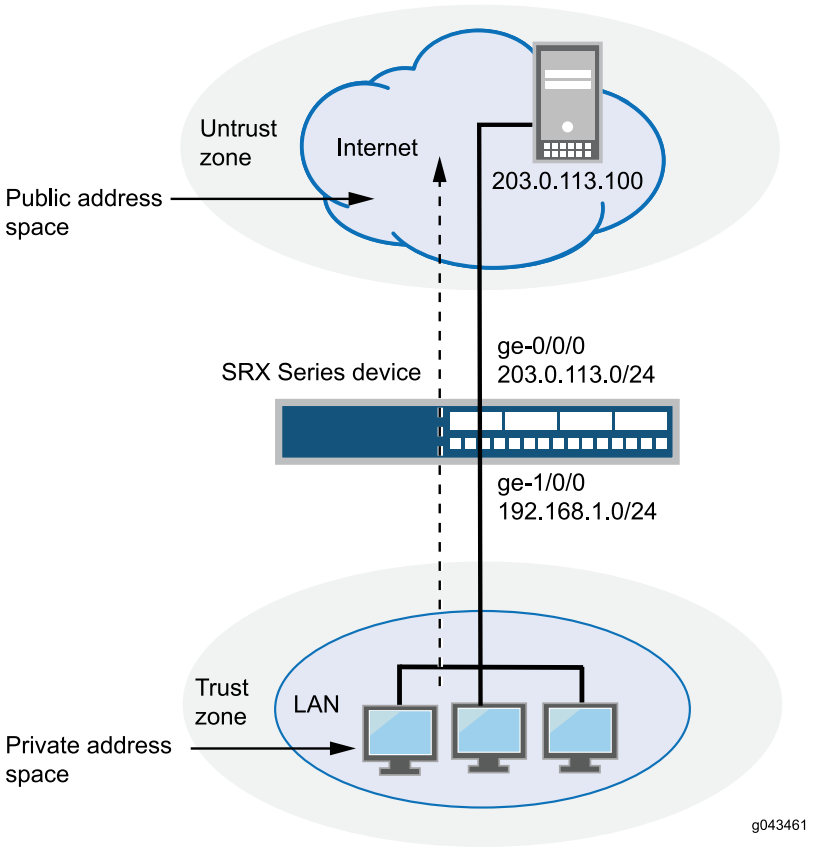

Copyright © 2016, Juniper Networks, Inc.

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address range 203.0.113.10 through 203.0.113.14.
- Source NAT rule set **rs1** with rule **r1** to match any packets from the trust zone to the untrust zone. For matching packets, the source address is translated to an IP address in the **src-nat-pool-1** pool.
- Destination NAT pool **dst-nat-pool-1** that contains the IP address 10.1.1.200/32.
- Destination NAT rule set **rs1** with rule **r1** to match packets from the trust zone with the destination IP address 203.0.113.100. For matching packets, the destination address is translated to the IP address in the **dst-nat-pool-1** pool.
- Proxy ARP for the addresses 203.0.113.10 through 203.0.113.14 and 203.0.113.100/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policy to permit traffic from the trust zone to the untrust zone.
- Security policy to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.10/32 to 203.0.113.14/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat destination pool dst-nat-pool-1 address 10.1.1.200/32
set security nat destination rule-set rs1 from zone untrust
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.100/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.10/32 to 203.0.113.24/32
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.100/32
set security policies from-zone trust to-zone untrust policy internet-access match
    source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
    destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
    application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
set security address-book global address dst-nat-pool-1 10.1.1.200/32
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match
    source-address any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match
    destination-address dst-nat-pool-1
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match
    application any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access then
    permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the source and destination NAT translations:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.10 to 203.0.113.14
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

3. Configure a rule that matches packets and translates the source address to an address in the source NAT pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```

4. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set pool dst-nat-pool-1 address 10.1.1.200/32
   ```

5. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 from zone untrust
   ```

6. Configure a rule that matches packets and translates the destination address to the address in the destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.100/32
   user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
   ```

7. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.10 to 203.0.113.14
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.100
   ```

8. Configure a security policy that allows traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy internet-access match source-address any
     destination-address any application any
   user@host# set policy internet-access then permit
   ```

9. Configure an address in the global address book.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security address-book global]
user@host# set address dst-nat-pool-1 10.1.1.200/32
```

10. Configure a security policy that allows traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy dst-nat-pool-1-access match source-address any
  destination-address dst-nat-pool-1 application any
user@host# set policy dst-nat-pool-1-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.10/32 to 203.0.113.14/32;
    }
  }
  rule-set rs1 {
    to zone untrust;
    rule r1 {
      match {
        source-address 0.0.0.0/0;
        destination-address 0.0.0.0/0;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
  destination {
    pool dst-nat-pool-1 {
      address 10.1.1.200/32;
    }
    rule-set rs1 {
      from zone untrust;
      rule r1 {
        match {
          destination-address 203.0.113.100/32;
        }
        then {
          destination-nat pool dst-nat-pool-1;
        }
      }
    }
  }
  proxy-arp {
```

Copyright © 2016, Juniper Networks, Inc.

```
        interface ge-0/0/0.0 {
          address {
            203.0.113.10/32 to 203.0.113.14/32;
            203.0.113.100/32;
          }
        }
      }
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
  }
  policy internet-access {
    then {
      permit;
    }
  }
}
  from-zone untrust to-zone trust {
    policy dst-nat-pool-1-access {
      match {
        source-address any;
        destination-address dst-nat-pool-1;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4878
- Verifying Source NAT Rule Usage on page 4879
- Verifying Destination NAT Pool Usage on page 4879
- Verifying Destination NAT Rule Usage on page 4879
- Verifying NAT Application to Traffic on page 4879

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Copyright © 2016, Juniper Networks, Inc.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying Destination NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the destination NAT pool.

Action From operational mode, enter the **show security nat destination pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose Verify that there is traffic matching the destination NAT rule.

Action From operational mode, enter the **show security nat destination rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Destination NAT on page 4923
- Destination NAT Configuration Overview on page 4926

Copyright © 2016, Juniper Networks, Inc.

Understanding Central Point Architecture Enhancements for NAT

System session capacity and session ramp-up rate are limited by central point memory capacity and CPU capacity. The central point architecture has been enhanced to handle higher system session capacity and session ramp-up rate for the SRX5000 line. Hence, the workload on the central point is reduced to increase the session capacity and to support more sessions to achieve higher connections per second (CPS). The following list describes the enhancements to NAT to improve performance:

- The central point architecture no longer supports central point sessions. Therefore, NAT needs to maintain a NAT tracker to track the IP address or port allocation and usage. NAT tracker is a global array for SPU session ID to NAT IP or port mapping that is used to manage NAT resources.
- By default, a NAT rule alarm and trap statistics counter update message is sent from the Services Processing Unit (SPU) to the central point at intervals of 1 second instead of updating the statistics based on each session trigger in the central point system.
- To support a specific NAT IP address or port allocated such that the 5-tuple hash after NAT is the same as the original 5-tuple hash before NAT, select a NAT port that results in the same hash as the original hash by the specific calculation. Hence, the forwarding session is reduced. When NAT is used, the reverse wing is hashed to a different SPU. A forward session has to be installed to forward reverse traffic to a session SPU. NAT tries to select a port that can be used by the hash algorithm to make the reverse wing be hashed to the same SPU as the initial wing. So, both NAT performance and throughput are improved with this approach.
- To improve NAT performance, IP shifting pool (non-PAT pool) management is moved from the central point to the SPU so that all local NAT resources for that pool are managed locally instead of sending the NAT request to the central point. Hence, IP address-shifting NAT pool connections per second and throughput are improved.

Related Documentation

- Understanding Reverse NAT Enhancements for Central Point Architecture on page 4880

Understanding Reverse NAT Enhancements for Central Point Architecture

To improve NAT performance, you can enable reverse NAT to select the port based on the predefined algorithms, that is, randomized and round-robin.

NOTE: **If the selected Services Processing Unit (SPU) of two wings of a session is different, then by default reverse NAT is disabled and forwarding session is installed.**

The device enables randomized and round-robin algorithms to select a port by default. However, the randomized algorithm has higher priority than the round robin algorithm. When both randomized and round-robin algorithms are enabled, the randomized algorithm is applied for port selection because randomization has higher priority.

Copyright © 2016, Juniper Networks, Inc.

When randomized algorithm is disabled, apply round-robin algorithm for port selection by using the following command:

set security nat source port-randomization disable

When round-robin algorithm is disabled, apply randomized algorithm for port selection by using the following command:

set security nat source port-round-robin disable

When both randomized and round-robin algorithms are disabled, activate reverse NAT enhancement or session affinity by using the following command:

set security nat source port-randomization disable

set security nat source port-round-robin disable

Related Documentation

- Understanding Central Point Architecture Enhancements for NAT on page 4880

Understanding Source NAT Rules

Source NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify combinations of **from interface**, **from zone**, or **from routing-instance** and **to interface**, **to zone**, or **to routing-instance**. You cannot configure the same **from** and **to** contexts for different rule sets.
- Packet information—Can be source and destination IP addresses or subnets, source port numbers or port ranges, destination port numbers or port ranges, protocols, or applications.

For all ALG traffic, except FTP, we recommend that you not use the **source-port** rule option. Data session creation can fail if this option is used because the IP address and the source port value, which is a random value, might not match the rule.

In addition, we recommend that you not use the **destination-port** option or the **application** option as matching conditions for ALG traffic. If these options are used, translation may fail because the port value in the application payload might not match the port value in the IP address.

If multiple source NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 to zone 2 and rule B specifies traffic from zone 1 to interface ge-0/0/0, rule B is used to perform source NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match. For more information about rule set matching, see "Understanding NAT Rule Sets and Rules" on page 4840.

The actions you can specify for a source NAT rule are:

- off—Do not perform source NAT.

Copyright © 2016, Juniper Networks, Inc.

- pool—Use the specified user-defined address pool to perform source NAT.
- interface—Use the egress interface's IP address to perform source NAT.

Source NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Source NAT rules are processed after static NAT rules, destination NAT rules, and reverse mapping of static NAT rules and after route and security policy lookup.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding NAT Rule Sets and Rules on page 4840

Example: Configuring Source NAT with Multiple Rules

This example describes how to configure source NAT mappings with multiple rules.

- Requirements on page 4882
- Overview on page 4882
- Configuration on page 4884
- Verification on page 4888

Requirements

Before you begin:

- Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
- Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 206, the following translations are performed on the Juniper Networks security device for the source NAT mapping for traffic from the trust zone to the untrust zones:

- The source IP address in packets sent by the 10.1.1.0/24 and 10.1.2.0/24 subnets to any address in the untrust zone is translated to a public address in the range from 192.0.2.1 to 192.0.2.24 with port translation.
- The source IP address in packets sent by the 192.168.1.0/24 subnet to any address in the untrust zone is translated to a public address in the range from 192.0.2.100 to 192.0.2.249 with no port translation.
- The source IP address in packets sent by the 192.168.1.250/32 host device is not translated.

Copyright © 2016, Juniper Networks, Inc.

Figure 206: Source NAT with Multiple Translation Rules

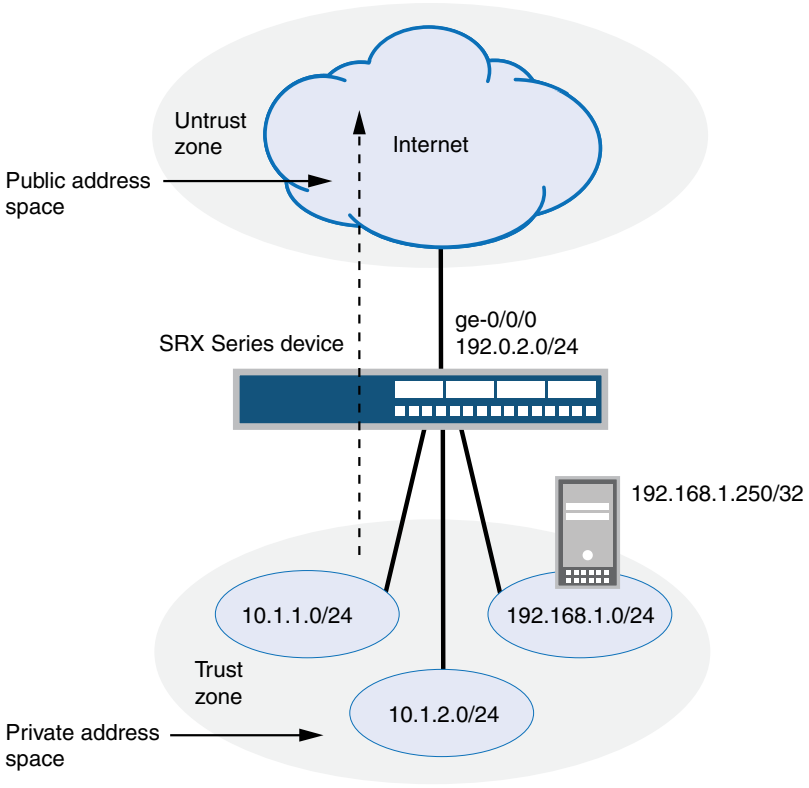

Original Source IP	Translated Source IP
10.1.1.0/24, 10.1.2.0/24	192.0.2.1 – 192.0.2.24 (w/port translation)
192.168.1.0/24	192.0.2.100 - 192.0.2.249 (no port translation)
192.168.1.250/32	(no source NAT translation)

g030673

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address range 192.0.2.1 through 192.0.2.24.
- Source NAT pool **src-nat-pool-2** that contains the IP address range 192.0.2.100 through 192.0.2.249, with port address translation disabled.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When port address translation is disabled, the number of translations that the source NAT pool can support concurrently is limited to the number of addresses in the pool, unless the **address-shared** option is enabled. Packets are dropped if there are no addresses available in the source NAT pool. You can optional specify an overflow pool from which IP addresses and port numbers are allocated when there are no addresses available in the original source NAT pool.

- Source NAT rule set **rs1** to match packets from the trust zone to the untrust zone. Rule set **rs1** contains multiple rules:
  - Rule **r1** to match packets with a source IP address in either the 10.1.1.0/24 or 10.1.2.0/24 subnets. For matching packets, the source address is translated to an IP address in the **src-nat-pool-1** pool.
  - Rule **r2** to match packets with a source IP address of 192.168.1.250/32. For matching packets, there is no NAT translation performed.
  - Rule **r3** to match packets with a source IP address in the 192.168.1.0/24 subnet. For matching packets, the source address is translated to an IP address in the **src-nat-pool-2** pool.

    NOTE: The order of rules in a rule set is important, as the first rule in the rule set that matches the traffic is used. Therefore, rule r2 to match a specific IP address must be placed before rule r3 that matches the subnet on which the device is located.

- Proxy ARP for the addresses 192.0.2.1 through 192.0.2.24 and 192.0.2.100 through 192.0.2.249 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 192.0.2.1/32 to 192.0.2.24/32
set security nat source pool src-nat-pool-2 address 192.0.2.100/32 to 192.0.2.249/32
set security nat source pool src-nat-pool-2 port no-translation
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 10.1.1.0/24
set security nat source rule-set rs1 rule r1 match source-address 10.1.2.0/24
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat source rule-set rs1 rule r2 match source-address 192.168.1.250/32
```


Copyright © 2016, Juniper Networks, Inc.

```
set security nat source rule-set rs1 rule r2 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r2 then source-nat off
set security nat source rule-set rs1 rule r3 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r3 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r3 then source-nat pool src-nat-pool-2
set security nat proxy-arp interface ge-0/0/0.0 address 192.0.2.1/32 to 192.0.2.24/32
set security nat proxy-arp interface ge-0/0/0.0 address 192.0.2.100/32 to 192.0.2.249/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure multiple source NAT rules in a rule set:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 192.0.2.1 to 192.0.2.24
   ```

2. Create a source NAT pool with no port translation.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-2 address 192.0.2.100 to 192.0.2.249
   user@host# set pool src-nat-pool-2 port no-translation
   ```

   NOTE: To configure an overflow pool for src-nat-pool-2 using the egress interface:

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-2 overflow-pool interface
   ```

3. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address [10.1.1.0/24 10.1.2.0/24]
   user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```

5. Configure a rule to match packets for which the source address is not translated.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r2 match source-address 192.168.1.250/32
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set rule-set rs1 rule r2 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r2 then source-nat off
```

6. Configure a rule to match packets and translate the source address to an address in the pool with no port translation.

```
[edit security nat source]
user@host# set rule-set rs1 rule r3 match source-address 192.168.1.0/24
user@host# set rule-set rs1 rule r3 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r3 then source-nat pool src-nat-pool-2
```

7. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 192.0.2.1 to 192.0.2.24
user@host# set proxy-arp interface ge-0/0/0.0 address 192.0.2.100 to 192.0.2.249
```

8. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any
  destination-address any application any
user@host# set policy internet-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      192.0.2.1/32 to 192.0.2.24/32;
    }
  }
  pool src-nat-pool-2 {
    address {
      192.0.2.100/32 to 192.0.2.249/32;
    }
    port no-translation;
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address [ 10.1.1.0/24 10.1.2.0/24 ];
        destination-address 0.0.0.0/0;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        rule r2 {
          match {
            source-address 192.168.1.250/32;
            destination-address 0.0.0.0/0;
          }
          then {
            source-nat {
              off;
            }
          }
        }
        rule r3 {
          match {
            source-address 192.168.1.0/24;
            destination-address 0.0.0.0/0;
          }
          then {
            source-nat {
              pool {
                src-nat-pool-2;
              }
            }
          }
        }
      }
    }
      proxy-arp {
        interface ge-0/0/0.0 {
          address {
            192.0.2.1/32 to 192.0.2.24/32;
            192.0.2.100/32 to 192.0.2.249/32;
          }
        }
      }
    user@host# show security policies
    from-zone trust to-zone untrust {
      policy internet-access {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4888
- Verifying Source NAT Rule Usage on page 4888
- Verifying NAT Application to Traffic on page 4888

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Rules on page 4881

Disabling Port Randomization for Source NAT (CLI Procedure)

For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT. When port randomization is disabled, ports are allocated, by default, on a round-robin basis: NAT first selects by IP address, then selects by port configuration. For example, if the source pool contains only one IP, when the first packet of a flow arrives, creating a session, it is translated to port N of IP1. Subsequent packets in that flow are allocated the same IP/port. When the first packet or a new flow arrives, it is translated to port N+1 of IP1, and so on. If the source pool contains two IPs, when the first packet of a flow arrives, creating a session, it is translated to port X of IP1. Subsequent packets in that flow are allocated the same IP/port. When the first packet of a new flow

Copyright © 2016, Juniper Networks, Inc.

arrives, it is translated to port Y of IP2; and when another packet in a new flow arrives, it is translated to port Y+1 of IP2, and so on.

You can disable port randomization by using the **port-randomization disable** statement at the [**edit security nat source**] hierarchy level. To re-enable port randomization, delete the **port-randomization** statement at the [**edit security nat source**] hierarchy level.

```
user@host# set security nat source port-randomization disable
```

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864

Monitoring Source NAT Information

Purpose Display configured information about source Network Address Translation (NAT) rules, pools, persistent NAT, and paired addresses.

Action Select **Monitor>NAT>Source NAT** in the J-Web user interface, or enter the following CLI commands:

- **show security nat source summary**
- **show security nat source pool** ***pool-name***
- **show security nat source persistent-nat-table**
- **show security nat source paired-address**

Table 440 describes the available options for monitoring source NAT.

Table 440: Source NAT Monitoring Page

Field	Description	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Name	Name of the rule .	–
From	Name of the routing instance/zone/interface from which the packet flows.	–
To	Name of the routing instance/zone/interface to which the packet flows.	–
Source address range	Source IP address range in the source pool.	–

Copyright © 2016, Juniper Networks, Inc.

Table 440: Source NAT Monitoring Page *(continued)*

Field	Description	Action
Destination address range	Destination IP address range in the source pool.	–
Source ports	Source port numbers.	–
Ip protocol	IP protocol.	–
Action	Action taken for a packet that matches a rule.	–
Persistent NAT type	Persistent NAT type.	–
Inactivity timeout	Inactivity timeout interval for the persistent NAT binding.	–
Alarm threshold	Utilization alarm threshold.	
Max session number	The maximum number of sessions.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule.	–
Translation Hits	Number of times a translation in the translation table is used for a source NAT rule.	–
Pools		
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	–
ID	ID of the pool.	–
Name	Name of the source pool.	–
Address range	IP address range in the source pool.	–

Copyright © 2016, Juniper Networks, Inc.

Table 440: Source NAT Monitoring Page *(continued)*

Field	Description	Action
Single/Twin ports	Number of allocated single and twin ports.	–
Port	Source port number in the pool.	–
Address assignment	Displays the type of address assignment.	–
Alarm threshold	Utilization alarm threshold.	–
Port overloading factor	Port overloading capacity.	–
Routing instance	Name of the routing instance.	–
Total addresses	Total IP address, IP address set, or address book entry.	–
Host address base	Host base address of the original source IP address range.	–
Translation hits	Number of times a translation in the translation table is used for source NAT.	–
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	–
Persistent NAT		
Persistent NAT table statistics		
binding total	Displays the total number of persistent NAT bindings for the FPC.	–
binding in use	Number of persistent NAT bindings that are in use for the FPC.	–
enode total	Total number of persistent NAT enodes for the FPC.	–
enode in use	Number of persistent NAT enodes that are in use for the FPC.	–
Persistent NAT table		
Source NAT pool	Name of the pool.	Select all pools or a specific pool to display from the list.

Copyright © 2016, Juniper Networks, Inc.

Table 440: Source NAT Monitoring Page *(continued)*

Field	Description	Action
Internal IP	Internal IP address.	Select all IP addresses or a specific IP address to display from the list.
Internal port	Displays the internal ports configured in the system.	Select the port to display from the list.
Internal protocol	Internal protocols .	Select all protocols or a specific protocol to display from the list.
Internal IP	Internal transport IP address of the outgoing session from internal to external.	–
Internal port	Internal transport port number of the outgoing session from internal to external.	–
Internal protocol	Internal protocol of the outgoing session from internal to external.	–
Reflective IP	Translated IP address of the source IP address.	–
Reflective port	Displays the translated number of the port.	–
Reflective protocol	Translated protocol.	–
Source NAT pool	Name of the source NAT pool where persistent NAT is used.	–
Type	Persistent NAT type.	–
Left time/Conf time	Inactivity timeout period that remains and the configured timeout value.	–
Current session num/Max session num	Number of current sessions associated with the persistent NAT binding and the maximum number of sessions.	–
Source NAT rule	Name of the source NAT rule to which this persistent NAT binding applies.	–
External node table		
Internal IP	Internal transport IP address of the outgoing session from internal to external.	–

Copyright © 2016, Juniper Networks, Inc.

Table 440: Source NAT Monitoring Page *(continued)*

Field	Description	Action
Internal port	Internal port number of the outgoing session from internal to external.	–
External IP	External IP address of the outgoing session from internal to external.	–
External port	External port of the outgoing session from internal to external.	–
Zone	External zone of the outgoing session from internal to external.	–
Paired Address		
Pool name	Name of the pool.	Select all pools or a specific pool to display from the list.
Specified Address	IP address.	Select all addresses, or select the internal or external IP address to display, and enter the IP address.
Pool name	Displays the selected pool or pools.	–
Internal address	Displays the internal IP address.	–
External address	Displays the external IP address.	–
Resource Usage		
Utilization for all source pools		
Pool name	Name of the pool.	To view additional usage information for Port Address Translation (PAT) pools, select a pool name. The information displays under Detail Port Utilization for Specified Pool.
Pool type	Pool type: PAT or Non-PAT.	–
Port overloading factor	Port overloading capacity for PAT pools.	–
Address	Addresses in the pool.	–
Used	Number of used resources in the pool. For Non-PAT pools, the number of used IP addresses is displayed. For PAT pools, the number of used ports is displayed.	–

Copyright © 2016, Juniper Networks, Inc.

Table 440: Source NAT Monitoring Page *(continued)*

Field	Description	Action
Available	Number of available resources in the pool. For Non-PAT pools, the number of available IP addresses is displayed. For PAT pools, the number of available ports is displayed.	–
Total	Number of used and available resources in the pool. For Non-PAT pools, the total number of used and available IP addresses is displayed. For PAT pools, the total number of used and available ports is displayed.	–
Usage	Percent of resources used. For Non-PAT pools, the percent of IP addresses used is displayed. For PAT pools, the percent of ports, including single and twin ports, is displayed.	–
Peak usage	Percent of resources used during the peak date and time.	–
Detail Port Utilization for Specified Pool		
Address Name	IP addresses in the PAT pool.	Select the IP address for which you want to display detailed usage information.
Factor-Index	Index number.	–
Port-range	Displays the number of ports allocated at a time.	–
Used	Displays the number of used ports.	–
Available	Displays the number of available ports.	–
Total	Displays the number of used and available ports.	–
Usage	Displays the percentage of ports used during the peak date and time.	–

Related Documentation

- Monitoring Destination NAT Information on page 4944
- Monitoring Static NAT Information on page 4965
- Monitoring Incoming Table Information on page 4859
- Monitoring Interface NAT Port Information on page 4860

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 201

Configuring Source NAT Pools

- Understanding Source NAT Pools on page 4895
- Understanding Source NAT Pool Capacities on page 4897
- Understanding Persistent Addresses on page 4898
- Example: Configuring Capacity for Source NAT Pools with PAT on page 4898
- Understanding Source NAT Pools with Address Pooling on page 4900
- Understanding Source NAT Pools with Address Shifting on page 4901
- Example: Configuring Source NAT with Address Shifting on page 4901
- Understanding Source NAT Pools with PAT on page 4906
- Example: Configuring Source NAT for Multiple Addresses with PAT on page 4907
- Understanding Source NAT Pools Without PAT on page 4912
- Example: Configuring a Single IP Address in a Source NAT Pool Without PAT on page 4913
- Example: Configuring Source NAT for Multiple Addresses Without PAT on page 4917
- Understanding Source NAT Pools with Shared Address on page 4921

Understanding Source NAT Pools

A NAT pool is a user-defined set of IP addresses that are used for translation. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with source NAT, you translate the original source IP address to an IP address in the address pool.

For source Network Address Translation (NAT) address pools, specify the following:

- Name of the source NAT address pool.
- Up to eight address or address ranges.

NOTE: **Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.**

- Routing instance—Routing instance to which the pool belongs (the default is the main **inet.0** routing instance).

Copyright © 2016, Juniper Networks, Inc.

- Port —The Port Address Translation (PAT) for a source pool. By default, PAT is performed with source NAT. If you specify the **no-translation** option, the number of hosts that the source NAT pool can support is limited to the number of addresses in the pool. If you specify **block-allocation**, a block of ports is allocated for translation, instead of individual ports being allocated. If you specify **deterministic**, an incoming (source) IP address and port always map to the specific destination address and port block, based on predefined, deterministic NAT algorithm. If you specify **port-overloading**, you can configure the port overloading capacity in source NAT. If you specify **range**, you can provide the port number range attached to each address in the pool, and the twin port range for source NAT pools.
- Overflow pool (optional)—Packets are dropped if there are no addresses available in the designated source NAT pool. To prevent that from happening when the **port no-translation** option is configured, you can specify an overflow pool. Once addresses from the original source NAT pool are exhausted, IP addresses and port numbers are allocated from the overflow pool. A user-defined source NAT pool or an egress interface can be used as the overflow pool. (When the overflow pool is used, the pool ID is returned with the address.)
- IP address shifting (optional)—A range of original source IP addresses can be mapped to another range of IP addresses, or to a single IP address, by shifting the IP addresses. Specify the **host-address-base** option with the base address of the original source IP address range.
- Address sharing (optional)—Multiple internal IP addresses can be mapped to the same external IP address. This option can be used only when the source NAT pool is configured with no port translation. Specify the **address-shared** option when a source NAT pool has few external IP addresses available, or only one external IP address. With a many-to-one mapping, use of this option increases NAT resources and improves traffic.
- Address pooling (optional)— Address pooling can be configured as paired or no-paired. Specify **address-pooling paired** for applications that require all sessions associated with one internal IP address to be mapped to the same external IP address for the duration of a session. This differs from the **persistent-address** option, in which the same internal address is translated to the same external address every time. Specify **address-pooling no-paired** for applications that can be can be assigned IP addresses in a round-robin fashion. If either **address-pooling paired** or **address-pooling no-paired** is configured for a source NAT pool with PAT, the persistent address option is disabled. If **address-shared** is configured on a source NAT pool without PAT, then the **persistent-address** option is enabled. Both **address-shared** and **address-pooling paired** can be configured on the same source NAT pool without PAT.
- Pool utilization alarm (optional)— When the **raise-threshold** option is configured for source NAT, an SNMP trap is triggered if the source NAT pool utilization rises above this threshold. If the optional **clear-threshold** option is configured, an SNMP trap is triggered if the source NAT pool utilization drops below this threshold. If **clear-threshold** is not configured, it is set by default to 80 percent of the **raise-threshold** value.

Copyright © 2016, Juniper Networks, Inc.

You can use the **show security nat resource usage source pool** command to view address use in a source NAT pool without PAT, and to view port use in a source NAT pool with PAT.

Related Documentation

- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools with PAT on page 4906
- Understanding Source NAT Pools Without PAT on page 4912
- Understanding Source NAT Pools with Address Shifting on page 4901
- Understanding Persistent Addresses on page 4898
- Understanding Source NAT Pools with Shared Address on page 4921
- Understanding Source NAT Pools with Address Pooling on page 4900

Understanding Source NAT Pool Capacities

Maximum capacities for source pools and IP addresses on all branch SRX Series devices are as follows:

Pool/PAT Maximum Address Capacity	SRX300 SRX320	SRX340 SRX345
Source NAT pools	1024	2048
IP addresses supporting port translation	1024	2048
PAT port number	64M	64M

Maximum capacities for source pools and IP addresses on all high-end SRX Series devices are as follows:

Pool/PAT Maximum Address Capacity	SRX1500	SRX5400 SRX5600 SRX5800
Source NAT pools	8192	12,288
IP addresses supporting port translation	8192	1M
PAT port number	256M	384M

Increasing the capacity of source NAT pools consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT does not consume too much memory. Use the **port-range** statement in configuration mode in

Copyright © 2016, Juniper Networks, Inc.

the CLI to assign a new port range or the **pool-default-port-range** statement to override the specified default.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For source pool with PAT in range (63,488 through 65,535), two ports are allocated at one time for RTP/RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (63,488 through 65,535) for ALG module use.

Related Documentation

- Understanding Source NAT Pools on page 4895

Understanding Persistent Addresses

By default, port address translation is performed with source NAT. However, an original source address may not be translated to the same IP address for different traffic that originates from the same host. The source NAT **address-persistent** option ensures that the same IP address is assigned from the source NAT pool to a specific host for multiple concurrent sessions.

This option differs from the address-pooling paired option, where the internal address is mapped to an external address within the pool on a first-come, first-served basis, and might be mapped to a different external address for each session.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools with PAT on page 4906
- Understanding Source NAT Pools with Address Pooling on page 4900

Example: Configuring Capacity for Source NAT Pools with PAT

This example describes how to configure the capacity of source NAT pools with Port Address Translation (PAT) if a default port range is not set or you want to override it. Translations are set for each IP address. When the source pool is increased, ports should be reassigned if the current port number exceeds limitations.

- Requirements on page 4899
- Overview on page 4899
- Configuration on page 4899
- Verification on page 4900

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example shows how to configure a PAT pool of 2048 IP addresses with 32K ports for each IP address.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit security nat source]
set pool src-nat-pat-addr address 192.168.0.0/32 to 192.168.3.255/32
set pool src-nat-pat-addr address 192.168.4.0/32 to 192.168.7.255/32
set pool-default-port-range 2001
set pool-default-port-range to 32720
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure capacity for a source NAT pool with PAT:

1. Specify a source NAT pool with PAT and an IP address range.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pat-addr address 192.168.0.0/32 to 192.168.3.255/32
   user@host#set pool src-nat-pat-addr address 192.168.4.0/32 to 192.168.7.255/32
   ```

2. Specify a default port range for the source pool.

   ```
   [edit security nat source]
   user@host# set pool-default-port-range 2001
   user@host# set pool-default-port-range to 32720
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat-source-summary** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> run show security nat source summary
Total port number usage for port translation pool: 16515072
Maximum port number for port translation pool: 134217728
Total pools: 1
```

Copyright © 2016, Juniper Networks, Inc.

```
Pool Address Routing PAT Total Name Range Instance Address pool2
203.0.113.1-203.0.113.3 default yes 2048
Name Range Instance Address
pool1 198.51.100.0-198.51.100.255 default yes 256

Total rules: 1
Rule name Rule set From To Action
rule 1 ruleset1 ge-2/2/2.0 ge-2/2/3.0 pool1
rule 1 ge-2/2/4.0 ge-2/2/5.0
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Capacity of Source NAT Pools

Purpose View port and pool information. Port limitations are automatically checked, so the configuration will not be committed if port limitations are exceeded.

Action From operational mode, enter the **show security nat source summary** command to view port and pool details.

Related Documentation

- Understanding Source NAT on page 4863
- Understanding Source NAT Pools with PAT on page 4906
- Source NAT Configuration Overview on page 4864
- Example: Configuring Source NAT for Multiple Addresses with PAT on page 4907

Understanding Source NAT Pools with Address Pooling

When a host initiates several sessions that match a policy that requires NAT, and is assigned an IP address from a source pool that has port address translation enabled, a different source IP address is used for each session.

Because some applications require the same source IP address for each session, you can use the **address-pooling paired** feature to enable all sessions associated with one internal IP address to map to the same external IP address for the duration of the sessions. When the sessions end, the mapping between the internal IP address and the external IP address cease. The next time the host initiates a session, a different IP address from the pool might be assigned to it.

This differs from the source NAT **persistent-address** feature, which keeps the mapping static; the same internal IP address is mapped to the same external IP address every time. It also differs from the **persistent-address** feature in that **address-pooling paired** is configured for a specific pool. The **persistent-address** feature is a global configuration that applies to all source pools.

Related Documentation

- Understanding Persistent Addresses on page 4898

Copyright © 2016, Juniper Networks, Inc.

Understanding Source NAT Pools with Address Shifting

The match conditions for a source NAT rule set do not allow you to specify an address range; only address prefixes may be specified in a rule. When configuring a source NAT pool, you can specify the **host-base-address** option; this option specifies the IP address where the original source IP address range begins.

The range of original source IP addresses that are translated is determined by the number of addresses in the source NAT pool. For example, if the source NAT pool contains a range of ten IP addresses, then up to ten original source IP addresses can be translated, starting with a specified base address. This type of translation is one-to-one, static, and without port address translation.

The match condition in a source NAT rule may define a larger address range than that specified in the source NAT pool. For example, a match condition might specify an address prefix that contains 256 addresses, but the source NAT pool might contain a range of only a few IP addresses, or only one IP address. A packet's source IP address can match a source NAT rule, but if the source IP address is not within the address range specified in the source NAT pool, the source IP address is not translated.

Related Documentation

- Understanding Source NAT on page 4863
- Understanding Source NAT Pools on page 4895
- Example: Configuring Source NAT with Address Shifting on page 4901

Example: Configuring Source NAT with Address Shifting

This example describes how to configure a source NAT mapping of a private address range to public addresses, with optional address shifting. This mapping is one-to-one between the original source IP addresses and translated IP addresses and no port translation is performed.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The match conditions for a source NAT rule set do not allow you to specify an address range; only address prefixes may be specified in a rule. When configuring a source NAT pool, you can specify the **host-base-address** option; this option specifies the IP address where the original source IP address range begins.

The range of original source IP addresses that are translated is determined by the number of addresses in the source NAT pool. For example, if the source NAT pool contains a range of ten IP addresses, then up to ten original source IP addresses can be translated, starting with a specified base address.

The match condition in a source NAT rule may define a larger address range than that specified in the source NAT pool. For example, a match condition might specify an address prefix that contains 256 addresses, but the source NAT pool contains a range of only ten IP addresses. A packet's source IP address can match a source NAT rule, but if the source IP address is not within the address range specified in the source NAT pool, the source IP address is not translated.

- Requirements on page 4902
- Overview on page 4902
- Configuration on page 4904
- Verification on page 4906

Requirements

Before you begin:

- Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
- Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 207, a range of private addresses in the trust zone is mapped to a range of public addresses in the untrust zone. For packets sent from the trust zone to the untrust zone, a source IP address in the range of 192.168.1.10/32 through 192.168.1.20/32 is translated to a public address in the range of 203.0.113.30/32 through 203.0.113.40/32.

Copyright © 2016, Juniper Networks, Inc.

Figure 207: Source NAT with Address Shifting

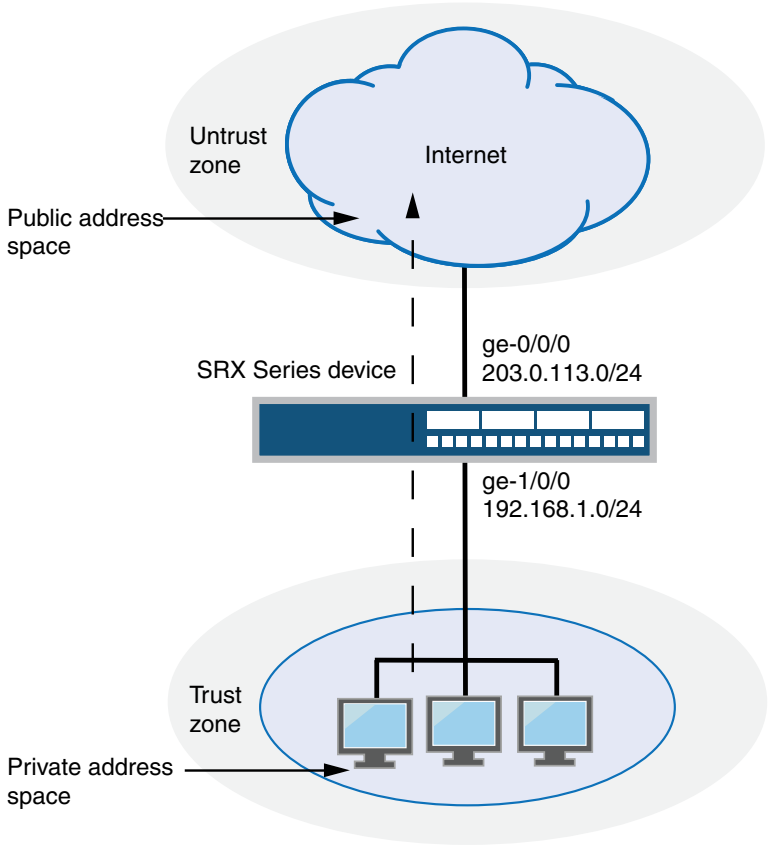

Original Source IP	Translated Source IP
192.168.1.10/32 - 192.168.1.20/32	203.0.113.30/32 - 203.0.113.40/32

g030672

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address range 203.0.113.30/32 through 203.0.113.40/32. For this pool, the beginning of the original source IP address range is 192.168.1.10/32 and is specified with the **host-address-base** option.
- Source NAT rule set **rs1** with rule **r1** to match packets from the trust zone to the untrust zone with a source IP address in the 192.168.1.0/24 subnet. For matching packets that fall within the source IP address range specified by the **src-nat-pool-1** configuration, the source address is translated to the IP address in **src-nat-pool-1** pool.

Copyright © 2016, Juniper Networks, Inc.

- Proxy ARP for the addresses 203.0.113.30/32 through 203.0.113.40/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.30/32 to 203.0.113.40/32
set security nat source pool src-nat-pool-1 host-address-base 192.168.1.10/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.30/32 to
  203.0.113.40/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping with address shifting:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.30/32 to 203.0.113.40/32
   ```

2. Specify the beginning of the original source IP address range.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 host-address-base 192.168.1.10/32
   ```

3. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 192.168.1.0/24
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```


Copyright © 2016, Juniper Networks, Inc.

5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.30/32 to
  203.0.113.40/32
```

6. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any
  destination-address any application any
user@host# set policy internet-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.30/32 to 203.0.113.40/32;
    }
    host-address-base 192.168.1.10/32;
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address 192.168.1.0/24;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/0.0 {
      address {
        203.0.113.30/32 to 203.0.113.40/32;
      }
    }
  }
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        application any;
      }
      then {
        permit;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4906
- Verifying Source NAT Rule Usage on page 4906
- Verifying NAT Application to Traffic on page 4906

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools with Address Shifting on page 4901

Understanding Source NAT Pools with PAT

Using the source pool with Port Address Translation (PAT), Junos OS translates both the source IP address and the port number of the packets. When PAT is used, multiple hosts can share the same IP address.

Copyright © 2016, Juniper Networks, Inc.

Junos OS maintains a list of assigned port numbers to distinguish what session belongs to which host. When PAT is enabled, up to 63,488 hosts can share a single IP address. Each source pool can contain multiple IP addresses, multiple IP address ranges, or both. For a source pool with PAT, Junos OS may assign different addresses to a single host for different concurrent sessions, unless the source pool or Junos OS has the persistent address feature or the paired address pooling feature enabled.

For interface source pool and source pool with PAT, range (1024, 65535) is available for port number mapping per IP address. Within range (1024, 63487) one port is allocated at a time, for a total of 62,464 ports. In range (63488, 65535), two ports are allocated at a time for RTP/RTCP applications such as SIP, H.323, and RTSP, for a total of 2,048 ports.

When a host initiates several sessions that match a policy that requires network address translation and is assigned an address from a source pool that has PAT enabled, the device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session. For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Message (AIM) client.

To ensure that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions, you can enable a persistent IP address per router. To ensure that the device assigns the same IP address from a source pool to a host for the duration of a single session, you can enable paired address pooling.

Related Documentation

- Understanding Source NAT on page 4863
- Understanding Source NAT Pools on page 4895
- Understanding Source NAT Pools Without PAT on page 4912
- Understanding Source NAT Pools with Address Shifting on page 4901
- Understanding Source NAT Pools with Address Pooling on page 4900
- Understanding Persistent Addresses on page 4898

Example: Configuring Source NAT for Multiple Addresses with PAT

This example describes how to configure a source NAT mapping of a private address block to a smaller public address block using port address translation.

- Requirements on page 4908
- Overview on page 4908
- Configuration on page 4910
- Verification on page 4912

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 208, the source IP address in packets sent from the trust zone to the untrust zone is mapped to a smaller block of public addresses in the range from 203.0.113.1/32 through 203.0.113.24/32. Because the size of the source NAT address pool is smaller than the number of potential addresses that might need to be translated, port address translation is used.

NOTE: Port address translation includes a source port number with the source IP address mapping. This allows multiple addresses on a private network to map to a smaller number of public IP addresses. Port address translation is enabled by default for source NAT pools.

Copyright © 2016, Juniper Networks, Inc.

Figure 208: Source NAT Multiple Addresses with PAT

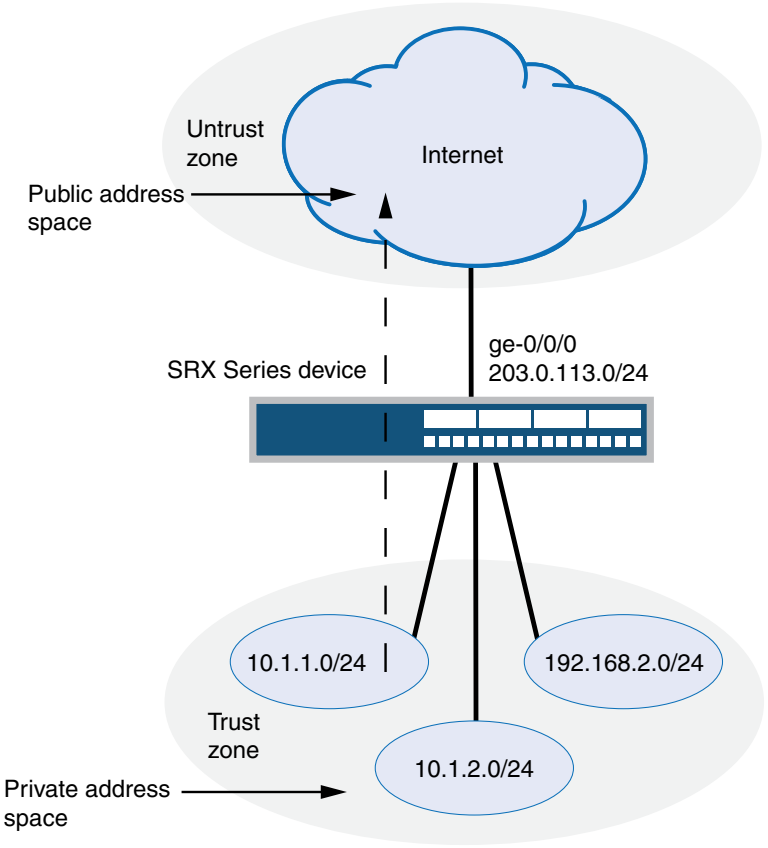

Original Source IP	Translated Source IP
10.1.1.0/24 10.1.2.0/24 192.168.1.0/24	203.0.113.1 (with port address translation)

g030670

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address range 203.0.113.1/32 through 203.0.113.24/32.
- Source NAT rule set **rs1** to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the **src-nat-pool-1** pool.

Copyright © 2016, Juniper Networks, Inc.

- Proxy ARP for the addresses 203.0.113.1/32 through 203.0.113.24/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/32 to 203.0.113.24/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 10.1.1.0/24
set security nat source rule-set rs1 rule r1 match source-address 10.1.2.0/24
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.24/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a smaller public address block using PAT:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.1 to 203.0.113.24
   ```

2. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

3. Configure a rule that matches packets and translates the source address to an address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address [10.1.1.0/24 10.1.2.0/24
     192.168.1.0/24]
   user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```


Copyright © 2016, Juniper Networks, Inc.

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1 to 203.0.113.24
```

5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any
  destination-address any application any
user@host# set policy internet-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.1/32 to 203.0.113.24/32;
    }
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address [10.1.1.0/24 10.1.2.0/24 192.168.1.0/24];
        destination-address 0.0.0.0/0;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
proxy-arp {
  interface ge-0/0/0.0 {
    address {
      203.0.113.1/32 to 203.0.113.24/32;
    }
  }
}
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          permit;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4912
- Verifying Source NAT Rule Usage on page 4912
- Verifying NAT Application to Traffic on page 4912

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools with PAT on page 4906

Understanding Source NAT Pools Without PAT

When you define a source pool, Junos OS enables PAT by default. To disable PAT, you must specify no port translation when you are defining a source pool.

When using a source pool without PAT, Junos OS performs source Network Address Translation for the IP address without performing PAT for the source port number. For

Copyright © 2016, Juniper Networks, Inc.

applications that require that a particular source port number remain fixed, you must use source pool without PAT.

The source pool can contain multiple IP addresses, multiple IP address ranges, or both. For source pool without PAT, Junos OS assigns one translated source address to the same host for all its concurrent sessions unless the address-pooling no-paired option is enabled.

The number of hosts that a source NAT pool without PAT can support is limited to the number of addresses in the pool. When you have a pool with a single IP address, only one host can be supported, and traffic from other hosts is blocked because there are no resources available. If a single IP address is configured for a source NAT pool without PAT when NAT resource assignment is not in active-backup mode in a chassis cluster, traffic through node 1 will be blocked.

Pool utilization for each source pool without PAT is computed. You can turn on pool utilization alarm by configuring alarm thresholds. An SNMP trap is triggered every time pool utilization rises above a threshold and goes below a threshold.

NOTE: If a static NAT rule is for one-to-one IP translation, avoid dividing the rule into a destination rule and a source rule when source no-pat pool without address sharing is used. If you choose to divide the rule, you will then have to use source pat-pool with single IP or source no-pat pool with multiple IP.

Related Documentation

- Understanding Source NAT on page 4863
- Understanding Source NAT Pools on page 4895
- Understanding Source NAT Pools with PAT on page 4906
- Understanding Source NAT Pools with Address Shifting on page 4901
- Understanding Persistent Addresses on page 4898
- Understanding Source NAT Pools with Address Pooling on page 4900

Example: Configuring a Single IP Address in a Source NAT Pool Without PAT

This example describes how to configure a private address block to a single public address in a source NAT pool without Port Address Translation.

Copyright © 2016, Juniper Networks, Inc.

NOTE: PAT is enabled by default for source NAT pools. When PAT is disabled, the number of translations that the source NAT pool can concurrently support is limited to the number of addresses in the pool. Packets are dropped if there are no addresses available in the source NAT pool. However, using the **address-shared** option, you can map more that one private IP address to a single public IP address as long as the traffic is from different source ports.

- Requirements on page 4914
- Overview on page 4914
- Configuration on page 4914
- Verification on page 4916

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. The source IP address of packets sent from the trust zone to the untrust zone are mapped to a single public address.

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address 203.0.113.1/30. The **port no-translation** option and the **address shared** option are specified for the pool.
- Source NAT rule set **rs1** to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the **src-nat-pool-1** pool.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/30
set security nat source pool src-nat-pool-1 port no-translation
set security nat source pool-src-nat-pool-1 address-shared
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
```


Copyright © 2016, Juniper Networks, Inc.

```
set security nat source rule-set rs1 rule1 match source address 192.0.2.0/24
set security nat source rule-set rs1 rule r1 then source src-nat-pool-1
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a single public address without PAT:

1. Create a source NAT pool with a single IP address for the shared address.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.1/30
   ```

   Specify the **port no-translation** option.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 port no-translation
   ```

2. Specify the **address-shared** option.

   ```
   [edit security nat source]
   user@host# set pool pool-src-nat-pool-1 address-shared
   ```

3. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 192.0.2.0/24
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```

5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy internet-access match source-address any
     destination-address any application any
   user@host# set policy internet-access then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat source pool** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.1/30
    }
    port no-translation;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    address-shared;
    rule-set rs1 {
      from zone trust;
      to zone untrust;
      rule r1 {
        match {
          source-address [192.0.2.0/24]
        }
        then {
          source-nat {
            pool {
              src-nat-pool-1;
            }
          }
        }
      }
    }
  }
}
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Shared Address on page 4916
- Verifying Shared Address Application to Traffic on page 4916

Verifying Shared Address

Purpose Verify that two internal IP addresses, with different source ports, share one external IP address.

Action From operational mode, enter the **show security nat source pool** command. View the **Address assignment** field to verify that it is shared.

Verifying Shared Address Application to Traffic

Purpose Verify that two sessions are using the same IP address.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools Without PAT on page 4912
- Understanding Source NAT Pools with Shared Address on page 4921

Example: Configuring Source NAT for Multiple Addresses Without PAT

This example describes how to configure a source NAT mapping of a private address block to a smaller public address block without port address translation.

NOTE: Port address translation is enabled by default for source NAT pools. When port address translation is disabled, the number of translations that the source NAT pool can concurrently support is limited to the number of addresses in the pool. Packets are dropped if there are no addresses available in the source NAT pool. You can optionally specify an overflow pool from which IP addresses and port numbers are allocated when there are no addresses available in the original source NAT pool.

- Requirements on page 4917
- Overview on page 4917
- Configuration on page 4919
- Verification on page 4921

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 209, the source IP address in packets sent from the trust zone to the untrust zone is mapped to a smaller block of public addresses in the range from 203.0.113.1/32 through 203.0.113.24/32.

Copyright © 2016, Juniper Networks, Inc.

Figure 209: Source NAT Multiple Addresses Without PAT

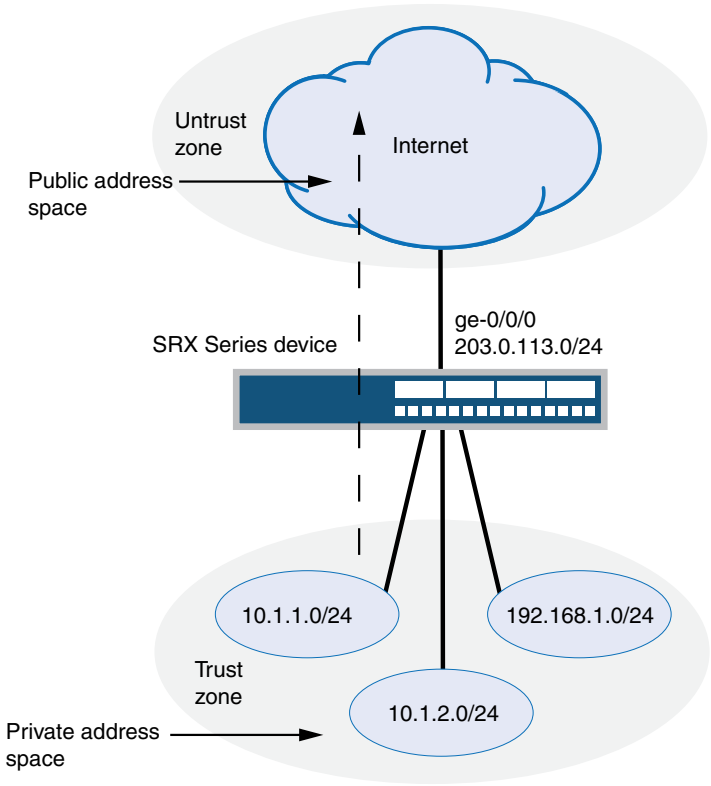

Original Source IP	Translated Source IP
10.1.1.0/24 10.1.2.0/24 192.168.1.0/24	203.0.113.1 (no port address translation)

g030671

This example describes the following configurations:

- Source NAT pool **src-nat-pool-1** that contains the IP address range 203.0.113.1/32 through 203.0.113.24/32. The **port no-translation** option is specified for the pool.
- Source NAT rule set **rs1** to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the **src-nat-pool-1** pool.
- Proxy ARP for the addresses 203.0.113.1/32 through 203.0.113.24/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/32 to 203.0.113.24/32
set security nat source pool src-nat-pool-1 port no-translation
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.24/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a smaller public address block without PAT:

1. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 address 203.0.113.1 to 203.0.113.24
   ```

2. Specify the **port no-translation** option.

   ```
   [edit security nat source]
   user@host# set pool src-nat-pool-1 port no-translation
   ```

3. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 from zone trust
   user@host# set rule-set rs1 to zone untrust
   ```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

   ```
   [edit security nat source]
   user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
   user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
   ```

5. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1 to 203.0.113.24
   ```

Copyright © 2016, Juniper Networks, Inc.

6. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any
  destination-address any application any
user@host# set policy internet-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.1/32 to 203.0.113.24/32;
    }
    port no-translation;
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address 0.0.0.0/0;
        destination-address 0.0.0.0/0;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/0.0 {
      address {
        203.0.113.1/32 to 203.0.113.24/32;
      }
    }
  }
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Source NAT Pool Usage on page 4921
- Verifying Source NAT Rule Usage on page 4921
- Verifying NAT Application to Traffic on page 4921

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Source NAT on page 4863
- Source NAT Configuration Overview on page 4864
- Understanding Source NAT Pools Without PAT on page 4912

Understanding Source NAT Pools with Shared Address

Source NAT pools with no port address translation perform static, one-to-one mappings from one source IP address to one external IP address. When there is only one external IP address, or very few available in a source no-pat pool or IP address shifting pool, the **address-shared** option enables you to map many source IP addresses to one external IP address, which increases NAT resources and improves traffic.

Copyright © 2016, Juniper Networks, Inc.

This option cannot be used on source NAT pools with port address translation because address sharing is already their default behavior.

Related Documentation

- Understanding Source NAT Pools Without PAT on page 4912
- Understanding Source NAT Pools with Address Shifting on page 4901

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 202

Configuring Destination NAT

- Understanding Destination NAT on page 4923
- Understanding Destination NAT Address Pools on page 4924
- Understanding Destination NAT Rules on page 4925
- Destination NAT Configuration Overview on page 4926
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939
- Monitoring Destination NAT Information on page 4944

Understanding Destination NAT

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

NOTE: **When destination NAT is performed, the destination IP address is translated according to configured destination NAT rules and then security policies are applied.**

Destination NAT allows connections to be initiated only for incoming network connections—for example, from the Internet to a private network. Destination NAT is commonly used to perform the following actions:

- Translate a single IP address to another address (for example, to allow a device on the Internet to connect to a host on a private network).
- Translate a contiguous block of addresses to another block of addresses of the same size (for example, to allow access to a group of servers).
- Translate a destination IP address and port to another destination IP address and port (for example, to allow access to multiple services using the same IP address but different ports).

Copyright © 2016, Juniper Networks, Inc.

The following types of destination NAT are supported:

- Translation of the original destination IP address to an IP address from a user-defined pool. This type of translation does not include Port Address Translation (PAT). If the original destination IP address range is larger than the address range in the user-defined address pool, any untranslated packets are dropped.
- Translation of the original destination IP address (and optional port number) to one specific IP address (and port number) from a user-defined pool.

Related Documentation

- Destination NAT Configuration Overview on page 4926
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939
- Understanding Destination NAT Address Pools on page 4924
- Understanding Destination NAT Rules on page 4925

Understanding Destination NAT Address Pools

A NAT pool is a user-defined set of IP addresses that are used for translation. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with destination NAT, you translate the original destination address to an IP address in the address pool.

For destination NAT address pools, specify the following:

- Name of the destination NAT address pool
- Destination address or address range

NOTE: Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.

- Destination port that is used for port forwarding
- Routing instance to which the pool belongs—A destination NAT pool that does not specify a specific routing instance will default to the routing instance of the ingress zone.

NOTE: You can configure a NAT pool to exist in the default routing instance. Configuration option to specify that a NAT pool exists in the default routing-instance is available. As a result, the NAT pool is reachable from zones in the default routing instance, and from zones in other routing instances.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Destination NAT on page 4923
- Destination NAT Configuration Overview on page 4926
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939

Understanding Destination NAT Rules

Destination NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify **from interface**, **from zone**, or **from routing-instance**.
- Packet information—Can be source IP addresses, destination IP address or subnet, destination port numbers or port ranges, protocols, or applications.

For ALG traffic, we recommend that you not use the **destination-port** option or the **application** option as matching conditions. If these options are used, translation may fail because the port value in the application payload might not match the port value in the IP address.

If multiple destination NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface **ge-0/0/0**, rule B is used to perform destination NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

The actions you can specify for a destination NAT rule are:

- off—Do not perform destination NAT.
- pool—Use the specified user-defined address pool to perform destination NAT.

Destination NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Destination NAT rules are processed after static NAT rules but before source NAT rules.

Related Documentation

- Destination NAT Configuration Overview on page 4926
- Understanding NAT Rule Sets and Rules on page 4840
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939

Copyright © 2016, Juniper Networks, Inc.

Destination NAT Configuration Overview

The main configuration tasks for destination NAT are as follows:

1. Configure a destination NAT address pool that aligns with your network and security requirements.
2. Configure destination NAT rules that align with your network and security requirements.
3. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Related Documentation

- Configuring Proxy ARP (CLI Procedure) on page 4857
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939
- Verifying NAT Configuration on page 4858

Example: Improving Security By Configuring Destination NAT for Single Address Translation

This example describes how to configure a destination NAT mapping of a single public address to a private address.

NOTE: Mapping one destination IP address to another can also be accomplished with static NAT. Static NAT mapping allows connections to be established from either side of the gateway device, whereas destination NAT only allows connections to be established from one side. However, static NAT only allows translations from one address to another or between blocks of addresses of the same size.

- Requirements on page 4926
- Overview on page 4927
- Configuration on page 4929
- Verification on page 4931

Requirements

This example uses the following hardware and software components:

- SRX Series device
- Server

Copyright © 2016, Juniper Networks, Inc.

Before you begin:

- Configure network interfaces on the device. See the *Interfaces Feature Guide for Security Devices*.
- Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

Destination NAT is commonly used to distribute a service located in a private network with a publicly accessible IP address. This allows users to use the private service with the public IP address. Destination NAT address pool and destination NAT rules configurations are used to align your network and improve security requirements.

In this example, first you configure the trust security zone for the private address space and then you configure the untrust security zone for the public address space. In Figure 210, devices in the untrust zone access a server in the trust zone by way of public address 203.0.113.200/32. For packets that enter the Juniper Networks security device from the untrust zone with the destination IP address 203.0.113.200/32, the destination IP address is translated to the private address 192.168.1.200/32.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 210: Destination NAT Single Address Translation

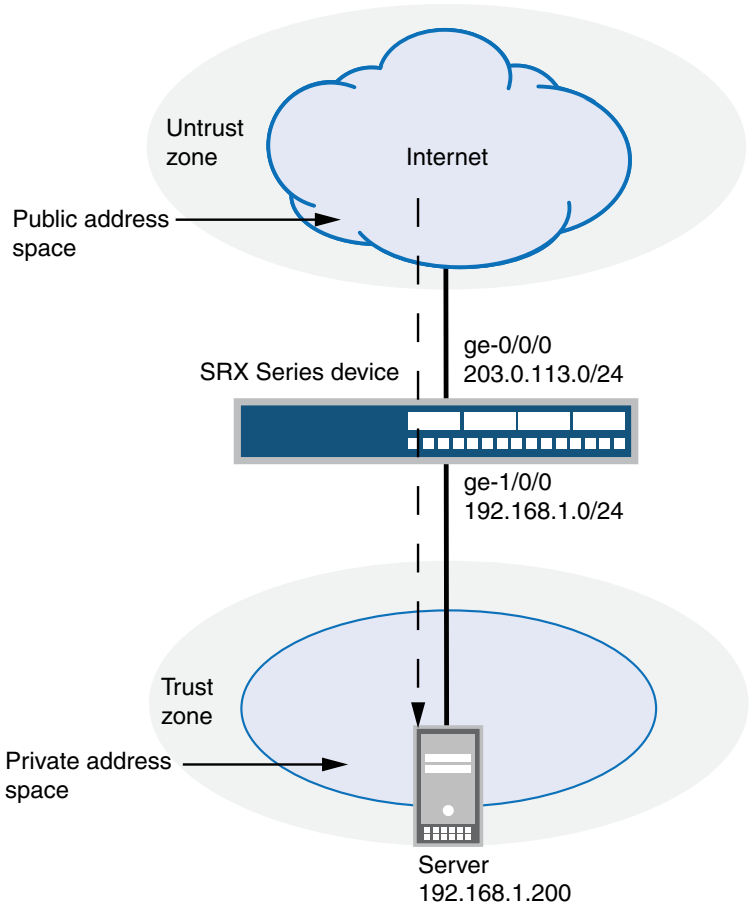

Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.200/32

g030665

Table 301 shows the parameters configured in this example.

Table 441: Interfaces, Zones, Server, and IP Address Information

Parameter	Description
Trust Zone	Security zone for the private address space.
Untrust Zone	Security zone for the public address space.

Copyright © 2016, Juniper Networks, Inc.

Table 441: Interfaces, Zones, Server, and IP Address Information *(continued)*

Parameter	Description
192.168.1.200/32	Translated destination NAT IP address.
192.168.1.0/24	Private subnet in private zone.
203.0.113.200/32	Public address of the server.
Server	Server address of the private address space.
ge-0/0/0 and ge-1/0/0	NAT interfaces for traffic direction.

This example describes the following configurations:

- Destination NAT pool **dst-nat-pool-1** that contains the IP address 192.168.1.200/32.
- Destination NAT rule set **rs1** with rule **r1** to match packets received from the ge-0/0/0.0 interface with the destination IP address 203.0.113.200/32. For matching packets, the destination address is translated to the address in the **dst-nat-pool-1** pool.
- Proxy ARP for the address 203.0.113.200/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the untrust zone to the translated destination IP address in the trust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.200/32
set security nat destination rule-set rs1 from interface ge-0/0/0.0
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone untrust to-zone trust policy server-access match
    source-address any
set security policies from-zone untrust to-zone trust policy server-access match
    destination-address server-1
set security policies from-zone untrust to-zone trust policy server-access match application
    any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a destination NAT mapping from a public address to a private address:

1. Create the destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set pool dst-nat-pool-1 address 192.168.1.200/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 from interface ge-0/0/0.0
   ```

3. Configure a rule that matches packets and translates the destination address to the address in the pool.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200/32
   user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
   ```

4. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
   ```

5. Configure an address in the global address book.

   ```
   [edit security address-book global]
   user@host# set address server-1 192.168.1.200/32
   ```

6. Configure a security policy that allows traffic from the untrust zone to the server in the trust zone.

   ```
   [edit security policies from-zone untrust to-zone trust]
   user@host# set policy server-access match source-address any
   user@host# set policy server-access match destination-address server-1
   user@host# set policy server-access match application any
   user@host# set policy server-access then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, and **show bridge-domains** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security nat
destination {
  pool dst-nat-pool-1 {
    address 192.168.1.200/32;
  }
  rule-set rs1 {
    from interface ge-0/0/0.0;
    rule r1 {
      match {
        destination-address 203.0.113.200/32;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          destination-nat pool dst-nat-pool-1;
        }
      }
    }
  }
    proxy-arp {
      interface ge-0/0/0.0 {
        address {
          203.0.113.200/32;
        }
      }
    }
[edit]
user@host# show security address-book
global {
  address server-1 192.168.1.200/32;
}
user@host# show security policies
from-zone untrust to-zone trust {
  policy server-access {
    match {
      source-address any;
      destination-address server-1;
      application any;
    }
    then {
      permit;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Destination NAT Pool Usage on page 4931
- Verifying Destination NAT Rule Usage on page 4932
- Verifying Destination NAT for a Single Address Translation on page 4932
- Verifying NAT Application to Traffic on page 4933

Verifying Destination NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the destination NAT pool.

Action From operational mode, enter the **show security nat destination pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

```
user@host>show security nat destination pool all
 Total destination-nat pools: 1

Pool name       : dst-nat-pool-1
```

```
Pool id         : 1
Total address   : 1
Translation hits: 71
Address range                        Port
    192.168.1.200 - 192.168.1.200           0
```

Meaning The **show security nat destination pool all** command displays the pool of translated addresses. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose Verify that there is traffic matching the destination NAT rule.

Action From operational mode, enter the **show security nat destination rule all** command.

```
user@host>show security nat destination rule all
Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0

Destination NAT rule: r1                  Rule-set: rs1
  Rule-Id                   : 1
  Rule position             : 1
  From interface            : ge-0/0/0.0
    Destination addresses   : 203.0.113.200 - 203.0.113.200
  Action                    : dst-nat-pool-1
  Translation hits          : 75
    Successful sessions     : 75
    Failed sessions         : 0
  Number of sessions        : 4
```

Meaning The **show security nat destination rule all** command displays the destination NAT rule. View the Translation hits field to check for traffic that matches the destination rule.

Verifying Destination NAT for a Single Address Translation

Purpose Verify the configuration of destination NAT for a single address translation.

Action From operational mode, enter the **show security nat destination summary** command.

```
user@host>show security nat destination summary
Total pools: 1
Pool name            Address  Range                Routing        Port  Total
                                                   Instance             Address
dst-nat-pool-1       192.168.1.200 - 192.168.1.200                  0     1

Total rules: 1
Rule name            Rule set       From                          Action
r1                  rs1           ge-0/0/0.0                    dst-nat-pool-1
```

Meaning The **show security nat destination summary** command displays information about destination NAT configuration. You can verify the following information:

- Rule sets
- Rules

Copyright © 2016, Juniper Networks, Inc.

- Address range
- NAT pool
- Port details

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session
Session ID: 26415, Policy name: server-access/11, Timeout: 2, Valid
  In: 203.0.113.219/30 --> 203.0.113.200/54850, If: ge-0/0/0.0, Pkts: 1, Bytes:
84
  Out: 192.168.1.200/54850 --> 203.0.113.219/30;icmp,If: ge-0/0/1.0, Pkts: 1,
Bytes: 84

Session ID: 26420, Policy name: server-access/11, Timeout: 2, Valid
  In: 203.0.113.219/31 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1,
Bytes: 84
  Out: 192.168.1.200/54850 --> 203.0.113.219/31;icmp, If: ge-0/0/1.0, Pkts: 1,
Bytes: 84

Session ID: 26425, Policy name: server-access/11, Timeout: 4, Valid
  In: 203.0.113.219/32 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1,
Bytes: 84
  Out: 192.168.1.200/54850 --> 203.0.113.219/32;icmp, If: ge-0/0/1.0, Pkts: 1,
Bytes: 84

Session ID: 26431, Policy name: server-access/11, Timeout: 4, Valid
  In: 203.0.113.219/33 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1,
Bytes: 84
  Out: 192.168.1.200/54850 --> 203.0.113.219/33;icmp, If: ge-0/0/1.0, Pkts: 1,
Bytes: 84
Total sessions: 9
```

Meaning The **show security flow session** command displays active sessions on the device and each session's associated security policy. The output shows traffic entering the device using the private source address 203.0.113.219/30 destined to a public host at 203.0.113.200. The return traffic from this flow travels to the translated public address 203.0.113.219.

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **server-access**—Policy name that permitted the traffic from the untrust zone to the translated destination IP address in the trust zone.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, the session is ICMP, and the source interface for this session is ge-0/0/0.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, the session is ICMP, and the destination interface for this session is ge-0/0/1.0).

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Destination NAT Configuration Overview on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934
- Example: Configuring Destination NAT for Subnet Translation on page 4939

Example: Configuring Destination NAT for IP Address and Port Translation

This example describes how to configure destination NAT mappings of a public address to private addresses, depending on the port number.

- Requirements on page 4934
- Overview on page 4934
- Configuration on page 4936
- Verification on page 4938

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 211, devices in the untrust zone access servers in the trust zone by way of public address 203.0.113.200 on port 80 or 8000. Packets entering the Juniper Networks security device from the untrust zone are mapped to the private addresses of the servers as follows:

- The destination IP address 203.0.113.200 and port 80 is translated to the private address 192.168.1.200 and port 80.
- The destination IP address 203.0.113.200 and port 8000 is translated to the private address 192.168.1.220 and port 8000.

Copyright © 2016, Juniper Networks, Inc.

Figure 211: Destination NAT Address and Port Translation

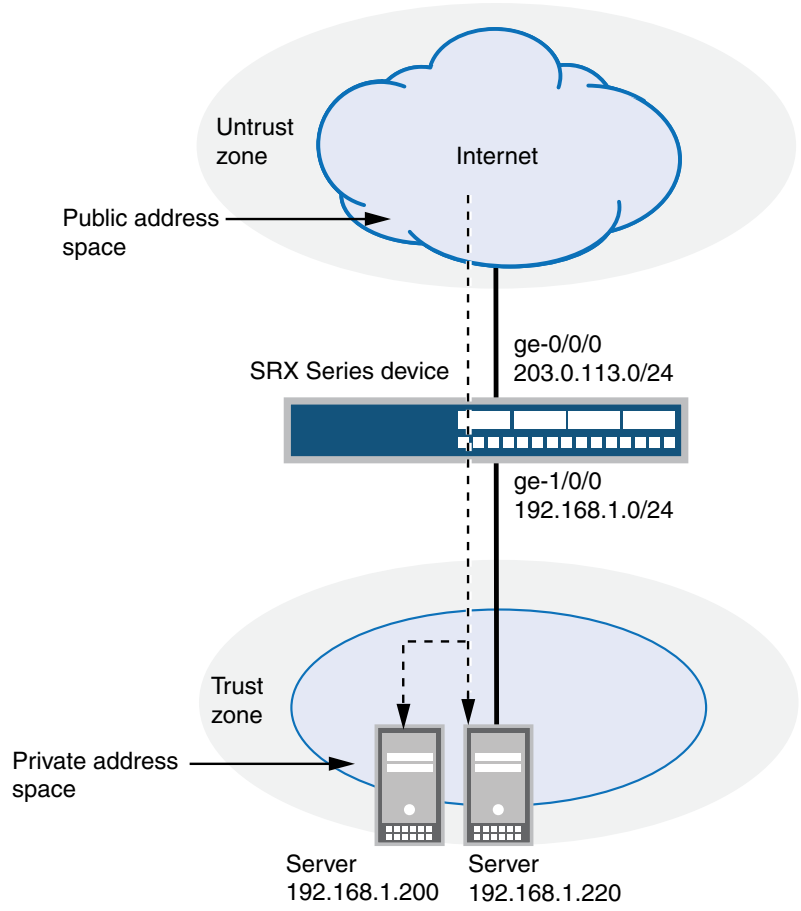

Original Destination IP	Translated Destination IP
203.0.113.200 port 80	192.168.1.200 port 80
203.0.113.200 port 8000	192.168.1.220 port 8000

g030666

This example describes the following configurations:

- Destination NAT pool **dst-nat-pool-1** that contains the IP address 192.168.1.200 port 80.
- Destination NAT pool **dst-nat-pool-2** that contains the IP address 192.168.1.220 and port 8000.

Copyright © 2016, Juniper Networks, Inc.

- Destination NAT rule set **rs1** with rule **r1** to match packets received from the untrust zone with the destination IP address 203.0.113.200 and destination port 80. For matching packets, the destination address is translated to the address in the **dst-nat-pool-1** pool.
- Destination NAT rule set **rs1** with rule **r2** to match packets received from the untrust zone with the destination IP address 203.0.113.200 and destination port 8000. For matching packets, the destination IP address and port are translated to the address and port in the **dst-nat-pool-2** pool.
- Proxy ARP for the address 203.0.113.200/32. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.200/32
set security nat destination pool dst-nat-pool-1 address port 80
set security nat destination pool dst-nat-pool-2 address 192.168.1.220/32
set security nat destination pool dst-nat-pool-2 address port 8000
set security nat destination rule-set rs1 from zone untrust
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r1 match destination-port 80
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat destination rule-set rs1 rule r2 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r2 match destination-port 8000
set security nat destination rule-set rs1 rule r2 then destination-nat pool dst-nat-pool-2
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-2 192.168.1.220/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone untrust to-zone trust policy server-access match
    source-address any
set security policies from-zone untrust to-zone trust policy server-access match
    destination-address server-1
set security policies from-zone untrust to-zone trust policy server-access match
    destination-address server-2
set security policies from-zone untrust to-zone trust policy server-access match application
    any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a destination NAT mapping from a public address to a private address:

1. Create destination NAT pools.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat destination]
user@host# set pool dst-nat-pool-1 address 192.168.1.200 port 80
user@host# set pool dst-nat-pool-2 address 192.168.1.220 port 8000
```

2. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs1 from zone untrust
```

3. Configure a rule that matches packets and translates the destination address to the address in the pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200
user@host# set rule-set rs1 rule r1 match destination-port 80
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r2 match destination-address 203.0.113.200
user@host# set rule-set rs1 rule r2 match destination-port 8000
user@host# set rule-set rs1 rule r2 then destination-nat pool dst-nat-pool-2
```

5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
```

6. Configure addresses in the global address book.

```
[edit security address-book global]
user@host# set address server-2 192.168.1.220/32
user@host# set address server-1 192.168.1.200/32
```

7. Configure a security policy that allows traffic from the untrust zone to the servers in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address
   [server-1 server-2] application any
user@host# set policy server-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
destination {
   pool dst-nat-pool-1 {
      address 192.168.1.200/32 port 80;
   }
   pool dst-nat-pool-2 {
      address 192.168.1.220/32 port 8000;
   }
   rule-set rs1 {
      from zone untrust;
```

Copyright © 2016, Juniper Networks, Inc.

```
        rule r1 {
          match {
            destination-address 203.0.113.200;
            destination-port 80;
          }
          then {
            destination-nat pool dst-nat-pool-1;
          }
        }
        rule r2 {
          match {
            destination-address 203.0.113.200
            destination-port 8000;
          }
          then {
            destination-nat pool dst-nat-pool-2;
          }
        }
      }
    }
      proxy-arp {
        interface ge-0/0/0.0 {
          address {
            203.0.113.200;
          }
        }
      }
    user@host# show security policies
    from-zone untrust to-zone trust {
      policy server-access {
        match {
          source-address any;
          destination-address [ server-1 server-2 ];
          application any;
        }
        then {
          permit;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Destination NAT Pool Usage on page 4938
- Verifying Destination NAT Rule Usage on page 4939
- Verifying NAT Application to Traffic on page 4939

Verifying Destination NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the destination NAT pool.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security nat destination pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose Verify that there is traffic matching the destination NAT rule.

Action From operational mode, enter the **show security nat destination rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Destination NAT Configuration Overview on page 4926
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for Subnet Translation on page 4939

Example: Configuring Destination NAT for Subnet Translation

This example describes how to configure a destination NAT mapping of a public subnet address to a private subnet address.

NOTE: Mapping addresses from one subnet to another can also be accomplished with static NAT. Static NAT mapping allows connections to be established from either side of the gateway device, whereas destination NAT allows connections to be established from only one side. However, static NAT only allows translations between blocks of addresses of the same size.

- Requirements on page 4939
- Overview on page 4940
- Configuration on page 4941
- Verification on page 4943

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.

Copyright © 2016, Juniper Networks, Inc.

2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 212, devices in the untrust zone access devices in the trust zone by way of public subnet address 203.0.113.0/24. For packets that enter the Juniper Networks security device from the untrust zone with a destination IP address in the 203.0.113.0/24 subnet, the destination IP address is translated to a private address on the 192.168.1.0/24 subnet.

Figure 212: Destination NAT Subnet Translation

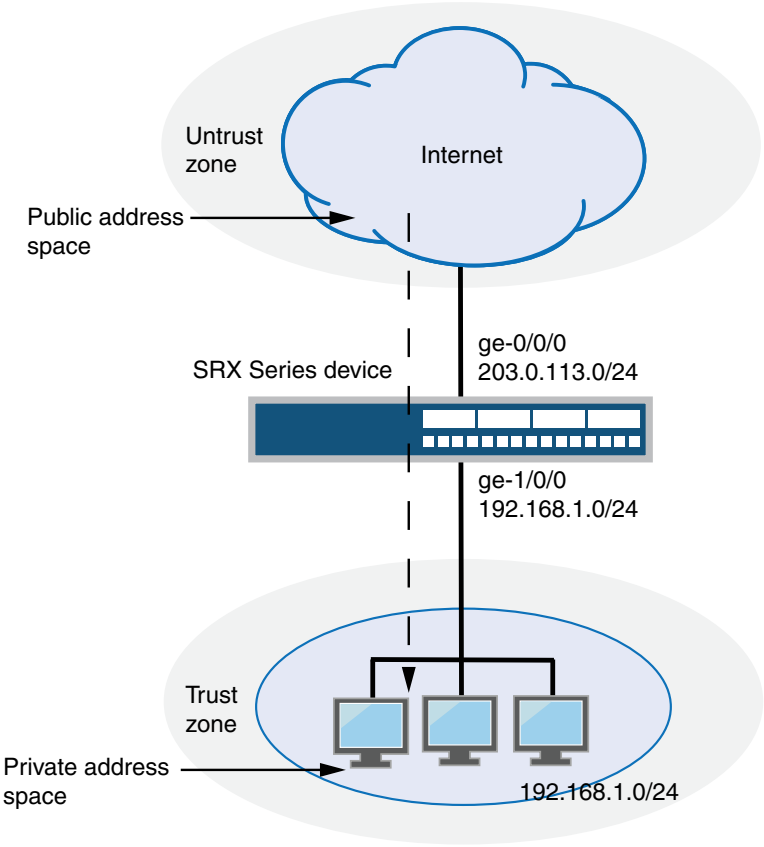

Original Destination IP	Translated Destination IP
203.0.113.0/16	192.168.1.0/24

g030667

Copyright © 2016, Juniper Networks, Inc.

This example describes the following configurations:

- Destination NAT pool **dst-nat-pool-1** that contains the IP address 192.168.1.0/24.
- Destination NAT rule set **rs1** with rule **r1** to match packets received from the ge-0/0/0.0 interface with the destination IP address on the 203.0.113.0/24 subnet. For matching packets, the destination address is translated to the address in the **dst-nat-pool-1** pool.
- Proxy ARP for the addresses 203.0.113.1/32 through 203.0.113.62/32 on the interface ge-0/0/0.0; these are the IP addresses of the hosts that should be translated from the 203.0.113.0/24 subnet. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses. The address 203.0.113.63/32 is assigned to the interface itself, so this address is not included in the proxy ARP configuration. The addresses that are not in the 203.0.113.1/32 through 203.0.113.62/32 range are not expected to be present on the network and would not be translated.
- Security policies to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.0/24
set security nat destination rule-set rs1 from interface ge-0/0/0.0
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.0/24
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.62/32
set security address-book global address internal-net 192.168.1.0/24
set security policies from-zone untrust to-zone trust policy internal-access match
  source-address any
set security policies from-zone untrust to-zone trust policy internal-access match
  destination-address internal-net
set security policies from-zone untrust to-zone trust policy internal-access match
  application any
set security policies from-zone untrust to-zone trust policy internal-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a destination NAT mapping from a public subnet address to a private subnet address:

1. Create the destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set pool dst-nat-pool-1 address 192.168.1.0/24
   ```

2. Create a destination NAT rule set.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat destination]
user@host# set rule-set rs1 from interface ge-0/0/0.0
```

3. Configure a rule that matches packets and translates the destination address to an address in the pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.0/24
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
```

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to
  203.0.113.62/32
```

5. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address internal-net 192.168.1.0/24
```

6. Configure a security policy that allows traffic from the untrust zone to the devices in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy internal-access match source-address any
  destination-address internal-net application any
user@host# set policy internal-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
destination {
  pool dst-nat-pool-1 {
    address 192.168.1.0/24;
  }
  rule-set rs1 {
    from interface ge-0/0/0.0;
    rule r1 {
      match {
        destination-address 203.0.113.0/24;
      }
      then {
        destination-nat pool dst-nat-pool-1;
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/0.0 {
      address {
        203.0.113.1/32 to 203.0.113.62/32;
      }
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security policies
from-zone untrust to-zone trust {
  policy internal-access {
    match {
      source-address any;
      destination-address internal-net;
      application any;
    }
    then {
      permit;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Destination NAT Pool Usage on page 4943
- Verifying Destination NAT Rule Usage on page 4943
- Verifying NAT Application to Traffic on page 4943

Verifying Destination NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the destination NAT pool.

Action From operational mode, enter the **show security nat destination pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose Verify that there is traffic matching the destination NAT rule.

Action From operational mode, enter the **show security nat destination rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Destination NAT Configuration Overview on page 4926
- Example: Improving Security By Configuring Destination NAT for Single Address Translation on page 4926
- Example: Configuring Destination NAT for IP Address and Port Translation on page 4934

Copyright © 2016, Juniper Networks, Inc.

Monitoring Destination NAT Information

Purpose View the destination Network Address Translation (NAT) summary table and the details of the specified NAT destination address pool information.

Action Select **Monitor>NAT> Destination NAT** in the J-Web user interface, or enter the following CLI commands:

- **show security nat destination summary**
- **show security nat destination pool *pool-name***

Table 442 summarizes key output fields in the destination NAT display.

Table 442: Summary of Key Destination NAT Output Fields

Field	Values	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Name	Name of the rule .	–
Ruleset Name	Name of the rule set.	–
From	Name of the routing instance/zone/interface from which the packet flows.	–
Source address range	Source IP address range in the source pool.	–
Destination address range	Destination IP address range in the source pool.	–
Destination port	Destination port in the destination pool.	–
IP protocol	IP protocol.	–
Action	Action taken for a packet that matches a rule.	–
Alarm threshold	Utilization alarm threshold.	–

Copyright © 2016, Juniper Networks, Inc.

Table 442: Summary of Key Destination NAT Output Fields *(continued)*

Field	Values	Action
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule.	–
Translation hits	Number of times a translation in the translation table is used for a destination NAT rule.	–
Pools		
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	–
ID	ID of the pool.	–
Name	Name of the destination pool.	–
Address range	IP address range in the destination pool.	–
Port	Destination port number in the pool.	–
Routing instance	Name of the routing instance.	–
Total addresses	Total IP address, IP address set, or address book entry.	–
Translation hits	Number of times a translation in the translation table is used for destination NAT.	–
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	–

Related Documentation

- Monitoring Source NAT Information on page 4889
- Monitoring Static NAT Information on page 4965
- Monitoring Incoming Table Information on page 4859
- Monitoring Interface NAT Port Information on page 4860

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 203

Configuring Static NAT

- Understanding Static NAT on page 4947
- Understanding Static NAT Rules on page 4948
- Static NAT Configuration Overview on page 4949
- Example: Configuring Static NAT for Single Address Translation on page 4949
- Example: Configuring Static NAT for Subnet Translation on page 4953
- Example: Configuring Static NAT for Port Mapping on page 4958
- Monitoring Static NAT Information on page 4965

Understanding Static NAT

Static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the reverse direction. From the NAT device, the original destination address is the virtual host IP address while the mapped-to address is the real host IP address.

Static NAT allows connections to be originated from either side of the network, but translation is limited to one-to-one or between blocks of addresses of the same size. For each private address, a public address must be allocated. No address pools are necessary.

Static NAT also supports the following types of translation:

- To map multiple IP addresses and specified ranges of ports to a same IP address and different range of ports
- To map a specific IP address and port to a different IP address and port

The port address translation (PAT) is also supported by giving static mapping between destination-port (range) and mapped-port (range).

NOTE: **The original destination address, along with other addresses in source and destination NAT pools, must not overlap within the same routing instance.**

In NAT rule lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules take precedence over source NAT rules.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Static NAT Configuration Overview on page 4949
- Example: Configuring Static NAT for Single Address Translation on page 4949
- Example: Configuring Static NAT for Subnet Translation on page 4953
- Introduction to NAT on page 4839
- Understanding Static NAT Rules on page 4948

Understanding Static NAT Rules

Static Network Address Translation (NAT) rules specify two layers of match conditions:

- Traffic direction—Allows you to specify **from interface**, **from zone**, or **from routing-instance**.
- Packet information—Can be source addresses and ports, and destination addresses and ports.

For all ALG traffic, except FTP, we recommend that you not use the static NAT rule options **source-address** or **source-port**. Data session creation can fail if these options are used because the IP address and the source port value, which is a random value, might not match the static NAT rule. For FTP ALG traffic, the **source-address** option can be used because an IP address can be provided to match the source address of a static NAT rule.

When both source and destination addresses are configured as match conditions for a rule, traffic is matched to both the source address and destination address. Because static NAT is bidirectional, traffic in the opposite direction reverse matches the rule, and the destination address of the traffic is matched to the configured source address.

If multiple static NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface **ge-0/0/0**, rule B is used to perform static NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

Because static NAT rules do not support overlapping addresses and ports, they should not be used to map one external IP address to multiple internal IP addresses for ALG traffic. For example, if different sites want to access two different FTP servers, the internal FTP servers should be mapped to two different external IP addresses.

For the static NAT rule action, specify the translated address and (optionally) the routing instance.

In NAT lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules takes precedence over source NAT rules.

Related Documentation

- Understanding Static NAT on page 4947
- Static NAT Configuration Overview on page 4949

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Static NAT for Single Address Translation on page 4949
- Example: Configuring Static NAT for Subnet Translation on page 4953
- Understanding NAT Rule Sets and Rules on page 4840

Static NAT Configuration Overview

The main configuration tasks for static NAT are as follows:

1. Configure static NAT rules that align with your network and security requirements.
2. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Related Documentation

- Understanding Static NAT on page 4947
- Configuring Proxy ARP (CLI Procedure) on page 4857
- Example: Configuring Static NAT for Single Address Translation on page 4949
- Example: Configuring Static NAT for Subnet Translation on page 4953
- Verifying NAT Configuration on page 4858

Example: Configuring Static NAT for Single Address Translation

This example describes how to configure a static NAT mapping of a single private address to a public address.

- Requirements on page 4949
- Overview on page 4949
- Configuration on page 4951
- Verification on page 4953

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space.

In Figure 213, devices in the untrust zone access a server in the trust zone by way of public address 203.0.113.200/32. For packets that enter the Juniper Networks security device

Copyright © 2016, Juniper Networks, Inc.

from the untrust zone with the destination IP address 203.0.113.200/32, the destination IP address is translated to the private address 192.168.1.200/32. For a new session originating from the server, the source IP address in the outgoing packet is translated to the public address 203.0.113.200/32.

Figure 213: Static NAT Single Address Translation

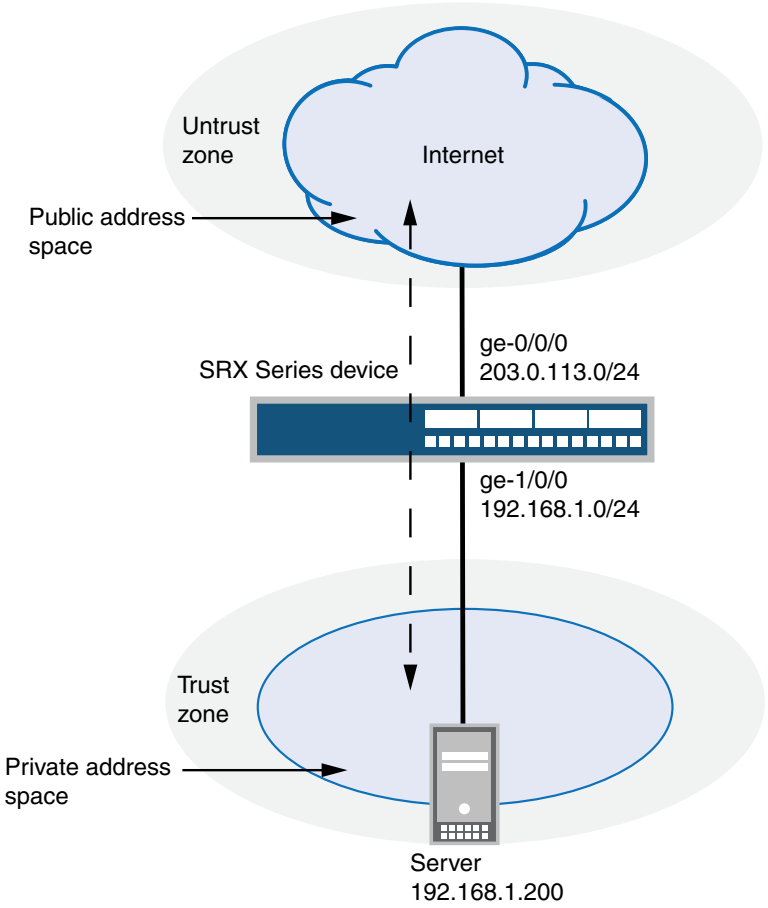

Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.200/32

g030663

This example describes the following configurations:

- Static NAT rule set **rs1** with rule **r1** to match packets from the untrust zone with the destination address 203.0.113.200/32. For matching packets, the destination IP address is translated to the private address 192.168.1.200/32.

Copyright © 2016, Juniper Networks, Inc.

- Proxy ARP for the address 203.0.113.200 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic to and from the 192.168.1.200 server.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set rs1 from zone untrust
set security nat static rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 192.168.1.200/32
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone trust to-zone untrust policy permit-all match
  source-address server-1
set security policies from-zone trust to-zone untrust policy permit-all match
  destination-address any
set security policies from-zone trust to-zone untrust policy permit-all match application
  any
set security policies from-zone trust to-zone untrust policy permit-all then permit
set security policies from-zone untrust to-zone trust policy server-access match
  source-address any
set security policies from-zone untrust to-zone trust policy server-access match
  destination-address server-1
set security policies from-zone untrust to-zone trust policy server-access match application
  any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static NAT mapping from a private address to a public address:

1. Create a static NAT rule set.

   ```
   [edit security nat static]
   user@host# set rule-set rs1 from zone untrust
   ```

2. Configure a rule that matches packets and translates the destination address in the packets to a private address.

   ```
   [edit security nat static]
   user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200/32
   user@host# set rule-set rs1 rule r1 then static-nat prefix 192.168.1.200/32
   ```

3. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200
   ```

4. Configure an address in the global address book.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security address-book global]
user@host# set address server-1 192.168.1.200/32
```

5. Configure a security policy that allows traffic from the untrust zone to the server in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address
  server-1 application any
user@host# set policy server-access then permit
```

6. Configure a security policy that allows all traffic from the server in the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-all match source-address server-1 destination-address
  any application any
user@host# set policy permit-all then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
    from zone untrust;
    rule r1 {
      match {
        destination-address 203.0.113.200/32;
      }
      then {
        static-nat prefix 192.168.1.200/32;
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/0.0 {
      address {
        203.0.113.200/32;
      }
    }
  }
user@host# show security policies
  from-zone trust to-zone untrust {
    policy permit-all {
      match {
        source-address server-1;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  from-zone untrust to-zone trust {
    policy server-access {
      match {
        source-address any;
        destination-address server-1;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static NAT Configuration on page 4953
- Verifying NAT Application to Traffic on page 4953

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Static NAT on page 4947
- Static NAT Configuration Overview on page 4949
- Example: Configuring Static NAT for Subnet Translation on page 4953

Example: Configuring Static NAT for Subnet Translation

This example describes how to configure a static NAT mapping of a private subnet address to a public subnet address.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Address blocks for static NAT mapping must be of the same size.

- Requirements on page 4954
- Overview on page 4954
- Configuration on page 4956
- Verification on page 4958

Requirements

Before you begin:

1. Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In Figure 214, devices in the untrust zone access devices in the trust zone by way of public subnet address 203.0.113.0/24. For packets that enter the Juniper Networks security device from the untrust zone with a destination IP address in the 203.0.113.0/24 subnet, the destination IP address is translated to a private address on the 192.168.1.0/24 subnet. For new sessions originating from the 192.168.1.0/24 subnet, the source IP address in outgoing packets is translated to an address on the public 203.0.113.0/24 subnet.

Copyright © 2016, Juniper Networks, Inc.

Figure 214: Static NAT Subnet Translation

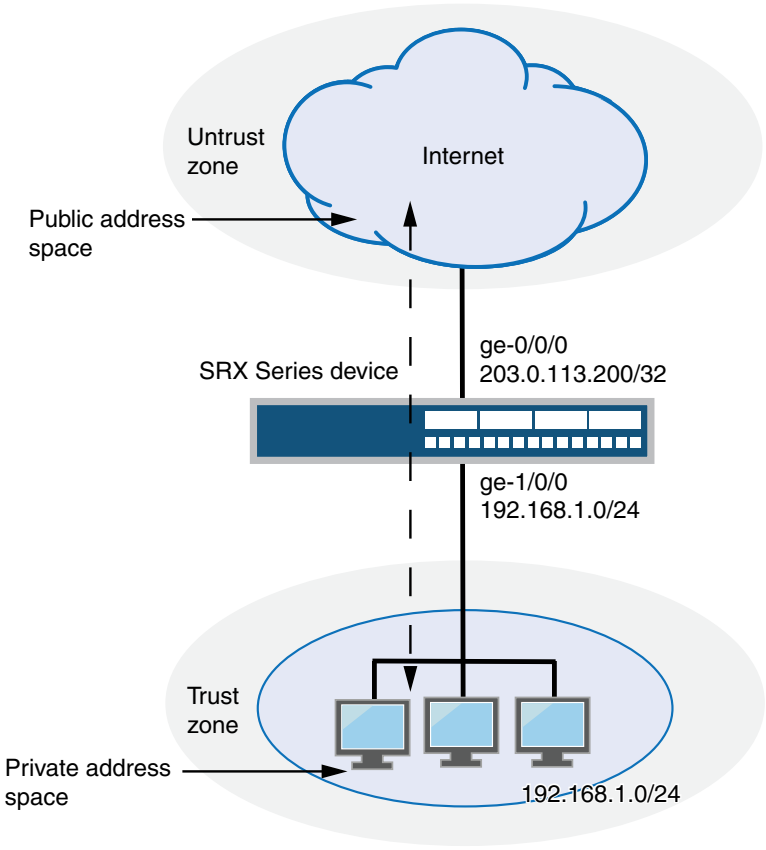

Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.0/24

g030664

This example describes the following configurations:

- Static NAT rule set **rs1** with rule **r1** to match packets received on interface ge-0/0/0.0 with a destination IP address in the 203.0.113.0/24 subnet. For matching packets, the destination address is translated to an address on the 192.168.1.0/24 subnet.
- Proxy ARP for the address ranges 203.0.113.1/32 through 203.0.113.249/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses. The address 203.0.113.250/32 is assigned to the interface itself, so this address is not included in the proxy ARP configuration.
- Security policies to permit traffic to and from the 192.168.1.0/24 subnet.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set rs1 from interface ge-0/0/0.0
set security nat static rule-set rs1 rule r1 match destination-address 203.0.113.0/24
set security nat static rule-set rs1 rule r1 then static-nat prefix 192.168.1.0/24
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.249/32
set security address-book global address server-group 192.168.1.0/24
set security policies from-zone trust to-zone untrust policy permit-all match
  source-address server-group
set security policies from-zone trust to-zone untrust policy permit-all match
  destination-address any
set security policies from-zone trust to-zone untrust policy permit-all match application
  any
set security policies from-zone trust to-zone untrust policy permit-all then permit
set security policies from-zone untrust to-zone trust policy server-access match
  source-address any
set security policies from-zone untrust to-zone trust policy server-access match
  destination-address server-group
set security policies from-zone untrust to-zone trust policy server-access match application
  any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static NAT mapping from a private subnet address to a public subnet address:

1. Create a static NAT rule set.

   ```
   [edit security nat static]
   user@host# set rule-set rs1 from interface ge-0/0/0.0
   ```

2. Configure a rule that matches packets and translates the destination address in the packets to an address in a private subnet.

   ```
   [edit security nat static]
   user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.0/24
   user@host# set rule-set rs1 rule r1 then static-nat prefix 192.168.1.0/24
   ```

3. Configure proxy ARP.

   ```
   [edit security nat]
   user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to
     203.0.113.249/32
   ```

4. Configure an address in the global address book.

   ```
   [edit security address-book global]
   user@host# set address server-group 192.168.1.0/24
   ```


Copyright © 2016, Juniper Networks, Inc.

5. Configure a security policy that allows traffic from the untrust zone to the subnet in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address
  server-group application any
user@host# set policy server-access then permit
```

6. Configure a security policy that allows all traffic from the subnet in the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-all match source-address server-group
  destination-address any application any
user@host# set policy permit-all then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
static {
  rule-set rs1 {
    from interface ge-0/0/0.0;
    rule r1 {
      match {
        destination-address 203.0.113.0/24;
      }
      then {
        static-nat prefix 192.168.1.0/24;
      }
    }
  }
}
  proxy-arp {
    interface ge-0/0/0.0 {
      address {
        203.0.113.1/32 to 203.0.113.249/32;
      }
    }
  }
user@host# show security policies
from-zone trust to-zone untrust {
  policy permit-all {
    match {
      source-address server-group;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone untrust to-zone trust {
```

Copyright © 2016, Juniper Networks, Inc.

```
    policy server-access {
      match {
        source-address any;
        destination-address server-group;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static NAT Configuration on page 4958
- Verifying NAT Application to Traffic on page 4958

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Action From operational mode, enter the **show security nat static rule** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding Static NAT on page 4947
- Static NAT Configuration Overview on page 4949
- Example: Configuring Static NAT for Single Address Translation on page 4949

Example: Configuring Static NAT for Port Mapping

This example describes how to configure static NAT mappings of a public address to private addresses on a specified range of ports.

This topic includes the following sections:

- Requirements on page 4959
- Overview on page 4959
- Configuration on page 4961

Copyright © 2016, Juniper Networks, Inc.

- Verification on page 4963
- Troubleshooting on page 4964

Requirements

Before you begin:

- Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
- Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space.

In Figure 215, devices in the untrust zone access a server in the trust zone by way of public addresses 203.0.113.1/32, 203.0.113.1/32, and 203.0.113.3/32. For packets that enter the Juniper Networks security device from the untrust zone with the destination IP addresses 203.0.113.1/32, 203.0.113.1/32, and 203.0.113.3/32, the destination IP address is translated to the private addresses 10.1.1.1/32,10.1.1.2/32, and 10.1.1.2/32.

Copyright © 2016, Juniper Networks, Inc.

Figure 215: Static NAT for Port Mapping

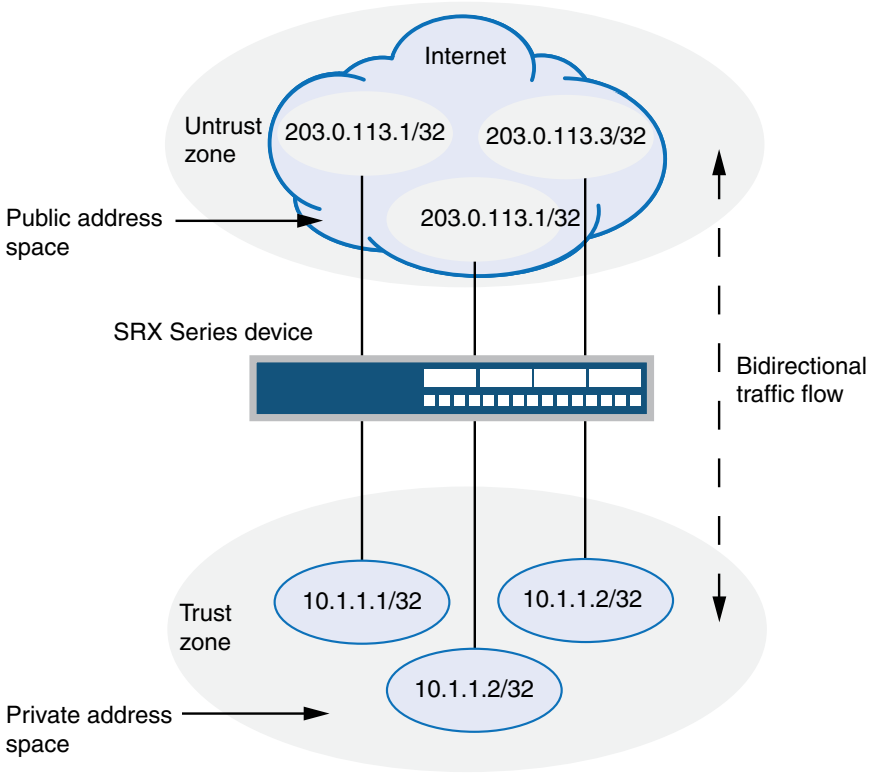

Original Source IP	Translated Source IP
203.0.113.1/32 (port 100 to 200) 203.0.113.1/32 (port 300 to 400) 203.0.113.3/32 (port 300)	10.1.1.1/32 (port 300 to 400) 10.1.1.2/32 (port 300 to 400) 10.1.1.2/32 (port 200)

g034403

Copyright © 2016, Juniper Networks, Inc.

NOTE:

- To configure the destination port, you must use an IP address for the destination address field instead of an IP address prefix.
- You must configure the destination port to configure the mapped port and vice versa.
- Use the same number range for the ports while configuring the destination port and the mapped port.
- If you do not configure the destination port and the mapped port, the IP mapping will be the one-to-one mapping.
- Any address overlapping or any address and port overlapping is not allowed.

This example describes the following configurations:

- Static NAT rule set rs1 with rule r1 to match packets from the untrust zone with the destination address 203.0.113.1/32 and destination port 100 to 200. For matching packets, the destination IP address is translated to the private address 10.1.1.1/32 and mapped to port 300 to 400.
- Static NAT rule set rs1 with rule r2 to match packets from the untrust zone with the destination address 203.0.113.1/32 and destination port 300 to 400. For matching packets, the destination IP address is translated to the private address 10.1.1.2/32 and mapped to port 300 to 400.
- Static NAT rule set rs1 with rule r3 to match packets from the untrust zone with the destination address 203.0.113.3/32 and destination port 300. For matching packets, the destination IP address is translated to the private address 10.1.1.2/32 and mapped to port 200.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set rs from zone untrust
set security nat static rule-set rs rule r1 match destination-address 203.0.113.1/32
set security nat static rule-set rs rule r1 match destination-port 100 to 200
set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1/32
set security nat static rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs rule r2 match destination-address 203.0.113.1/32
set security nat static rule-set rs rule r2 match destination-port 300 to 400
set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs rule r3 match destination-address 203.0.113.3/32
set security nat static rule-set rs rule r3 match destination-port 300
set security nat static rule-set rs rule r3 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs rule r3 then static-nat prefix mapped-port 200
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see Using the CLI Editor in Configuration Mode.

To configure a static NAT mapping from a private subnet address to a public subnet address:

1. Create a static NAT rule set.

   ```
   [edit security nat static]
   user@host# set rule-set rs from zone untrust
   ```

2. Configure a rule that matches packets and translates the destination address in the packets to a private address.

   ```
   [edit security nat static]
   user@host# set rule-set rs rule r1 match destination-address 203.0.113.1/32
   user@host# set rule-set rs rule r1 match destination-port 100 to 200
   user@host# set rule-set rs rule r1 then static-nat prefix 10.1.1.1/32
   user@host# set rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400
   ```

3. Configure a rule that matches packets and translates the destination address in the packets to a private address.

   ```
   [edit security nat static]
   user@host# set rule-set rs rule r2 match destination-address 203.0.113.1/32
   user@host# set rule-set rs rule r2 match destination-port 300 to 400
   user@host# set rule-set rs rule r2 then static-nat prefix 10.1.1.2/32
   user@host# set rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400
   ```

4. Configure a rule that matches packets and translates the destination address in the packets to a private address.

   ```
   [edit security nat static]
   user@host# set rule-set rs rule r3 match destination-address 203.0.113.3/32
   user@host# set rule-set rs rule r3 match destination-port 300
   user@host# set rule-set rs rule r3 then static-nat prefix 10.1.1.2/32
   user@host# set rule-set rs rule r3 then static-nat prefix mapped-port 200
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]

user@host# show security nat

  security {
    nat {
      static {
        rule-set rs {
          from zone untrust;
          rule r1 {
            match {
```


Copyright © 2016, Juniper Networks, Inc.

```
            destination-address 203.0.113.1/32;
            destination-port 100 to 200;
          }
          then {
            static-nat {
              prefix {
                10.1.1.1/32;
                mapped-port 300 to 400;
              }
            }
          }
        }
        rule r2 {
          match {
            destination-address 203.0.113.1/32;
            destination-port 300 to 400;
          }
          then {
            static-nat {
              prefix {
                10.1.1.2/32;
                mapped-port 300 to 400;
              }
            }
          }
        }
        rule r3 {
          match {
            destination-address 203.0.113.3/32;
            destination-port 300;
          }
          then {
            static-nat {
              prefix {
                10.1.1.2/32;
                mapped-port 200;
              }
            }
          }
        }
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Static NAT Configuration

Purpose Verify that there is traffic matching the static NAT rule set.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security nat static rule** command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat static rule all
Total static-nat rules: 3

Static NAT rule: r2                   Rule-set: rs
  Rule-Id                   : 3
  Rule position             : 2
  From zone                 : untrust
  Destination addresses     : 203.0.113.1
  Destination ports         : 300 - 400
  Host addresses            : 10.1.1.2
  Host ports                : 300 - 400
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0

Static NAT rule: r3                   Rule-set: rs
  Rule-Id                   : 4
  Rule position             : 3
  From zone                 : untrust
  Destination addresses     : 203.0.113.3
  Destination ports         : 300 - 300
  Host addresses            : 10.1.1.2
  Host ports                : 200 - 200
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0

Static NAT rule: r1                   Rule-set: rs
  Rule-Id                   : 9
  Rule position             : 1
  From zone                 : untrust
  Destination addresses     : 203.0.113.1
  Destination ports         : 100 - 200
  Host addresses            : 10.1.1.1
  Host ports                : 300 - 400
  Netmask                   : 32
  Host routing-instance     : N/A
  Translation hits          : 0
```

Troubleshooting

- Troubleshooting Static NAT Port Configuration on page 4964

Troubleshooting Static NAT Port Configuration

Problem Static NAT port mapping configuration failures occur during a commit.

Invalid configurations with overlapped IP addresses and ports result in commit failure.

The following example shows invalid configurations with overlapped addresses and ports:

- **set security nat static rule-set rs rule r1 match destination-address 203.0.113.1**

  set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1

- **set security nat static rule-set rs rule r2 match destination-address 203.0.113.1**

Copyright © 2016, Juniper Networks, Inc.

set security nat static rule-set rs rule r2 match destination-port 300 to 400

set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.2

set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400

- **set security nat static rule-set rs rule r1 match destination-address 203.0.113.1**

  set security nat static rule-set rs rule r1 match destination-port 100 to 200

  set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1

  set security nat static rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400

- **set security nat static rule-set rs rule r2 match destination-address 203.0.113.2**

  set security nat static rule-set rs rule r2 match destination-port 300 to 400

  set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.1

  set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 390 to 490

The following error message was displayed when the aforementioned configuration was submitted for commit:

```
error: 'prefix/mapped-port' of static nat rule r2 overlaps with
'prefix/mapped-port' of static nat rule r1
error: configuration check-out failed
```

Solution To configure the destination port, you must avoid any address overlapping or any address and port overlapping. For an example of valid configuration, see "Configuration" on page 4961.

Related Documentation

- Understanding Static NAT on page 4947
- Understanding Static NAT Rules on page 4948
- Static NAT Configuration Overview on page 4949

Monitoring Static NAT Information

Purpose View static NAT rule information.

Action Select **Monitor>NAT>Static NAT** in the J-Web user interface, or enter the following CLI command:

show security nat static rule

Table 443 summarizes key output fields in the static NAT display.

Table 443: Summary of Key Static NAT Output Fields

Field	Values	Action
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.

Copyright © 2016, Juniper Networks, Inc.

Table 443: Summary of Key Static NAT Output Fields *(continued)*

Field	Values	Action
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Position	Position of the rule that indicates the order in which it applies to traffic.	–
Name	Name of the rule.	–
Ruleset Name	Name of the rule set.	–
From	Name of the routing instance/interface/zone from which the packet comes	–
Source addresses	Source IP addresses.	–
Source ports	Source port numbers.	–
Destination addresses	Destination IP address and subnet mask.	–
Destination ports	Destination port numbers .	–
Host addresses	Name of the host addresses.	–
Host ports	Host port numbers.	
Netmask	Subnet IP address.	–
Host routing instance	Name of the routing instance from which the packet comes.	–
Alarm threshold	Utilization alarm threshold.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule.	–
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.	–

Top 10 Translation Hits

Copyright © 2016, Juniper Networks, Inc.

Table 443: Summary of Key Static NAT Output Fields *(continued)*

Graph	Displays the graph of top 10 translation hits.	–

Related Documentation

- Monitoring Source NAT Information on page 4889
- Monitoring Destination NAT Information on page 4944
- Monitoring Incoming Table Information on page 4859
- Monitoring Interface NAT Port Information on page 4860

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 204

Configuring Persistent NAT and NAT64

- Understanding Persistent NAT and NAT64 on page 4969
- Understanding Session Traversal Utilities for NAT (STUN) Protocol on page 4971
- Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation on page 4972
- Persistent NAT and NAT64 Configuration Overview on page 4974
- Example: Configuring Address Persistent NAT64 Pools on page 4975
- Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT on page 4977
- Example: Configuring Persistent NAT with Source NAT Address Pool (CLI) on page 4983
- Example: Configuring Address-Dependent Filtering for IPv6 Clients on page 4984
- Example: Configuring Endpoint-Independent Filtering for IPv6 Clients on page 4987
- Example: Setting Maximum Persistent NAT Bindings on page 4990

Understanding Persistent NAT and NAT64

Persistent NAT allows applications to use the Session Traversal Utilities for NAT (STUN) protocol when passing through NAT firewalls (see "Understanding Session Traversal Utilities for NAT (STUN) Protocol" on page 4971). Persistent NAT ensures that all requests from the same internal transport address (internal IP address and port) are mapped to the same reflexive transport address (the public IP address and port created by the NAT device closest to the STUN server).

NAT64 is a mechanism for translating IPv6 packets to IPv4 packets and vice versa that allows IPv6 clients to contact IPv4 servers using unicast UDP, TCP, or ICMP. It is an enhancement of Network Address Translation-Protocol Translation (NAT-PT).

NAT64 supports the following:

- Endpoint-independent mappings
- Endpoint-independent filtering and address-dependent filtering

NOTE: The mapping and filtering behaviors of NAT64 and persistent NAT are identical.

Copyright © 2016, Juniper Networks, Inc.

The following types of persistent NAT can be configured on the Juniper Networks device:

- Any remote host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. Any external host can send a packet to the internal host by sending the packet to the reflexive transport address.
- Target host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address.
- Target host port—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.

NOTE: The target-host-port configuration is not supported for NAT64 when configured with IPv6 address.

You configure any of the persistent NAT types with source NAT rules. The source NAT rule action can use a source NAT pool (with or without port translation) or an egress interface. Persistent NAT is not applicable for destination NAT, because persistent NAT bindings are based on outgoing sessions from internal to external.

NOTE: Port overloading is used in Junos OS only for normal interface NAT traffic. Persistent NAT does not support port overloading, and you must explicitly disable port overloading with one of the following options at the [edit security nat source] hierarchy level:

- port-overloading off
- port-overloading-factor 1

To configure security policies to permit or deny persistent NAT traffic, you can use two new predefined services—**junos-stun** and **junos-persistent-nat**.

NOTE: Persistent NAT is different from the persistent address feature (see "Understanding Persistent Addresses" on page 4898). The persistent address feature applies to address mappings for source NAT pools configured on the device. The persistent NAT feature applies to address mappings on an external NAT device, and is configured for a specific source NAT pool or egress interface. Also, persistent NAT is intended for use with STUN client/server applications.

Related Documentation

- Understanding Session Traversal Utilities for NAT (STUN) Protocol on page 4971
- Persistent NAT and NAT64 Configuration Overview on page 4974

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Persistent NAT with Source NAT Address Pool (CLI) on page 4983
- Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT on page 4977

Understanding Session Traversal Utilities for NAT (STUN) Protocol

Many video and voice applications do not work properly in a NAT environment. For example, Session Initiation Protocol (SIP), used with VoIP, encodes IP addresses and port numbers within application data. If a NAT firewall exists between the requestor and receiver, the translation of the IP address and port number in the data invalidates the information.

Also, a NAT firewall does not maintain a pinhole for incoming SIP messages. This forces the SIP application to either constantly refresh the pinhole with SIP messages or use an ALG to track registration, a function that may or may not be supported by the gateway device.

The Session Traversal Utilities for NAT (STUN) protocol, first defined in *RFC 3489, Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)* and then later in *RFC 5389, Session Traversal Utilities for NAT*, is a simple client/server protocol. A STUN client sends requests to a STUN server, which returns responses to the client. A STUN client is usually part of an application that requires a public IP address and/or port. STUN clients can reside in an end system such as a PC or in a network server whereas STUN servers are usually attached to the public Internet.

NOTE: **Both the STUN client and STUN server must be provided by the application. Juniper Networks does not provide a STUN client or server.**

The STUN protocol allows a client to:

- Discover whether the application is behind a NAT firewall.
- Determine the type of NAT binding being used (see "Understanding Persistent NAT and NAT64" on page 4969).
- Learn the reflexive transport address, which is the IP address and port binding allocated by NAT device closest to the STUN server. (There may be multiple levels of NAT between the STUN client and the STUN server.)

The client application can use the IP address binding information within protocols such as SIP and H.323.

Related Documentation

- Persistent NAT and NAT64 Configuration Overview on page 4974
- Understanding Persistent NAT and NAT64 on page 4969

Copyright © 2016, Juniper Networks, Inc.

Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation

The NAT64 mechanism enables IPv6 clients to contact IPv4 servers by translating IPv6 addresses to IPv4 addresses (and vice versa). However, some IPv4 applications and services cannot work correctly over IPv6-only networks with standard NAT64 in a dual-translation scenario, such as 464XLAT. In those scenarios, address-persistent translation is required.

Figure 216 illustrates the 464XLAT architecture, whereby IPv4 packets are translated to IPv6 packets on the customer-side translator (CLAT), then go across the IPv6-only network, and are translated back to IPv4 packets on the provider-side translator (PLAT) to access global IPv4-only content in the core network. This architecture uses a combination of stateless translation on the CLAT and stateful translation on the PLAT.

Figure 216: 464XLAT Architecture

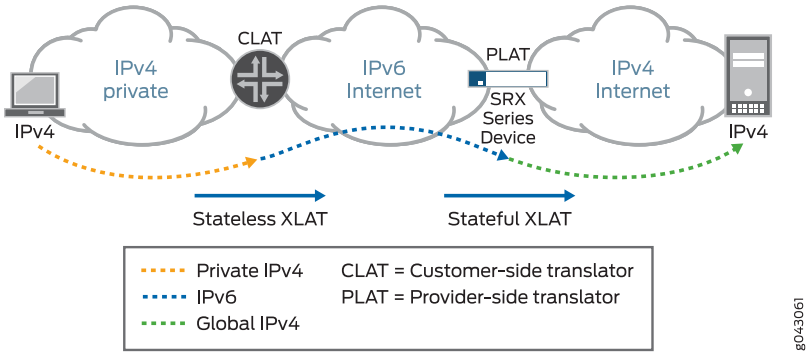

When an SRX Series device functions as a PLAT, it is responsible for keeping the sticky mapping relationship between one specific IPv6 prefix and one translated IPv4 address. The SRX Series device treats the IPv6 prefix as a single user. This mapping is accomplished by configuring the specific IPv6 prefix length in an IPv4 source NAT pool using the **address-persistent** feature.

Copyright © 2016, Juniper Networks, Inc.

Figure 217 illustrates a NAT rule configured in the CLAT, which translates an IPv4 address to an IPv6 address with an address-persistent prefix. With stateless NAT46 translation on the CLAT and stateful NAT64 translation on the PLAT, the traffic from IPv4 host 192.168.1.2 reaches the global server 198.51.100.1 over an IPv6-only network.

Figure 217: NAT64 Translation on the PLAT (SRX Series Device)

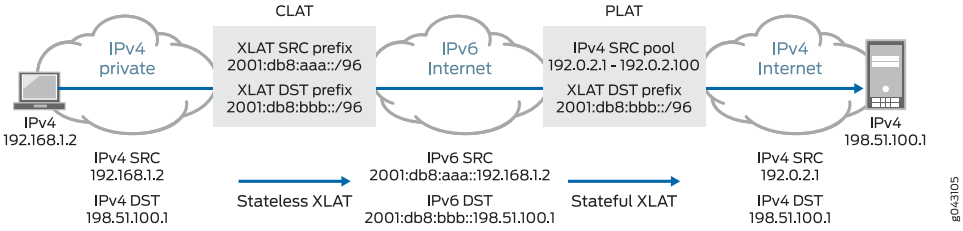

Table 444 lists other NAT features and their compatibility with the address-persistent feature.

Table 444: NAT Feature Compatibility with the Address Persistent Feature

Feature			Compatible
PAT pools	IPv4	NAT IPv4 to IPv6	No
		NAT IPv6 to IPv4	Yes
	IPv6	NAT IPv4 to IPv6	No
		NAT IPv6 to IPv4	No
Non-PAT pools			No
Port-overloading			Yes
Persistent NAT in PAT pool			Yes
Port block allocation			Yes
Deterministic NAT			No
Address pooling paired			No
ALG (Existing ALG NAT translations , such as FTP/PPTP/RTSP/DNS/SIP from native IPv6 clients.)			Yes

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

Persistent NAT and NAT64 Configuration Overview

To configure persistent NAT, specify the following options with the source NAT rule action (for either a source NAT pool or an egress interface):

- The type of persistent NAT—One of the following: any remote host, target host, or target host port (see "Understanding Persistent NAT and NAT64" on page 4969).
- (Optional) Address mapping—This option allows requests from a specific internal IP address to be mapped to the same reflexive IP address; internal and reflexive ports can be any ports. An external host using any port can send a packet to the internal host by sending the packet to the reflexive IP address (with a configured incoming policy that allows external to internal traffic). If this option is not configured, the persistent NAT binding is for specific internal and reflexive transport addresses.

  You can only specify the **address-mapping** option when the persistent NAT type is any remote host and the source NAT rule action is one of the following actions:

  - Source NAT pool with IP address shifting
  - Source NAT pool with no port translation and no overflow pool

- (Optional) Inactivity timeout—Time, in seconds, that the persistent NAT binding remains in the device's memory when all the sessions of the binding entry have expired. When the configured timeout is reached, the binding is removed from memory. The default value is 300 seconds. Configure a value from 60 through 7200 seconds.

  When all sessions of a persistent NAT binding have expired, the binding remains in a query state in the SRX Series device's memory for the specified inactivity timeout period. The query binding is automatically removed from memory when the inactivity timeout period expires (the default is 300 seconds). You can explicitly remove all or specific persistent NAT query bindings with the **clear security nat source persistent-nat-table** command.

- (Optional) Maximum session number—Maximum number of sessions with which a persistent NAT binding can be associated. The default is 30 sessions. Configure a value from 8 through 100.

For interface NAT, you need to explicitly disable port overloading with one of the following options at the [**edit security nat source**] hierarchy level:

- port-overloading off
- port-overloading-factor 1

Finally, there are two predefined services that you can use in security policies to permit or deny STUN and persistent NAT traffic:

- **junos-stun**—STUN protocol traffic.
- **junos-persistent-nat**—Persistent NAT traffic.

Copyright © 2016, Juniper Networks, Inc.

For the **any remote host** persistent NAT type, the direction of the security policy is from external to internal. For target host or target host port persistent NAT types, the direction of the security policy is from internal to external.

Related Documentation

- Example: Configuring Persistent NAT with Source NAT Address Pool (CLI) on page 4983
- Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT on page 4977
- Understanding Persistent NAT and NAT64 on page 4969
- Understanding Session Traversal Utilities for NAT (STUN) Protocol on page 4971

Example: Configuring Address Persistent NAT64 Pools

This example shows how to configure address persistent NAT64 pools to ensure a sticky mapping relationship between one specific IPv6 prefix, which is calculated by the configured IPv6 prefix length, and one translated IPv4 address.

- Requirements on page 4975
- Overview on page 4975
- Configuration on page 4975
- Verification on page 4977

Requirements

Before you begin, be sure the existing NAT rules and pool configuration do not conflict with the new one.

Overview

In this example, you configure an IPv6 prefix length of /64 in an IPv4 source NAT pool for NAT IPv6 to IPv4 translations. Traffic matching the NAT rule and NAT pool perform address persistent translation between the IPv6 prefix and the IPv4 translated address. This configuration can be used on the provider-side translator (PLAT) in a dual-translation scenario, 464XLAT, to enable IPv4 services to work over IPv6-only networks.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool NAT64 address 198.51.100.240/32 to 198.51.100.254/32
set security nat source pool NAT64 address-persistent subscriber ipv6-prefix-length 64
set security nat source rule-set RS1 from zone trust
set security nat source rule-set RS1 to zone untrust
set security nat source rule-set RS1 rule R1 match source-address 2001:db8::/32
set security nat source rule-set RS1 rule R1 match destination-address 198.51.100.198/32
set security nat source rule-set RS1 rule R1 then source-nat pool NAT64
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool NAT64 address 198.51.100.240/32 to 198.51.100.254/32
```

2. Specify the IPv6 prefix length for the source NAT pool.

```
[edit security nat source]
user@host# set pool NAT64 address-persistent subscriber ipv6-prefix-length 64
```

3. Create a rule set.

```
[edit security nat source]
user@host# set rule-set RS1 from zone trust
user@host# set rule-set RS1 to zone untrust
```

4. Match the rule.

```
[edit security nat source]
user@host# set rule-set RS1 rule R1 match source-address 2001:db8::/32
user@host# set rule-set RS1 rule R1 match destination-address 198.51.100.198/32
```

5. Provide the action to be performed when the rule matches.

```
[edit security nat source]
user@host# set security nat source rule-set RS1 rule R1 then source-nat pool NAT64
```

Results

From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool NAT64 {
    address {
      198.51.100.240/32 to 198.51.100.254/32;
    }
    address-persistent subscriber ipv6-prefix-length 64;
  }
  rule-set RS1 {
    from zone trust;
    to zone untrust;
    rule R1 {
      match {
        source-address 2001:db8::/32;
        destination-address 198.51.100.198/32;
      }
      then {
        source-nat {
          pool {
            NAT64;
          }
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying NAT Application to Traffic

Purpose Verify that the same IPv6 prefix is translated to the persistent IPv4 address.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation on page 4972

Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT

You can configure any of the persistent NAT types with source NAT rules. This example illustrates how to apply persistent NAT with an interface IP address and how to use an interface IP address as a NAT IP address to perform persistent NAT for a specific internal host. It also shows how to maintain persistent address port mapping behavior and persistent NAT filter behavior for the host. You must disable port overloading for interface NAT.

- Requirements on page 4977
- Overview on page 4977
- Configuration on page 4979
- Verification on page 4981

Requirements

This example uses the following hardware and software components:

- 1 SRX Series device
- 4 PCs

Before you begin:

- Understand the concepts of persistent NAT. See "Persistent NAT and NAT64 Configuration Overview" on page 4974 .

Overview

In a Carrier Grade NAT (CGN) network deployment, you can configure the interface IP address as a NAT address to perform persistent network address translation. In this way, the internal host can create one source NAT mapping relationship by the outgoing traffic

Copyright © 2016, Juniper Networks, Inc.

initiated from internal to external. Then the external host sends traffic back to this internal host by sending the traffic to this interface NAT address through the shared NAT mapping relationship.

In this example, you first configure the interface NAT rule set int1 to match traffic from interface ge-0/0/1 to interface ge-0/0/2, and then you configure the NAT rule in1 to match the specific source and destination addresses to perform persistent NAT. You configure the **any remote host** persistent NAT type when interface NAT is performed.

For packets with source address 192.0.2.0/24 (internal phones) and destination address 198.51.100.0/24 (including STUN server, SIP proxy server, and external phones), you configure interface NAT with the **any remote host** persistent NAT type. Then you disable port overloading for interface NAT.

Next, you configure a security policy to allow persistent NAT traffic from the external network (external zone) to the internal network (internal zone) for any of the remote host persistent NAT types.

Topology

Figure 218 shows an interface persistent NAT topology.

Figure 218: Interface Persistent NAT Topology

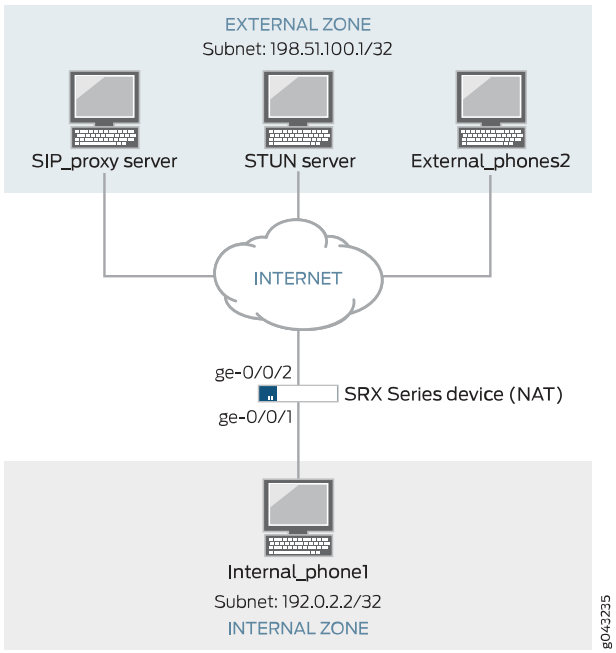

Table 445 shows the parameters configured in this example.

Copyright © 2016, Juniper Networks, Inc.

Table 445: Interfaces, Zones, Servers, and IP Address Information

Parameter	Description
External Zone	External network
Internal Zone	Internal network
External_phones2	Phone2 address of external network
Internal_phone1	Phone1 address of internal network
SIP_proxy server	SIP proxy server address of external network
STUN server	STUN server address of external network
Subnet 198.51.100.1/32	Destination IP address
Subnet 192.0.2.2/32	Source IP address
ge-0/0/1 and ge-0/0/2	NAT interfaces for traffic direction

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source rule-set int1 from interface ge-0/0/1.0
set security nat source rule-set int1 to interface ge-0/0/2.0
set security nat source rule-set int1 rule in1 match source-address 192.0.2.0/24
set security nat source rule-set int1 rule in1 match destination-address 198.51.100.0/24
set security nat source rule-set int1 rule in1 then source-nat interface persistent-nat permit
any-remote-host
set security nat source interface port-overloading off
set security policies from-zone internal to-zone external policy stun_traffic match source-address
 internal_phones destination-address stun_server application junos-stun
set security policies from-zone internal to-zone external policy sip_proxy_traffic match
source-address internal_phones destination-address sip_proxy_server application junos-sip
set security policies from-zone internal to-zone external policy sip_traffic match source-address
 internal_phones destination-address external_phones application junos-persistent-nat
set security policies from-zone internal to-zone external policy sip_traffic then permit
set security policies from-zone internal to-zone external policy stun_traffic then permit
set security policies from-zone internal to-zone external policy sip_proxy_traffic then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an interface NAT rule set:

1. Create a persistent NAT rule for an interface NAT.

   [edit security nat source rule-set int1]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from interface ge-0/0/1.0
user@host# set to interface ge-0/0/2.0
user@host# set rule in1 match source-address 192.0.2.0/24
user@host# set rule in1 match destination-address 198.51.100.0/24
user@host# set rule in1 then source-nat interface persistent-nat permit
  any-remote-host
```

2. Disable port overloading for interface NAT.

```
[edit security]
user@host# set nat source interface port-overloading off
```

3. Configure a security policy to allow STUN traffic from internal SIP phones to an external STUN server.

```
[edit security policies]
user@host# set from-zone internal to-zone external policy stun_traffic match
  source-address internal_phones destination-address stun_server application
  junos-stun
```

4. Configure a security policy to allow SIP proxy traffic from internal SIP phones to an external SIP proxy server.

```
[edit security policies]
user@host# set from-zone internal to-zone external policy sip_proxy_traffic match
  source-address internal_phones destination-address sip_proxy_server application
  junos-sip
```

5. Configure a security policy to allow SIP traffic from external SIP phones to internal SIP phones.

```
[edit security policies]
user@host# set from-zone internal to-zone external policy sip_traffic match
  source-address internal_phones destination-address external_phones application
  junos-persistent-nat
user@host# set from-zone internal to-zone external policy sip_traffic then permit
user@host#set from-zone internal to-zone external policy stun_traffic then permit
user@host#set from-zone internal to-zone external policy sip_proxy_traffic then
  permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security nat
source {
  interface {
    port-overloading off;
  }
  rule-set int1 {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule in1 {
      match {
        source-address 192.0.2.0/24;
        destination-address 198.51.100.0/24;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
        then {
          source-nat {
            interface {
              persistent-nat {
                permit any-remote-host;
              }
            }
          }
        }
      }
    }
  }
[edit]
user@host# show security policies
from-zone internal to-zone external {
  policy stun_traffic {
    match {
      source-address internal_phones;
      destination-address stun_server;
      application junos-stun;
    }
    then {
      permit;
    }
  }
  policy sip_proxy_traffic {
    match {
      source-address internal_phones;
      destination-address sip_proxy_server;
      application junos-sip;
    }
    then {
      permit;
    }
  }
  policy sip_traffic {
    match {
      source-address internal_phones;
      destination-address external_phones;
      application junos-persistent-nat;
    }
    then {
      permit;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That Rules Are Matched and Used on page 4982
- Verifying That NAT Traffic Sessions Are Established on page 4982

Copyright © 2016, Juniper Networks, Inc.

Verifying That Rules Are Matched and Used

Purpose Verify that all the rules are matched and used.

Action From operational mode, enter the **show security nat source persistent-nat-table all** command.

```
user@host>show security nat source persistent-nat-table all
Internal  Reflective                 Source     Type
Left_time/Curr_Sess_Num/   Source
      In_IP     In_Port I_Proto Ref_IP     Ref_Port R_Proto NAT Pool  Conf_time
Max_Sess_Num    NAT Rule
     192.0.2.12  17012   udp    198.51.100.1  28153     udp   interface
any-remote-host  3528/3600     -/-           in1
     192.0.2.12  7078    udp    198.51.100.1  6133      udp   interface
any-remote-host     -/300      1/30          in1
```

Meaning The output displays a summary of persistent NAT information.

Verifying That NAT Traffic Sessions Are Established

Purpose Verify that the sessions are established on the device.

Action From operational mode, enter the **show security flow session** command.

```
user@host>show security flow session

Session ID: 6992, Policy name: sip_proxy_traffic/5, Timeout: 16, Valid
  In: 192.0.2.12/17012 --> 198.51.100.45, If: ge-0/0/1.0, Pkts: 4, Bytes: 1850
  Out: 198.51.100.45/5060 --> 198.51.100.1, If: ge-0/0/2.0, Pkts: 5, Bytes: 2258

Session ID: 7382, Policy name: stun_traffic/4, Timeout: 16, Valid
  In: 192.0.2.12/7078 --> 198.51.100.49, If: ge-0/0/1.0, Pkts: 20, Bytes: 1040
  Out: 198.51.100.49/3478 --> 198.51.100.1, If: ge-0/0/2.0, Pkts: 0, Bytes: 0
```

Meaning The **show security flow session** command displays active sessions on the device and each session's associated security policy. The output shows traffic entering the device using the private source address 192.0.2.12 destined to a public host at 198.51.100.45. The return traffic from this flow travels to the translated public address 198.51.100.1.

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **sip_proxy_traffic**— Policy name that permitted the SIP traffic from the internal SIP phones to the external SIP proxy server.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers. The session is UDP, and the source interface for this session is ge-0/0/1.0).

Copyright © 2016, Juniper Networks, Inc.

- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers. The session is UDP, and the destination interface for this session is ge-0/0/2.0).
- **stun_traffic**—Policy name that permitted the STUN traffic from the internal SIP phones to the external STUN server.

Related Documentation

- Understanding Persistent NAT and NAT64 on page 4969
- Persistent NAT and NAT64 Configuration Overview on page 4974

Example: Configuring Persistent NAT with Source NAT Address Pool (CLI)

You can configure any of the persistent NAT types with source NAT rules. The example in this section shows how to configure persistent NAT when source NAT is performed with a user-defined address pool.

The following example configures the target host persistent NAT type when source NAT is performed. In the following configuration, the source NAT address pool **sp1** consists of the address 203.0.113.5/32. The source NAT rule set **srs1** configures the following:

- Traffic direction is from zone **internal** to zone **external**.
- For packets with source address in the 192.0.2.0/24 subnet (internal phones) and destination address 198.51.100.0/24 (including STUN server, SIP proxy server and external phones), use the source NAT pool **sp1** to perform source NAT with the target host persistent NAT type.
- Set the persistent NAT **inactivity-timeout** to 180 seconds.

To configure the source NAT address pool:

```
user@host# set security nat source pool sp1 address 203.0.113.5/32
```

To configure the source NAT rule set:

```
user@host# set security nat source rule-set srs1 from zone internal
user@host# set security nat source rule-set srs1 to zone external
user@host# set security nat source rule-set srs1 rule sr1 match source-address 192.0.2.0/24
user@host# set security nat source rule-set srs1 rule sr1 match destination-address
  198.51.100.0/24
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool sp1
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool persistent-nat
  permit target-host
user@host# set security nat source rule-set srs1 rule sr1 then source-nat pool persistent-nat
  inactivity-timeout 180
```

For the target host persistent NAT type, configure a security policy to allow persistent NAT traffic from the internal network (internal zone) to the external network (external zone).

To configure a security policy to allow STUN traffic from internal SIP phones to an external STUN server:

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security policies from-zone internal to-zone external policy stun_traffic
  match source-address internal_phones destination-address stun_server application
  junos-stun
user@host# set security policies from-zone internal to-zone external policy stun_traffic
  then permit
```

To configure a security policy to allow SIP proxy traffic from internal SIP phones to an external SIP proxy server:

```
user@host# set security policies from-zone internal to-zone external policy
  sip_proxy_traffic match source-address internal_phones destination-address
  sip_proxy_server application junos-sip
user@host# set security policies from-zone internal to-zone external policy
  stun_proxy_traffic then permit
```

To configure a security policy to allow SIP traffic from internal to external SIP phones:

```
user@host# set security policies from-zone internal to-zone external policy sip_traffic
  match source-address internal_phones destination-address external_phones application
  junos-persistent-nat
user@host# set security policies from-zone internal to-zone external policy sip_traffic
  then permit
```

Related Documentation

- Understanding Persistent NAT and NAT64 on page 4969
- Persistent NAT and NAT64 Configuration Overview on page 4974

Example: Configuring Address-Dependent Filtering for IPv6 Clients

This example shows how to configure address-dependent filtering for IPv6 clients using NAT64.

- Requirements on page 4984
- Overview on page 4984
- Configuration on page 4985
- Verification on page 4987

Requirements

Before you begin:

- Ensure that IPv6 is enabled on the device.
- Ensure that the existing NAT rule and pool configuration do not conflict with the new ones.

Overview

In this example you use NAT64 to send packets from the IPv6 internal host to the IPv4 external host and from the IPv4 external host to the IPv4 internal host.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address
  2001:db8::/128
set security nat static rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address
  2001:db8::/96
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address
  10.2.2.15
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
  permit target-host
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure address-dependent filtering for IPv6 clients:

1. Create a set of rules for NAT64.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs from interface ge-0/0/1
   ```

2. Match the rule.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs rule test_rule match destination-address
     2001:db8::/128
   ```

3. Provide the action to be performed when the rule matches.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
   ```

4. Define a source address pool and add the address to the pool.

   ```
   [edit security nat]
   user@host# set source pool myipv4 address 203.0.113.2
   ```

5. Create another set of rules for NAT64.

   ```
   [edit security nat]
   user@host# set source rule-set myipv4_rs from interface ge-0/0/1
   ```

6. Match the rule with the source address.

   ```
   [edit security nat]
   user@host# set source rule-set myipv4_rs rule ipv4_rule match source-address
     2001:db8::/96
   ```

Copyright © 2016, Juniper Networks, Inc.

7. Match the rule with the destination address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match destination-address
  10.2.2.15
```

8. Provide the action to be performed when the rules match.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

9. Configure persistent NAT.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool
  persistent-nat permit target-host
```

Results From configuration mode, confirm your configuration by entering the **show nat source** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit security]
user@host#show nat source
  pool myipv4 {
      address {
        203.0.113.2/32;
      }
    }
    rule-set test_rs {
      rule test_rule {
        match {
          destination-address 2001:db8::/128;
        }

      }

    }
    rule-set myipv4_rs {
      from interface ge-0/0/1.0;
      to interface ge-0/0/2.0;
      rule ipv4_rule {
        match {
          source-address 2001:db8::/96;
          destination-address 10.2.2.15/32;
        }
        then {
          source-nat {
            pool {
              myipv4;
              persistent-nat {
                permit target-host;
              }
            }
          }
        }
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly:

- Verifying that the Configuration Is Enabled and Working on page 4987
- Verifying that Rules Are Matched and Used on page 4987

Verifying that the Configuration Is Enabled and Working

Purpose Verify that the configuration is enabled and working.

Action From operational mode, enter the following commands.

- **show security nat static rule test_rule**
- **show security nat source rule ipv4_rule**
- **show security nat source pool myipv4**

Verifying that Rules Are Matched and Used

Purpose Verify that all the rules are matched and used.

Action From operational mode, enter the **show security nat source persistent-nat-table all** command.

Related Documentation

- Understanding Persistent NAT and NAT64 on page 4969
- Persistent NAT and NAT64 Configuration Overview on page 4974

Example: Configuring Endpoint-Independent Filtering for IPv6 Clients

This example shows how to configure endpoint-independent filtering for IPv6 clients using NAT64.

- Requirements on page 4987
- Overview on page 4988
- Configuration on page 4988
- Verification on page 4990

Requirements

Before you begin:

- Ensure that IPv6 is enabled on the device

Copyright © 2016, Juniper Networks, Inc.

- Ensure that the existing NAT rules and pool configuration do not conflict with the new ones.

Overview

In this example you use NAT64 to send packets from the IPv6 internal host to the IPv4 external host and from the IPv4 external host to the IPv4 internal host.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address
  2001:db8::/128
set security nat static rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address
  2001:db8::/96
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address
  10.2.2.15
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
  permit any-remote-host
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure endpoint-independent filtering for IPv6 clients:

1. Create a set of rules for NAT64.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs from interface ge-0/0/1
   ```

2. Match the rule.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs rule test_rule match destination-address
     2001:db8::/128
   ```

3. Provide the action to be performed when the rule matches.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
   ```

4. Define a source address pool and add the address to the pool.

   ```
   [edit security nat]
   user@host# set source pool myipv4 address 203.0.113.2
   ```


Copyright © 2016, Juniper Networks, Inc.

5. Create another set of rules for NAT64.

```
[edit security nat]
user@host# set source rule-set myipv4_rs from interface ge-0/0/1
```

6. Match the rule with the source address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match source-address
  2001:db8::/96
```

7. Match the rule with the destination address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match destination-address
  10.2.2.15
```

8. Provide the action to be performed when the rules match.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

9. Configure persistent NAT.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool
  persistent-nat permit any-remote-host
```

Results From configuration mode, confirm your configuration by entering the **show nat source** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit security]
user@host#show nat source
    pool myipv4 {
      address {
        203.0.113.2/32;
      }
    }
    rule-set test_rs {
      rule test_rule {
        match {
          destination-address 2001:db8::/128;
        }

      }

    }
    rule-set myipv4_rs {
      from interface ge-0/0/1.0;
      to interface ge-0/0/2.0;
      rule ipv4_rule {
        match {
          source-address 2001:db8::/96;
          destination-address 10.2.2.15/32;
        }
        then {
          source-nat {
```

Copyright © 2016, Juniper Networks, Inc.

```
                pool {
                  myipv4;
                  persistent-nat {
                    permit any-remote-host;
                  }
                }
              }
            }
          }
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly:

- Verifying that Configuration is Enabled and Working on page 4990
- Verifying that Rules Are Matched and Used. on page 4990

Verifying that Configuration is Enabled and Working

Purpose Verify that the configuration is enabled and working.

Action From operational mode, enter the following commands.

- **show security nat static rule test_rule**
- **show security nat source rule ipv4_rule**
- **show security nat source pool myipv4**

Verifying that Rules Are Matched and Used.

Purpose Verify that all the rules are matched and used.

Action From operational mode, enter the **show security nat source persistent-nat-table all** command.

Related Documentation

- Understanding Persistent NAT and NAT64 on page 4969
- Persistent NAT and NAT64 Configuration Overview on page 4974

Example: Setting Maximum Persistent NAT Bindings

This example shows how to increase the persistent NAT capacity.

Requirements

Before you begin, see “Understanding Persistent NAT and NAT64” on page 4969.

 Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you enable the maximize persistent NAT capacity option. To enable this option, the supported central point maximum binding capacity can be approximately increased to 1/8 of the central point session capacity up to 2M and the supported SPU maximum binding capacity can be approximately increased to 1/4 of each SPU session capacity. Accordingly, the flow session capacity will decrease by 1/4 on both the CP and each of the SPU.

By default, the persistent NAT binding capacity on both the central point and the SPU of a high-end SRX Series device is 64K. In this example, you enable the session capacity to maximum 20,000,000 on the central point and maximum 1100K on each of the SPU with maximum session configuration. If you enable the maximize persistent NAT capacity option, a high-end SRX Series device with 4 GB of memory can support maximum 2M persistent NAT bindings on the central point and 275K bindings on each of the SPU.

Configuration

Step-by-Step Procedure

To increase the persistent NAT capacity:

1. Set maximize persistent NAT capacity option.

   ```
   [edit]
   user@host# set security forwarding-process application-services
     maximize-persistent-nat-capacity
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

3. Restart the system from operational mode.

   ```
   [edit]
   user@host# request system reboot
   ```

   NOTE: When switching to maximize persistent NAT capacity mode or back to regular mode, you must restart the device.

4. If you want to switch the device back to regular mode, delete the maximize persistent NAT capacity mode configuration.

   ```
   [edit]
   user@host# delete security forwarding-process application-services
     maximize-persistent-nat-capacity
   ```

Verification

To verify the configuration is working properly, enter the **show security forwarding-process application-services** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 205

Configuring NAT Hairpinning

- Persistent NAT Hairpinning Overview on page 4993
- Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting on page 4994

Persistent NAT Hairpinning Overview

When traffic is sent between two hosts, the source host of the traffic may only know the destination host by its public IP address. In reality, the destination host may be in the same private address space as the source host. Hairpinning is the process of returning the traffic in the direction from where it came from as a way to get it to its destination host in a private subnetwork.

Generally, a source host in a subnetwork may not recognize that the traffic is intended for a destination host within the same subnetwork, because it identifies the destination host only by its public IP address. The NAT analyzes the IP packets and routes the packet back to the correct host.

NAT hairpinning support is required if two hosts on the internal network want to communicate with each other by using a binding on the NAT device. In this case, the NAT device receives a packet from the internal network and forwards it back to the internal network. If hairpinning is not supported, forwarding the packet will fail and it will be dropped.

NOTE:

- NAT hairpinning behavior is not supported by target host persistent NAT and target host port persistent NAT. Only any remote host persistent NAT supports hairpinning behavior.
- To allow hairpinning, you must configure a security policy to allow traffic between endpoints in the same zone. Actually the two endpoints can be located in two different zones as well as long as either of the two hosts can only see the public address of the peer.
- Persistent NAT hairpinning applies only to any remote host persistent NAT type.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting on page 4994
- Introduction to NAT on page 4839

Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting

This example shows how to configure persistent NAT hairpinning.

Requirements

Before you begin:

- Configure network interfaces on the device. See *Interfaces Feature Guide for Security Devices*.
- Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.

Overview

Hairpinning allows packets from the private network to be translated and then looped back to the private network rather than being passed through to the public network. Hairpinning feature enables using a corresponding record in the NAT table to recognize that a packet is addressed to a host in the local network. Then it translates the destination IP address and sends the packet back to the local network (as well as in case of port mapping). This ensures that traffic between the two hosts work properly.

Hairpinning enables two endpoints (Host 1 and Host 2) on the private network to communicate even if they only use each other's external IP addresses and ports. This is explained in Figure 219.

When Host 1 sends traffic to Host 3, a NAT binding between Host 1's internal source IP address and port is associated in the NAT table with its external IP address and port. The same thing happens when Host 2 sends traffic to Host 3. In this way, when Host 1 and Host 2 want to communicate, they can identify each other's external IP addresses.

For example, if Host 1 communicates with Host 2, NAT (with hairpinning support) is used to route the packets, which contain Host 2's external address, back to Host 2's internal address.

Copyright © 2016, Juniper Networks, Inc.

Figure 219: Persistent NAT Hairpinning

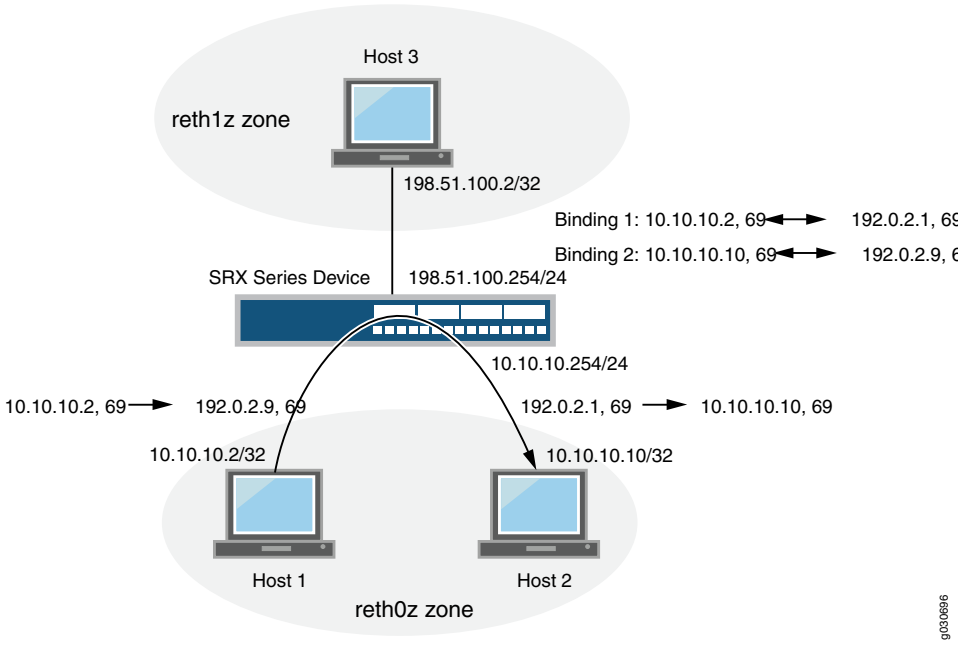

In Figure 219, the following parameters are used:

- Host 1 IP address - **10.10.10.2/32**
- Host 2 IP address - **10.10.10.10/32**
- Intra-zone IP address - **10.10.10.254/24**
- Host 3 IP address - **198.51.100.2/32**
- Inter-zone IP address - **198.51.100.254/24**
- Host 1 and Host 2 are in zone **reht0z**, and Host 3 is in **reth1z** zone

Table 446 shows the binding table used in this example.

Table 446: Persistent NAT Binding Table

Original Source IP Address	Translated Source IP Address
10.10.10.2/32 to 10.10.10.11/32	192.0.2.1/32 to 192.0.2.10/32

Configuration

Step-by-Step Procedure

To configure persistent NAT hairpinning:

1. Configure interfaces.

```
[edit]
user@host# set interfaces ge-11/0/0 unit 0 family inet address 10.10.10.254/24
user@host# set interfaces ge-11/0/1 unit 0 family inet address 198.51.100.254/24
```

Copyright © 2016, Juniper Networks, Inc.

2. Create zones (reth0z and reth1z).

```
[edit]
user@host# set security zones security-zone reth0z host-inbound-traffic
  system-services all
user@host# set security zones security-zone reth0z host-inbound-traffic protocols
  all
user@host# set security zones security-zone reth0z interfaces ge-11/0/0.0
user@host# set security zones security-zone reth1z host-inbound-traffic
  system-services all
user@host# set security zones security-zone reth1z host-inbound-traffic protocols
  all
user@host# set security zones security-zone reth1z interfaces ge-11/0/1.0
```

3. Create policies for zones reth0z and reth1z.

```
[edit]
user@host# set security address-book global address subnet10 10.10.10.0/24
user@host# set security address-book global address subnet 198.51.100.0/24
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match
  source-address subnet10
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match
  destination-address subnet20
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match
  application any
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 then
  permit
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match
  source-address subnet10
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match
  destination-address subnet10
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match
  application any
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 then
  permit
```

4. Add same zone policy to do persistent NAT hairpinning.

```
user@host# set security policies default-policy deny-all
```

5. Create a source NAT pool for Host 1 and Host 2 (src1).

```
[edit]
user@host# set security nat source pool src1 address 192.0.2.1/32 to 192.0.2.10/32
```

6. Specify the beginning of the original source IP address range for Host 1 and Host 2 (src1).

```
[edit]
user@host# set security nat source pool src1 host-address-base 10.10.10.2/32
```

7. Configure the source NAT rule set r1.

```
[edit]
user@host# set security nat source rule-set r1 from zone reth0z
user@host# set security nat source rule-set r1 to zone reth1z
user@host# set security nat source rule-set r1 to zone reth0z
user@host# set security nat source rule-set r1 rule rule1 match source-address
  10.10.10.0/24
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security nat source rule-set r1 rule rule1 match destination-address
  10.10.10.0/24
user@host# set security nat source rule-set r1 rule rule1 match destination-address
  198.51.100.0/24
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool src1
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool
  persistent-nat permit any-remote-host
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool
  persistent-nat inactivity-timeout 900
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool
  persistent-nat max-session-number 20
```

Results From configuration mode, enter the **show security nat** command to confirm your configuration. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src1 {
    address {
      192.0.2.1/32 to 192.0.2.10/32;
    }
    host-address-base 10.10.10.2/32;
  }
  rule-set r1 {
    from zone reth0z;
    to zone [ reth0z reth1z ];
    rule rule1 {
      match {
        source-address 10.10.10.0/24;
        destination-address [10.10.10.0/24 198.51.100.0/24];
      }
      then {
        source-nat {
          pool {
            src1;
            persistent-nat {
              permit any-remote-host;
              inactivity-timeout 900;
              max-session-number 20;
            }
          }
        }
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Traffic Sent Between the Hosts Creating Binding 1

Purpose Verify traffic sent from between the hosts (Host 1 and Host 3) creating binding 1.

Action **user@host>show security nat source persistent-nat-table all**

```
 sendip -d r28 -p ipv4 -iv 4 -is 10.10.10.2 -id 198.51.100.2 -p udp -us 69 -ud
69 198.51.100.2

        Source-IP: 10.10.10.2
        Source-port: 69
        Dst-IP: 198.51.100.2
    Dst-port: 69
        Binding1 is below:

Internal   Reflective       Source            Type          Left_time/
Curr_Sess_Num/ Source
 In_IP     In_Port Ref_IP   Ref_Port NAT Pool                Conf_time
Max_Sess_Num  NAT Rule
10.10.10.2 69      192.0.2.1  69     src1     any-remote-host  -/900      1/20
        rule1
```

Traffic Sent Between the Hosts Creating Binding 2

Purpose Verify traffic sent from between the hosts (Host 2 and Host 3) creating binding 2.

Action **user@host>show security nat source persistent-nat-table all**

```
sendip -d r28 -p ipv4 -iv 4 -is 10.10.10.10 -id 198.51.100.2 -p udp -us 69 -ud
69 198.51.100.2
        Source-IP: 10.10.10.10
        Source-port: 69
        Dst-IP: 198.51.100.2
        Dst-port: 69
        Binding2 is below:

Internal   Reflective          Source            Type            Left_time/
Curr_Sess_Num/  Source
 In_IP       In_Port Ref_IP    Ref_Port NAT Pool                  Conf_time
Max_Sess_Num   NAT Rule
10.10.10.2   69      192.0.2.1  69       src1     any-remote-host   -/900
1/20          rule1
10.10.10.10  69      192.0.2.9  69       src1     any-remote-host   -/900
1/20          rule1
```

Traffic Sent Between Two Hosts

Purpose Verify the traffic sent from Host 1 to Host 2:

Copyright © 2016, Juniper Networks, Inc.

Action

user@host>show security flow session

```
sendip -d r28 -p ipv4 -iv 4 -is 10.10.10.2 -id 192.0.2.9 -p udp -us 69 -ud 69
192.0.2.9

  Session ID: 100007628, Policy name: default-policy/2, Timeout: 52, Valid
  In: 10.10.10.2/69 --> 192.0.2.9/69;udp, If: ge-0/0/0.0, Pkts: 2, Bytes: 112
  Out: 10.10.10.10/69 --> 192.0.2.1/69;udp, If: ge-0/0/0.0, Pkts: 0, Bytes: 0
   Total sessions: 1
```

Related Documentation

- Persistent NAT Hairpinning Overview on page 4993

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 206

Configuring NAT for Multicast Flows

- Understanding NAT for Multicast Flows on page 5001
- Example: Configuring NAT for Multicast Flows on page 5002

Understanding NAT for Multicast Flows

Network Address Translation (NAT) can be used to translate source addresses in IPv4 multicast flows and to translate IPv4 multicast group destination addresses.

Either static NAT or destination NAT can be used to perform multicast group address translation. Static NAT allows connections to be originated from either side of the network, but translation is limited to one-to-one addresses or between blocks of addresses of the same size. No address pools are necessary. Use the **static** configuration statement at the [**edit security nat**] hierarchy level to configure static NAT rule sets for multicast traffic. Destination NAT allows connections to be initiated only for incoming network connections—for example, from the Internet to a private network. Use the **destination** configuration statement at the [**edit security nat**] hierarchy level to configure destination NAT pools and rule sets.

Source NAT for multicast traffic is supported only by using IP address shifting to translate the original source IP address to an IP address from a user-defined address pool. This type of translation is one-to-one, static, and without port address translation. If the original source IP address range is larger than the IP address range in the user-defined pool, untranslated packets are dropped. The mapping does not provide bidirectional mapping, which static NAT provides. Use the **source** configuration statement at the [**edit security nat**] hierarchy level to configure source NAT pools and rule sets. When you define the source NAT pool for this type of source NAT, use the **host-address-base** option to specify the start of the original source IP address range.

Related Documentation

- Understanding Destination NAT on page 4923
- Understanding Destination NAT on page 4923
- Understanding Source NAT Pools with Address Shifting on page 4901
- Example: Configuring NAT for Multicast Flows on page 5002

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring NAT for Multicast Flows

This example shows how to configure a Juniper Networks device for address translation of multicast flows.

- Requirements on page 5002
- Overview on page 5002
- Configuration on page 5004
- Verification on page 5009

Requirements

Before you begin:

1. Configure network interfaces on the device. See the *Interfaces Feature Guide for Security Devices*.
2. Create security zones and assign interfaces to them. See "Understanding Security Zones" on page 1050.
3. Configure the device for multicast forwarding. See *Multicast Feature Guide for Security Devices*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. Figure 220 depicts a typical deployment of the Juniper Networks device for multicast forwarding. The source router R1 sends multicast packets with source addresses in the range 203.0.113.100 through 203.0.113.110 and the group address 233.252.0.1/32 toward the Juniper Networks device. The source router R1 is in the private network (trust zone) upstream of the Juniper Networks device. There are several receivers in the public network (untrust zone) downstream of the device.

The Juniper Networks device translates incoming multicast packets from R1 before forwarding them out on the downstream interfaces. The following translations are applied:

- For the interface to R2, the source address is untranslated, and the group address is translated to 233.252.0.2/32.
- For the interface to R3, the source address is translated to an address in the range 198.51.100.200 through 198.51.100.210, and the group address is translated to 233.252.0.2/32.
- For the interface to R4, the source address is translated to an address in the range 10.10.10.100 through 10.10.10.110, and the group address is translated to 233.252.0.2/32.

Copyright © 2016, Juniper Networks, Inc.

Figure 220: NAT Translations for Multicast Flows

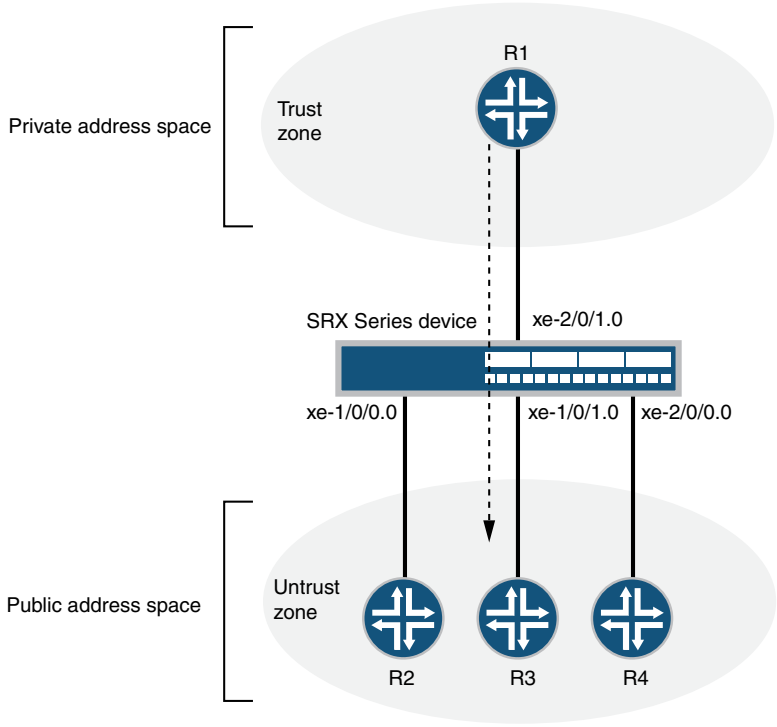

From R1	To R2	To R3	To R4
Original Group IP	Group IP	Group IP	Group IP
233.252.0.1/32	233.252.0.2/32	233.252.0.2/32	233.252.0.2/32
Original Source IP	Source IP	Source IP	Source IP
203.0.113.100 - 203.0.113.110	203.0.113.100 - 203.0.113.110	198.51.100.200 - 198.51.100.210	10.10.10.100 - 10.10.10.110

g030689

This example describes the following configurations:

- Destination NAT pool **dst-nat-pool** that contains the IP address 233.252.0.2/32.
- Destination NAT rule set **rs1** with rule **r1** to match packets arriving on interface xe-2/0/1.0 with the destination IP address 233.252.0.1/32. For matching packets, the destination address is translated to the IP address in the **dst-nat-pool** pool.
- Source NAT pool **src-nat-shift-1** that contains the IP address range 198.51.100.200/32 through 198.51.100.210/32. For this pool, the beginning of the original source IP address range is 203.0.113.100/32 and is specified with the **host-address-base** option.

Copyright © 2016, Juniper Networks, Inc.

- Source NAT rule set **rs-shift1** with rule **r1** to match packets from the trust zone to interface xe-1/0/1.0 with a source IP address in the 203.0.113.96/28 subnet. For matching packets that fall within the source IP address range specified by the **src-nat-shift-1** configuration, the source address is translated to the IP address in the **src-nat-shift-1** pool.
- Source NAT pool **src-nat-shift-2** that contains the IP address range 10.10.10.100/32 through 10.10.10.110/32. For this pool, the beginning of the original source IP address range is 1203.0.113.100/32 and is specified with the **host-address-base** option.
- Source NAT rule set **rs-shift2** with rule **r1** to match packets from the trust zone to interface xe-2/0/0.0 with a source IP address in the 203.0.113.96/28 subnet. For matching packets that fall within the source IP address range specified by the **src-nat-shift-2** configuration, the source address is translated to the IP address in the **src-nat-shift-2** pool.
- Proxy ARP for the addresses 203.0.113.100 through 203.0.113.110 on interface xe-1/0/0.0, addresses 198.51.100.200 through 198.51.100.210 on interface xe-1/0/1.0, and addresses 10.10.10.100 through 10.10.10.110 on interface xe-2/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policy to permit traffic from the trust zone to the untrust zone.
- Security policy to permit traffic from the untrust zone to the translated destination IP address in the trust zone.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat source pool src-nat-shift-1 address 5198.51.100.200/32 to
  198.51.100.210/32
set security nat source pool src-nat-shift-1 host-address-base 203.0.113.100/32
set security nat source pool src-nat-shift-2 address 10.10.10.100/32 to 10.10.10.110/32
set security nat source pool src-nat-shift-2 host-address-base 203.0.113.100/32
set security nat source rule-set rs-shift1 from zone trust
set security nat source rule-set rs-shift1 to interface xe-1/0/1.0
set security nat source rule-set rs-shift1 rule r1 match source-address 203.0.113.96/28
set security nat source rule-set rs-shift1 rule r1 then source-nat pool src-nat-shift1
set security nat source rule-set rs-shift2 from zone trust
set security nat source rule-set rs-shift2 to interface xe-2/0/0.0
set security nat source rule-set rs-shift2 rule r2 match source-address 203.0.113.96/28
set security nat source rule-set rs-shift2 rule r2 then source-nat pool src-nat-shift2
set security nat destination pool dst-nat-pool address 233.252.0.1/32
set security nat destination rule-set rs1 from interface xe-2/0/1.0
set security nat destination rule-set rs1 rule r1 match destination-address 233.252.0.1/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool
set security nat proxy-arp interface xe-1/0/0.0 address 203.0.113.100/32 to
  203.0.113.110/32
```


Copyright © 2016, Juniper Networks, Inc.

```
set security nat proxy-arp interface xe-1/0/1.0 address 198.51.100.200/32 to
  198.51.100.210/32
set security nat proxy-arp interface xe-2/0/0.0 address 10.10.10.100/32 to 10.10.10.110/32
set security policies from-zone trust to-zone untrust policy internet-access match
  source-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match
  application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match
  source-address any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match
  destination-address 233.252.0.1/21
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match
  application any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access then
  permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the destination and source NAT translations for multicast flows:

1. Create a destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set pool dst-nat-pool address 233.252.0.1/32
   ```

2. Create a destination NAT rule set.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 from interface xe-2/0/1.0
   ```

3. Configure a rule that matches packets and translates the destination address to the address in the destination NAT pool.

   ```
   [edit security nat destination]
   user@host# set rule-set rs1 rule r1 match destination-address 233.252.0.1/32
   user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool
   ```

4. Create a source NAT pool.

   ```
   [edit security nat source]
   user@host# set pool src-nat-shift-1 address 198.51.100.200 to 198.51.100.210
   ```

5. Specify the beginning of the original source IP address range.

   ```
   [edit security nat source]
   user@host# set pool src-nat-shift-1 host-address-base 203.0.113.100
   ```

6. Create a source NAT rule set.

   ```
   [edit security nat source]
   user@host# set rule-set rs-shift1 from zone trust
   user@host# set rule-set rs-shift1 to interface xe-1/0/1.0
   ```

7. Configure a rule that matches packets and translates the destination address to the address in the source NAT pool.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security nat source]
user@host# set rule-set rs-shift1 rule r1 match source-address 203.0.113.96/28
user@host# set rule-set rs-shift1 rule r1 then source-nat pool src-nat-shift1
```

8. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-shift-2 address 10.10.10.100 to 10.10.10.110
```

9. Specify the beginning of the original source IP address range.

```
[edit security nat source]
user@host# set pool src-nat-shift-2 host-address-base 203.0.113.100/32
```

10. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs-shift2 from zone trust
user@host# set rule-set rs-shift2 to interface xe-2/0/0.0
```

11. Configure a rule that matches packets and translates the destination address to the address in the source NAT pool.

```
[edit security nat source]
user@host# set rule-set rs-shift2 rule r2 match source-address 203.0.113.96/28
user@host# set rule-set rs-shift2 rule r2 then source-nat pool src-nat-shift2
```

12. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface xe-1/0/0.0 address 203.0.113.100 to 203.0.113.110
user@host# set proxy-arp interface xe-1/0/1.0 address 198.51.100.200 to
  198.51.100.210
user@host# set proxy-arp interface xe-2/0/0.0 address 10.10.10.100 to 10.10.10.110
```

13. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any
  destination-address any application any
user@host# set policy internet-access then permit
```

14. Configure a security policy that allows traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy dst-nat-pool-access match source-address any
  destination-address 233.252.0.1/32 application any
user@host# set policy dst-nat-pool-access then permit
```

Results From configuration mode, confirm your configuration by entering the **show security nat** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-shift-1 {
    address {
      198.51.100.200/32 to 198.51.100.210/32;
    }
```

 Copyright © 2016, Juniper Networks, Inc.

```
        host-address-base 203.0.113.100/32;
      }
      pool src-nat-shift-2 {
        address {
          10.10.10.100/32 to 10.10.10.110/32;
        }
        host-address-base 203.0.113.100/32;
      }
      rule-set trust-to-untrust {
        from zone trust;
        to zone untrust;
        rule source-nat-rule {
          match {
            source-address 0.0.0.0/0;
          }
          then {
            source-nat {
              interface;
            }
          }
        }
      }
      rule-set rs-shift1 {
        from zone trust;
        to interface xe-1/0/1.0;
        rule r1 {
          match {
            source-address 203.0.113.96/28;
          }
          then {
            source-nat {
              pool {
                src-nat-shift1;
              }
            }
          }
        }
      }
      rule-set rs-shift2 {
        from zone trust;
        to interface xe-2/0/0.0;
        rule r2 {
          match {
            source-address 203.0.113.96/28;
          }
          then {
            source-nat {
              pool {
                src-nat-shift2;
              }
            }
          }
        }
      }
    }
    destination {
```

```
    pool dst-nat-pool {
      address 233.252.0.1/32;
    }
    rule-set rs1 {
      from interface xe-2/0/1.0;
      rule r1 {
        match {
          destination-address 233.252.0.1/32;
        }
        then {
          destination-nat pool dst-nat-pool;
        }
      }
    }
  }
  proxy-arp {
    interface xe-1/0/0.0 {
      address {
        203.0.113.100/32 to 203.0.113.110/32;
      }
    }
    interface xe-1/0/1.0 {
      address {
        198.51.100.200/32 to 198.51.100.210/32;
      }
    }
    interface xe-2/0/0.0 {
      address {
        10.10.10.100/32 to 10.10.10.110/32;
      }
    }
  }

[edit]
user@host# show security policies
from-zone trust to-zone untrust {
    policy trust-to-untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
    policy internet-access {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    from-zone untrust to-zone trust {
      policy dst-nat-pool-access {
        match {
          source-address any;
          destination-address 233.252.0.1/21;
          application any;
        }
        then {
          permit;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Destination NAT Pool Usage on page 5009
- Verifying Destination NAT Rule Usage on page 5009
- Verifying Source NAT Pool Usage on page 5009
- Verifying Source NAT Rule Usage on page 5009
- Verifying NAT Application to Traffic on page 5010

Verifying Destination NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the destination NAT pool.

Action From operational mode, enter the **show security nat destination pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose Verify that there is traffic matching the destination NAT rule.

Action From operational mode, enter the **show security nat destination rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying Source NAT Pool Usage

Purpose Verify that there is traffic using IP addresses from the source NAT pool.

Action From operational mode, enter the **show security nat source pool all** command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose Verify that there is traffic matching the source NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security nat source rule all** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose Verify that NAT is being applied to the specified traffic.

Action From operational mode, enter the **show security flow session** command.

Related Documentation

- Understanding NAT for Multicast Flows on page 5001
- Source NAT Configuration Overview on page 4864
- Destination NAT Configuration Overview on page 4926

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 207

Configuring IPv6 NAT

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026

IPv6 NAT Overview

IPv6 has a vastly larger address space than the impending exhausted IPv4 address space. IPv4 has been extended using techniques such as Network Address Translation (NAT), which allows for ranges of private addresses to be represented by a single public address, and temporary address assignment. There are a lot of technologies to provide the transition mechanism for the legacy IPv4 host to keep the connection to the Internet. IPv6 NAT provides address translation between IPv4 and IPv6 addressed network devices. It also provides address translation between IPv6 hosts. NAT between IPv6 hosts is done in a similar manner and for similar purposes as IPv4 NAT.

IPv6 NAT in Junos OS provides the following NAT types:

- Source NAT
- Destination NAT
- Static NAT

Source NAT Translations Supported by IPv6 NAT

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

Copyright © 2016, Juniper Networks, Inc.

IPv6 NAT in Junos OS supports the following source NAT translations:

- Translation of one IPv6 subnet to another IPv6 subnet without port address translation
- Translation of IPv4 addresses to IPv6 prefix + IPv4 addresses
- Translation of IPv6 hosts to IPv6 hosts with or without port address translation
- Translation of IPv6 hosts to IPv4 hosts with or without port address translation
- Translation of IPv4 hosts to IPv6 hosts with or without port address translation

Destination NAT Mappings Supported by IPv6 NAT

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

IPv6 NAT in Junos OS supports the following destination NAT translations:

- Prefix translation between IPv4 and IPv6 prefix
- Mapping of one IPv6 subnet to another IPv6 subnet
- Mapping of one IPv6 subnet to an IPv6 host
- Mapping of one IPv6 subnet to one IPv4 subnet
- Mapping of one IPv4 subnet to one IPv6 subnet
- Mapping of one IPv6 host (and optional port number) to one special IPv6 host (and optional port number)
- Mapping of one IPv6 host (and optional port number) to one special IPv4 host (and optional port number)
- Mapping of one IPv4 host (and optional port number) to one special IPv6 host (and optional port number)

Static NAT Mappings Supported by IPv6 NAT

Static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the reverse direction. From the NAT device, the original destination address is the virtual host IP address while the mapped-to address is the real host IP address.

IPv6 NAT in Junos OS supports the following static NAT translations:

- Translation of one IPv6 subnet to another IPv6 subnet
- Translation of one IPv6 host to another IPv6 host
- Translation of one IPv4 address a.b.c.d to IPv6 address Prefix::a.b.c.d
- Translation of IPv4 hosts to IPv6 hosts

Copyright © 2016, Juniper Networks, Inc.

See "Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping" on page 5015.

- Translation of IPv6 hosts to IPv4 hosts

  See "Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping" on page 5022.

- Mapping of one IPv6 prefix to one IPv4 prefix

  See "Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping" on page 5026.

- Mapping of one IPv4 prefix to one IPv6 prefix

  See "Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping" on page 5018.

Related Documentation

- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014

IPv6 NAT PT Overview

IPv6 Network Address Translation- Protocol Translation (NAT-PT) provides address allocation and protocol translation between IPv4 and IPv6 addressed network devices. The translation process is based on the Stateless IP/ICMP Translation (SIIT) method; however, the state and the context of each communication are retained during the session lifetime. IPv6 NAT-PT supports Internet Control Message Protocol (ICMP), TCP, and UDP packets.

IPv6 NAT-PT supports the following types of NAT-PT:

- Traditional NAT-PT— In traditional NAT-PT, the sessions are unidirectional and outbound from the IPv6 network . Traditional NAT-PT allows hosts within an IPv6 network to access hosts in an IPv4 network. There are two variations to traditional NAT-PT: basic NAT-PT and NAPT-PT.

  In basic NAT-PT, a block of IPv4 addresses at an IPv4 interface is set aside for translating addresses as IPv6 hosts as they initiate sessions to the IPv4 hosts. The basic NAT-PT translates the source IP address and related fields such as IP, TCP, UDP, and ICMP header checksums for packets outbound from the IPv6 domain . For inbound packets, it translates the the destination IP address and the checksums.

  Network Address Port Translation and Protocol Translation (NAPT-PT) can be combined with basic NAT-PT so that a pool of external addresses is used in conjunction with port translation. NAPT-PT allows a set of IPv6 hosts to share a single IPv4 address. NAPT-PT translates the source IP address, source transport identifier, and related fields such as IP, TCP, UDP, and ICMP header checksums, for packets outbound from the IPv6 network. The transport identifier can be a TCP/UDP port or an ICMP query ID. For inbound packets, it translates the destination IP address, destination transport identifier, and the IP and the transport header checksums.

Copyright © 2016, Juniper Networks, Inc.

- Bidirectional NAT-PT— In bidirectional NAT-PT, sessions can be initiated from hosts in the IPv4 network as well as the IPv6 network. IPv6 network addresses are bound to IPv4 addresses, either statically or dynamically as connections are established in either direction. The static configuration is similar to static NAT translation. Hosts in IPv4 realm access hosts in the IPv6 realm using DNS for address resolution. A DNS ALG must be employed in conjunction with bidirectional NAT-PT to facilitate name-to-address mapping. Specifically, the DNS ALG must be capable of translating IPv6 addresses in DNS queries and responses into their IPv4 address bindings, and vice versa, as DNS packets traverse between IPv6 and IPv4 realms.

NOTE: The SRX Series devices partially supports the Bidirectional NAT-PT specification. It supports flow of bidirectional traffic assuming that there are other ways to convey the mapping between the IPv6 address and the dynamically allocated IPv4 address. For example, a local DNS can be configured with the mapped entries for IPv4 nodes to identify the addresses.

NAT- PT Operation—The SRX Series devices support the Traditional NAT-PT and allows static mapping for the user to communicate from IPv4 to IPv6 . The user needs to statically configure the DNS server with an IPv4 address for the host name and then create a static NAT on the device for the IPv6-only node to communicate from an IPv4-only node to an IPv6-only node based on the DNS.

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT-PT Communication Overview on page 5014
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015

IPv6 NAT-PT Communication Overview

NAT-PT communication with static mapping— Network Address Translation-Protocol Translation (NAT-PT) can be done in two directions, from IPv6 to IPv4 and vice versa. For each direction, static NAT is used to map the destination host to a local address and a source address NAT is used to translate the source address. There are two types of static NAT and source NAT mapping: one-to-one mapping and prefix-based mapping.

NAT- PT communication with DNS ALG—A DNS-based mechanism dynamically maps IPv6 addresses to IPv4-only servers. NAT-PT uses the DNS ALG to transparently do the translations. For example, a company using an internal IPv6 network needs to be able to communicate with external IPv4 servers that do not yet have IPv6 addresses.

Copyright © 2016, Juniper Networks, Inc.

To support the dynamic address binding, a DNS should be used for name resolution. The IPv4 host looks up the name of the IPv6 node in its local configured IPv4 DNS server, which then passes the query to the IPv6 DNS server through an SRX Series device using NAT-PT.

The DNS ALG in NAT device :

- Translates the IPv6 address resolution back to IPv4 address resolution.
- Allocates an IPv6 address for the mapping.
- Stores a mapping of the allocated IPv4 address to the IPv6 address returned in the IPv6 address resolution so that the session can be established from any-IPv4 hosts to the IPv6 host.

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping

This example shows how to configure an IPv4-initiated connection to an IPv6 node using default destination address prefix static mapping.

- Requirements on page 5015
- Overview on page 5015
- Configuration on page 5016
- Verification on page 5018

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes how to configure an IPv4-initiated connection to an IPv6 node that has a static mapping /96-based IPv6 address defined on its interface and static mapping /96 set up on the device. This example assumes the IPv6 addresses to be mapped IPv4 addresses, making the IPv4 addresses a part of the IPv6 address space.

Copyright © 2016, Juniper Networks, Inc.

Configuring an IPv4-initiated connection to an IPv6 node is useful when the devices on the IPv4 network must be interconnected to the devices on the IPv6 network and during migration of an IPv4 network to an IPv6 network. The mapping can be used for DNS ALG for reverse lookup of IPv4 addresses from IPv6 addresses, for the traffic initiated from the IPv6 network. This also provides connectivity for sessions initiated from IPv4 nodes with IPv6 nodes on the other side of the NAT/PT device.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 10.1.1.45/30
set security nat static rule-set test_rs rule test_rule then static-nat prefix 2001:db8::/96
set security nat source pool myipv6_prefix address 2001:db8::/96
set security nat source rule-set myipv6_rs from interface ge-0/0/1
set security nat source rule-set myipv6_rs to interface ge-0/0/2
set security nat source rule-set myipv6_rs rule ipv6_rule match source-address 10.1.1.1/30
set security nat source rule-set myipv6_rs rule ipv6_rule match destination-address
  2001:db8:a0a:a2d/126
set security nat source rule-set myipv6_rs rule ipv6_rule then source-nat pool myipv6_prefix
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs from interface ge-0/0/1
   ```

2. Define the rule to match the destination address prefix.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule match destination-address 10.1.1.45/30
   ```

3. Define the static NAT prefix for the device.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule then static-nat prefix 2001:db8::/96
   ```

4. Configure the source NAT pool with an IPv6 address prefix.

   ```
   [edit security nat source]
   user@host# set pool myipv6_prefix address 2001:db8::/96
   ```

5. Configure the source NAT rule set for the interface.

   ```
   [edit security nat source]
   user@host# set rule-set myipv6_rs from interface ge-0/0/1
   user@host# set rule-set myipv6_rs to interface ge-0/0/2
   ```


Copyright © 2016, Juniper Networks, Inc.

6. Configure the IPv6 source NAT source address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match source-address 10.1.1.1/30
```

7. Configure the IPv6 source NAT destination address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match destination-address 2001:db8:a0a:a2d/126
```

8. Define the configured source NAT IPv6 pool in the rule.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule then source-nat pool myipv6_prefix
```

Results From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool myipv6_prefix {
    address {
      2001:db8::/96;
    }
  }
  rule-set myipv6_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv6_rule {
      match {
        source-address 10.1.1.1/30;
        destination-address 2001:db8:a0a:a2d/126;
      }
      then {
        source-nat {
          pool {
            myipv6_prefix;
          }
        }
      }
    }
  }
}
static {
  rule-set test_rs {
    from interface ge-0/0/1.0;
    rule test_rule {
      match {
        destination-address 10.1.1.45/30;
      }
      then {
        static-nat prefix 2001:db8::/96;
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Static NAT Is Configured on page 5018
- Verifying That Source NAT Is Configured on page 5018

Verifying That Static NAT Is Configured

Purpose Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action From operational mode, enter the **show security nat static** command.

Verifying That Source NAT Is Configured

Purpose Verify whether source NAT is configured.

Action From operational mode, enter the **show security nat source** command.

Related Documentation

- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping

This example shows how to configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping.

- Requirements on page 5018
- Overview on page 5019
- Configuration on page 5019
- Verification on page 5021

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Copyright © 2016, Juniper Networks, Inc.

Overview

The following example describes how to configure an IPv4 node to communicate with an IPv6 node using one-to-one static NAT on the device.

The communication of an IPv4 node with an IPv6 node is useful for IPv4 hosts accessing an IPv6 server, for new servers that support IPv6 only and that need to be connected to the IPv6 network, and for migrating of old hosts to the new server when most of the machines have already moved to IPv6. For example, you can use this feature to connect an IPv4-only node to an IPv6-only printer. This mapping can also be used for DNS ALG for reverse lookup of IPv4 addresses from IPv6 addresses for traffic that is initiated from the IPv6 network.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 10.1.1.25/32
set security nat static rule-set test_rs rule test_rule then static-nat prefix 2001:db8::25/128
set security nat source pool myipv6_prefix address 2001:db8::/96
set security nat source rule-set myipv6_rs from interface ge-0/0/1
set security nat source rule-set myipv6_rs to interface ge-0/0/2
set security nat source rule-set myipv6_rs rule ipv6_rule match source-address
  10.10.10.1/30
set security nat source rule-set myipv6_rs rule ipv6_rule match destination-address
  2001:db8::25
set security nat source rule-set myipv6_rs rule ipv6_rule then source-nat pool myipv6_prefix
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs from interface ge-0/0/1
   ```

2. Define the rule and the destination address.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule match destination-address 10.1.1.25/32
   ```

3. Define the static NAT prefix.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule then static-nat prefix 2001:db8::25/128
   ```

4. Configure a source NAT pool with an IPv6 prefix address.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security]
user@host# set nat source pool myipv6_prefix address 2001:db8::/96
```

5. Configure the source NAT rule set.

```
[edit security nat source]
user@host# set rule-set myipv6_rs from interface ge-0/0/1
user@host# set rule-set myipv6_rs from interface ge-0/0/2
```

6. Configure the source NAT source address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match source-address 10.10.10.1/30
```

7. Configure the source NAT destination address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match destination-address 2001:db8::25
```

8. Define a configured source NAT IPv6 pool in the rule.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule then source-nat pool myipv6_prefix
```

Results From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    pool myipv6_prefix {
        address {
            2001:db8::/96;
        }
    }
    rule-set myipv6_rs {
        from interface ge-0/0/1.0;
        to interface ge-0/0/2.0;
        rule ipv6_rule {
            match {
                source-address 10.10.10.1/30;
                destination-address 2001:db8::25;
            }
            then {
                source-nat {
                    pool {
                        myipv6_prefix;
                    }
                }
            }
        }
    }
}
static {
    rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
```

 Copyright © 2016, Juniper Networks, Inc.

```
            match {
                destination-address 10.1.1.25/32;
            }
            then {
                static-nat prefix 2001:db8::25/128;
            }
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Static NAT Is Configured on page 5021
- Verifying That Source NAT Is Configured on page 5021

Verifying That Static NAT Is Configured

Purpose Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action From operational mode, enter the **show security nat static** command.

Verifying That Source NAT Is Configured

Purpose Verify whether source NAT is configured.

Action From operational mode, enter the **show security nat source** command.

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping

This example shows how to configure an IPv6-initiated connection to an IPv4 node using default destination address prefix static mapping. This example does not show how to configure the NAT translation for the reverse direction.

- Requirements on page 5022
- Overview on page 5022
- Configuration on page 5022
- Verification on page 5024

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes the communication of an IPv6 node with an IPv4 node that has prefix-based static NAT defined on the device. The static NAT assumes that the IPv4 network is a special IPv6 network (that is, an IPv4-mapped IPv6 network), and hides the entire IPv4 network behind an IPv6 prefix.

The communication of an IPv6 node with an IPv4 node is useful when IPv6 is used in the network and must be connected to the IPv4 network, or when both IPv4 and IPv6 are used in the network and a mechanism is required to interconnect the two networks during migration. This also provides connectivity for sessions initiated from IPv6 nodes with IPv4 nodes on the other side of the NAT/PT device.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address
  2001:db8::1/96
set security nat static rule-set test_rs rule test_rule then static-nat inet
set security nat source pool myipv4 address 203.0.113.2 to 203.0.113.5
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address
  10.1.1.15/30
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address
  2001:db8::2/96
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv6-initiated connection to an IPv4 node using default destination address prefix static mapping:

1. Configure the static NAT for an interface.

   ```
   [edit security nat static]
   user@host# set rule test_rs from interface ge-0/0/1
   ```

2. Define the rule and destination address with the prefix for the static NAT translation defined on the device.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule match destination-address 2001:db8::1/96
   ```

3. Define the static NAT as inet to translate to an IPv4 address.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule then static-nat inet
   ```

4. Configure the IPv4 source NAT pool address.

   ```
   [edit security nat source]
   user@host# set pool myipv4 address 203.0.113.2 to 203.0.113.5
   ```

5. Configure the source NAT rule set.

   ```
   [edit security nat source ]
   user@host# set rule-set myipv4_rs from interface ge-0/0/1
   user@host# set rule-set myipv4_rs from interface ge-0/0/2
   ```

6. Configure the IPv4 source NAT destination address.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# set rule ipv4_rule match destination-address 10.1.1.15/30
   ```

7. Define the source address with the prefix for the source NAT defined on the device.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# set rule ipv4_rule match source-address 2001:db8::/96
   ```

8. Define a configured source NAT IPv4 pool in the rule.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# sset rule ipv4_rule then source-nat pool myipv4
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool myipv4 {
    address {
      1203.0.113.2/32 to 203.0.113.5/32;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    rule-set myipv4_rs {
      from interface ge-0/0/1.0;
      to interface ge-0/0/2.0;
      rule ipv4_rule {
        match {
          source-address 2001:db8::/96;
          destination-address 10.1.1.15/30;
        }
        then {
          source-nat {
            pool {
              myipv4;
            }
          }
        }
      }
    }
  }
  static {
    rule-set test_rs {
      from interface ge-0/0/1.0;
      rule test_rule {
        match {
          destination-address 2001:db8::1/96;
        }
        then {
          static-nat inet;
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Static NAT Is Configured on page 5024
- Verifying That Source NAT Is Configured on page 5025

Verifying That Static NAT Is Configured

Purpose Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action From operational mode, enter the **show security nat static rule** command.

```
user@host> show security nat static rule test_rule
Static NAT rule: test_rule            Rule-set: test_rs
  Rule-Id                   : 2
  Rule position             : 2
  From interface            : ge-0/0/1.0
  Destination addresses     : 2001:db8::1
```


Copyright © 2016, Juniper Networks, Inc.

```
  Host addresses             : 0.0.0.0
  Netmask                    : 96
  Host routing-instance      : N/A
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

Verifying That Source NAT Is Configured

Purpose Verify whether source NAT is configured.

Action From operational mode, enter the **show security nat source rule** command.

```
user@host> show security nat source rule ipv4_rule
source NAT rule: ipv4_rule            Rule-set: myipv4_rs
  Rule-Id                    : 2
  Rule position              : 2
  From interface             : ge-0/0/1.0
  To interface               : ge-0/0/2.0
  Match
    Source addresses         : 2001:db8:: - 2001:db8::ffff:ffff
    Destination addresses    : 10.1.1.15       - 10.1.1.15
  Action                       : myipv4
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits           : 0
    Successful sessions      : 0
    Failed sessions          : 0
  Number of sessions         : 0
```

From operational mode, enter the **show security nat source pool** command.

```
user@host> show security nat source pool myipv4
Pool name          : myipv4
Pool id            : 5
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 65535]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 4
Translation hits   : 0
Address range                              Single Ports   Twin Ports
            203.0.113.2 - 203.0.113.5             0              0
```

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping on page 5026
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping

This example shows how to configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping.

- Requirements on page 5026
- Overview on page 5026
- Configuration on page 5026
- Verification on page 5028

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes the communication of an IPv6 node with an IPv4 node that has a one-to-one static NAT address defined on the device. The communication of an IPv6 node with an IPv4 node allows IPv6 hosts to access an IPv4 server when neither of the devices has a dual stack and must depend on the NAT/PT device to communicate. This enables some IPv4 legacy server applications to work even after the network has migrated to IPv6.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security nat static rule test_rs from interface ge-0/0/1
set security nat static rule test_rs rule test_rule match destination-address
  2001:db8::15/128
set security nat static rule test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2 to 203.0.113.3
set security nat source rule myipv4_rs from interface ge-0/0/1
set security nat source rule myipv4_rs to interface ge-0/0/2
set security nat source rule myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
set security nat source rule myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
set security nat source rule myipv4_rs rule ipv4_rule then source-nat pool myipv4
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

   ```
   [edit security nat static]
   user@host# set rule-set test_rs from interface ge-0/0/1
   ```

2. Define a rule to match the destination address.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule match destination-address 2001:db8::15/128
   ```

3. Define the static NAT prefix to the rule.

   ```
   [edit security nat static rule-set test_rs]
   user@host# set rule test_rule then static-nat prefix 10.2.2.15/32
   ```

4. Configure a source NAT pool with an IPv4 addresses.

   ```
   [edit security nat]
   user@host# set source pool myipv4 address 203.0.113.2 203.0.113.3
   ```

5. Configure the IPv4 address for the interface.

   ```
   [edit security nat source ]
   user@host# set rule-set myipv4_rs from interface ge-0/0/1
   ```

6. Configure the source address to the IPv4 source NAT address.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# set rule ipv4_rule match source-address 2001:db8::/96
   ```

7. Configure the destination address to IPv4 source NAT address.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# set rule ipv4_rule match destination-address 10.2.2.15
   ```

8. Define the configured source NAT IPv4 pool in the rule.

   ```
   [edit security nat source rule-set myipv4_rs]
   user@host# set rule ipv4_rule then source-nat pool myipv4
   ```

Results

From configuration mode, confirm your configuration by entering the **show security nat** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
   pool myipv4 {
      address {
         203.0.113.2/32 to 203.0.113.3/32;
      }
   }
   rule-set myipv4_rs {
```

Copyright © 2016, Juniper Networks, Inc.

```
        from interface ge-0/0/1.0;
        to interface ge-0/0/2.0;
        rule ipv4_rule {
          match {
            source-address 2001:db8::/96;
            destination-address 10.2.2.15/32;
          }
          then {
            source-nat {
              pool {
                myipv4;
              }
            }
          }
        }
      }
    }
    static {
      rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
          match {
            destination-address 2001:db8::15/128;
          }
          then {
            static-nat prefix 10.2.2.15/32;
          }
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Static NAT Is Configured on page 5028
- Verifying That Source NAT Is Configured on page 5028

Verifying That Static NAT Is Configured

Purpose Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action From operational mode, enter the **show security nat static** command.

Verifying That Source NAT Is Configured

Purpose Verify whether source NAT is configured.

Action From operational mode, enter the **show security nat source** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPv6 NAT Overview on page 5011
- IPv6 NAT PT Overview on page 5013
- IPv6 NAT-PT Communication Overview on page 5014
- Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping on page 5022
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping on page 5018
- Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping on page 5015

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 208

Configuring IPv6 Dual-Stack

- Understanding IPv6 Dual-Stack Lite on page 5031
- Example: Configuring IPv6 Dual-Stack Lite on page 5034

Understanding IPv6 Dual-Stack Lite

IPv6 dual-stack lite (DS-Lite) is a technology that enables Internet service providers to move to an IPv6 network while simultaneously handling IPv4 address depletion.

IPv4 addresses are becoming depleted; therefore, broadband service providers (DSL, cable, and mobile) need new addresses to support new users. Providing IPv6 addresses alone is often not workable because most of the systems that make up the public Internet are still enabled and support only IPv4, and many users' systems do not yet fully support IPv6.

DS-Lite allows service providers to migrate to an IPv6 access network without changing end-user software. The device that accesses the Internet remains the same, thus allowing IPv4 users to continue accessing IPv4 internet content with minimum disruption to their home networks, while enabling IPv6 users to access IPv6 content.

Figure 91 illustrates the DS-Lite architecture which uses IPv6-only links between the provider and the user while maintaining the IPv4 (or dual-stack) hosts in the user network.

Copyright © 2016, Juniper Networks, Inc.

Figure 221: DS-Lite NAT (IPv4-in-IPv6)

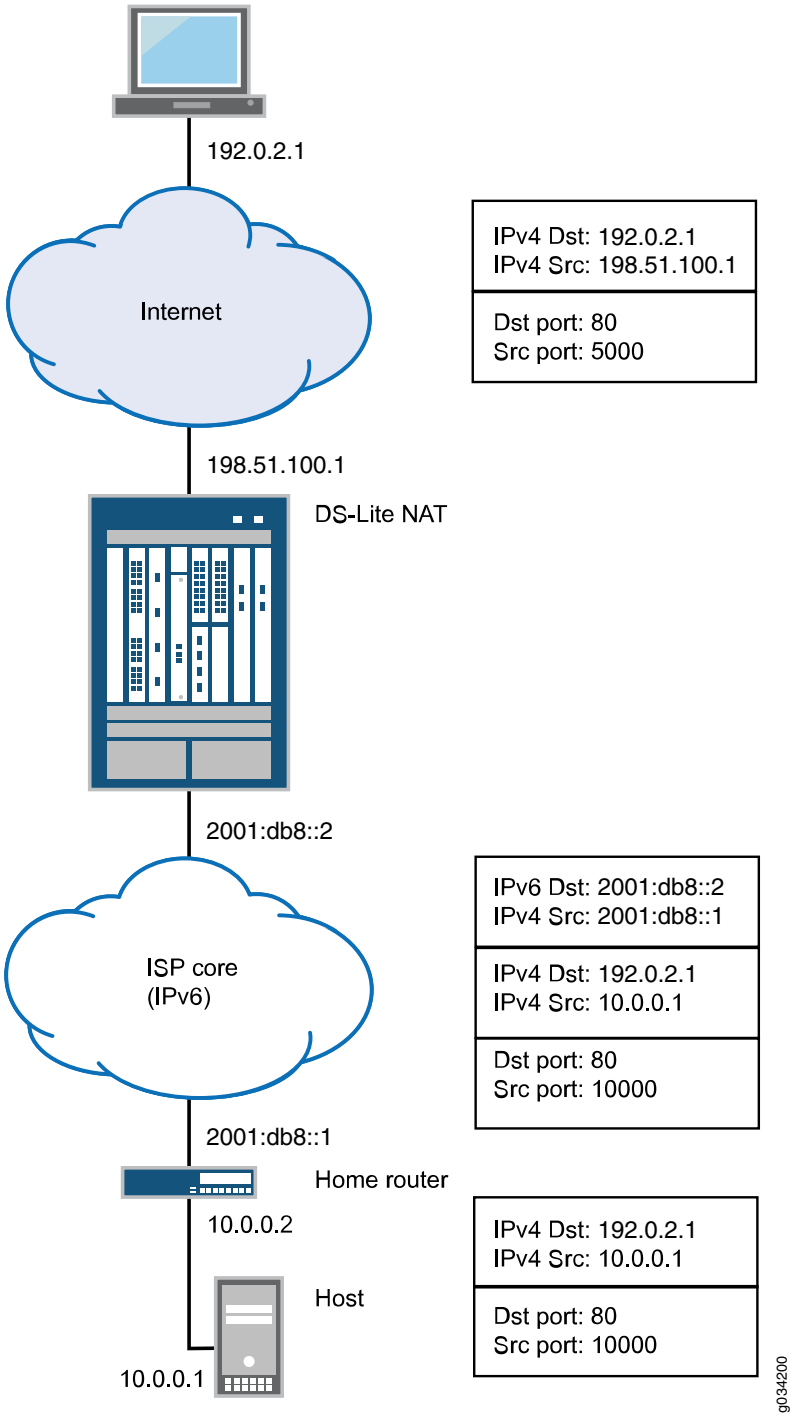

The DS-Lite deployment model consists of the following components:

Copyright © 2016, Juniper Networks, Inc.

- Softwire initiator for the DS-Lite home router--Encapsulates the IPv4 packet and transmits it across an IPv6 tunnel.
- Softwire concentrator for DS-Lite carrier-grade Network Address Translation (NAT)—Decapsulates the IPv4-in-IPv6 packet and also performs IPv4-IPv4 NAT translations.

When a user's device sends an IPv4 packet to an external destination, DS-Lite encapsulates the IPv4 packet in an IPv6 packet for transport into the provider network. These IPv4-in-IPv6 tunnels are called *softwires*. Tunneling IPv4 over IPv6 is simpler than translation and eliminates performance and redundancy concerns.

The softwires terminate in a softwire concentrator at some point in the service provider network, which decapsulates the IPv4 packets and sends them through a carrier-grade Network Address Translation (NAT) device. There, the packets undergo source NAT processing to hide the original source address.

IPv6 packets originated by hosts in the subscriber's home network are transported natively over the access network.

The DS-Lite carrier-grade NAT translates IPv4-to-IPv4 addresses to multiple subscribers through a single global IPv4 address. Overlapping address spaces used by subscribers are disambiguated through the identification of tunnel endpoints. One concentrator can be the endpoint of multiple softwires.

The IPv4 packets originated by the end hosts have private (and possibly overlapping) IP addresses. Therefore, NAT must be applied to these packets. If end hosts have overlapping addresses, Network Address Port Translation (NAPT) is needed.

Using NAPT, the system adds the source address of the encapsulating IPv6 packet in the subscriber network to the inside IPv4 source address and port. Because each user's IPv6 address is unique, the combination of the IPv6 source address with the IPv4 source address and port creates an unambiguous mapping.

The system takes the following actions when it receives a responding IPv4 packet from outside the subscriber network:

- Encapsulates the IPv4 packet in an IPv6 packet using the mapped IPv6 address as the IPv6 destination address.
- Forwards the packet to the user.

Table 175 lists the maximum number of softwire initiators and softwire concentrators per device.

Table 447: Softwire Initiator and Softwire Concentrator Capacity

Description	SRX1500	SRX5400	SRX5600	SRX5800
Maximum softwire initiators connected per device	300	100,000	100,000	100,000
Maximum softwire concentrator numbers per device	32	32	32	32

Copyright © 2016, Juniper Networks, Inc.

NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:

- The term softwire initiator has been replaced by B4.
- The term softwire concentrator has been replaced by AFTR.

Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the CLI statements used to configure DS-Lite.

For more information, see the following documents:

- draft-ietf-softwire-dual-stack-lite-06, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*, August 2010.
- RFC 2473, *Generic Packet Tunneling in IPv6 Specification*, December 1998.
- RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, August 1999.
- RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, BCP 127*, January 2007.
- RFC 4925, *Softwire Problem Statement*, July 2007.
- RFC 5382, *NAT Behavioral Requirements for TCP, BCP 142*, October 2008.
- RFC 5508, *NAT Behavioral Requirements for ICMP, BCP 148*, April 2009.
- http://www.potaroo.net/tools/ipv4/index.html
- http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml

Related Documentation

- Example: Configuring IPv6 Dual-Stack Lite on page 1776
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- Understanding How SRX Series Devices Handle ICMPv6 Packets on page 1769
- About the IPv6 Basic Packet Header on page 1764

Example: Configuring IPv6 Dual-Stack Lite

When an ISP begins to allocate IPv6 addresses and IPv6-capable equipment to new subscriber homes, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 CE WAN equipment to reach the IPv4 network. DS-Lite enables IPv4 customers to continue to access the Internet using their current hardware by using a softwire initiator at the customer edge to encapsulate IPv4 packets into IPv6 packets with minimum disruption to their home network, while enabling IPv6 customers to access IPv6 content. The softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4-IPv4 NAT translations.

Copyright © 2016, Juniper Networks, Inc.

This example shows you how to configure a softwire concentrator for IPv4-in-IPv6 addresses.

- Requirements on page 5035
- Overview on page 5035
- Configuration on page 5035
- Verification on page 5036

Requirements

Before you begin:

- Review the overview section on DS-Lite. See "Understanding IPv6 Dual-Stack Lite" on page 1773.
- Review how ICMPv6 packets are handled by the SRX Series devices. See "Understanding How SRX Series Devices Handle ICMPv6 Packets" on page 1769.

Overview

This configuration example shows how to configure a softwire concentrator, the softwire name, the concentrator address, and the softwire type.

NOTE: The softwire concentrator IPv6 address can match an IPv6 address configured on a physical interface or an IPv6 address configured on a loopback interface.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security softwires softwire-name my_sc1 softwire-concentrator 2001:db8::1
  softwire-type IPv4-in-IPv6
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a DS-Lite softwire concentrator to convert IPv4 packets into IPv6 packets:

1. Assign a name for the softwire concentrator.

   ```
   [edit security]
   user@host# edit softwires softwire-name my_sc1
   ```

2. Specify the address of the softwire concentrator.

   ```
   [edit security softwires softwire-name my_sc1]
   user@host# set softwire-concentrator 2001:db8::1
   ```

Copyright © 2016, Juniper Networks, Inc.

3. Specify the softwire type for IPv4 to IPv6.

```
[edit security softwires softwire-name my_sc1 softwire-concentrator 2001:db8::1
user@host# set softwire-type IPv4-in-IPv6
```

Results From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit security softwires softwire-name my_sc1]
user@host# show
softwire-concentrator 2001:db8::1;
softwire-type ipv4-in-ipv6;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

From operational mode, enter the **show security softwires** command. If a softwire is not connected, the operational output looks like the following sample:

```
user@host# show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  my-sc1                2001:db8::1         Active      0
```

If a softwire is connected, the operational output looks like the following sample:

```
user@host# show security softwires
Softwire Name            SC Address         Status   Number of SI connected
  my-sc1                2001:db8::1         Connected      1
```

Related Documentation

- Understanding IPv6 Dual-Stack Lite on page 1773
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 2455
- About the IPv6 Basic Packet Header on page 1764

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 209

Configuration Statements

- Security Configuration Statement Hierarchy on page 5039
- [edit security nat] Hierarchy Level on page 5041
- address (Security ARP Proxy) on page 5045
- address (Security Destination NAT) on page 5045
- address (Security NDP Proxy) on page 5046
- address-mapping on page 5046
- address-persistent (Security Source NAT) on page 5047
- address-persistent (Security Source NAT Pool) on page 5047
- address-pooling (Security Source NAT) on page 5048
- address-shared (Security Source NAT) on page 5049
- application (Security Destination NAT) on page 5049
- application (Security Source NAT) on page 5050
- application-services (Security Forwarding Process) on page 5051
- clear-threshold on page 5052
- description (Security NAT Pool) on page 5053
- description (Security NAT Rule) on page 5054
- description (Security NAT Rule Set) on page 5055
- destination (Security Destination NAT) on page 5056
- destination-address (Security Destination NAT) on page 5057
- destination-address (Security Source NAT) on page 5058
- destination-address (Security Static NAT) on page 5058
- destination-address-name (Security Destination NAT) on page 5059
- destination-address-name (Security Source NAT) on page 5059
- destination-address-name (Security Static NAT) on page 5060
- destination-nat on page 5060
- destination-port (Security Destination NAT) on page 5061
- destination-port (Security Source NAT) on page 5061
- destination-port (Security Static NAT) on page 5062

Copyright © 2016, Juniper Networks, Inc.

- from (Security NAT) on page 5062
- host-address-base on page 5063
- inactivity-timeout (Security Persistent NAT) on page 5063
- inet (Security Static NAT) on page 5064
- interface (Security NAT ARP Proxy) on page 5065
- interface (Security NAT NDP Proxy) on page 5065
- interface (Security Source NAT) on page 5066
- interface (Security Source NAT Rule Set) on page 5066
- mapped-port (Security Static NAT) on page 5067
- match (Security Destination NAT) on page 5068
- match (Security Source NAT) on page 5069
- match (Security Static NAT) on page 5070
- max-session-number on page 5070
- overflow-pool on page 5071
- nptv6-prefix on page 5072
- nptv6-prefix-name on page 5072
- permit (Security Persistent NAT) on page 5073
- persistent-nat on page 5074
- pool (Security Destination NAT) on page 5075
- pool (Security Source NAT) on page 5076
- pool (Security Source NAT Rule Set) on page 5077
- pool-default-port-range on page 5078
- pool-default-twin-port-range on page 5079
- pool-utilization-alarm on page 5080
- pool-utilization-alarm (Security Source NAT Pool) on page 5081
- port (Security Source NAT) on page 5082
- port-overloading (Security Source NAT Interface) on page 5083
- port-overloading-factor (Security Source NAT Interface) on page 5084
- port-overloading-factor (Security Source NAT Pool) on page 5085
- port-randomization on page 5086
- port-round-robin on page 5086
- prefix (Security Static NAT) on page 5087
- prefix-name (Security Static NAT) on page 5088
- protocol (Security Destination NAT) on page 5088
- protocol (Security Source NAT) on page 5089
- proxy-arp (Security NAT) on page 5089
- proxy-ndp (Security NAT) on page 5090

Copyright © 2016, Juniper Networks, Inc.

- raise-threshold on page 5090
- routing-instance (Security Destination NAT) on page 5091
- routing-instance (Security Source NAT) on page 5091
- rule (Security Destination NAT) on page 5092
- rule (Security Source NAT) on page 5093
- rule (Security Static NAT) on page 5095
- rule-session-count-alarm (Security Destination NAT Rule Set) on page 5096
- rule-session-count-alarm (Security Source NAT Rule Set) on page 5097
- rule-session-count-alarm (Security Static NAT Rule Set) on page 5098
- rule-set (Security Destination NAT) on page 5099
- rule-set (Security Source NAT) on page 5100
- rule-set (Security Static NAT) on page 5102
- source (Security Source NAT) on page 5104
- source-address (Security Destination NAT) on page 5106
- source-address (Security Source NAT) on page 5107
- source-address (Security Static NAT Rule Set) on page 5107
- source-address-name (Security Destination NAT) on page 5108
- source-address-name (Security Source NAT) on page 5108
- source-address-name (Security Static NAT Rule Set) on page 5109
- source-nat on page 5110
- source-port (Security Source NAT Rule Set) on page 5111
- source-port (Security Static NAT Rule Set) on page 5111
- static (Security NAT) on page 5112
- static-nat on page 5113
- to (Security Source NAT) on page 5114
- then (Security Destination NAT) on page 5114
- then (Security Source NAT) on page 5115
- then (Security Static NAT) on page 5116
- traceoptions (Security NAT) on page 5117

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Copyright © 2016, Juniper Networks, Inc.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- *CLI User Guide*
- CLI Explorer

[edit security nat] Hierarchy Level

```
security {
  nat {
    destination {
      pool pool-name {
        address <ip-address> {
          (port port-number | to ip-address);
        }
        description text;
        routing-instance (routing-instance-name | default);
      }
      rule-set rule-set-name {
        description text;
        from {
          interface [interface-name];
          routing-instance [routing-instance-name];
          zone [zone-name];
        }
        rule rule-name {
          description text;
          match {
            application {
              [application];
              any;
            }
            (destination-address ip-address | destination-address-name address-name);
            destination-port (port-or-low <to high>);
            protocol [protocol-name-or-number];
            source-address [ip-address];
            source-address-name [address-name];
          }
          then {
            destination-nat (off | pool pool-name | rule-session-count-alarm
              (clear-threshold value | raise-threshold value));
          }
        }
      }
    }
    proxy-arp interface interface-name address ip-address;
      to ip-address;
    }
    proxy-ndp interface interface-name address ip-address;
      to ip-address;
    }
    source {
      address-persistent;
      interface (port-overloading off | port-overloading-factor number);
      pool pool-name {
        address ip-address {
```

Copyright © 2016, Juniper Networks, Inc.

```
      to ip-address;
    }
    address-persistent subscriber ipv6-prefix-length prefix-length;
    address-pooling (paired | no-paired);
    address-shared;
    description text;
    host-address-base ip-address;
    overflow-pool (pool-name | interface);
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port {
      block-allocation {
        active-block-timeouttimeout-interval;
        block-size block-size;
        log disable;
        maximum-blocks-per-host maximum-block-number
      }
      deterministic {
        block-size block-size;
        host {
          address ip-address;
          address-name address-name;
      }
        no-translation;
        port-overloading-factor number;
        range {
          port-low <to port-high>;
          to port-high;
          twin-port port-low <to port-high>;
        }
    }
    routing-instance routing-instance-name;
}
pool-default-port-range lower-port-range to upper-port-range;
pool-default-twin-port-range lower-port-range to upper-port-range;
pool-utilization-alarm (clear-threshold value | raise-threshold value);
port-randomization disable;
port-round-robin disable;
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      application {
        [application];
        any;
      }
      (destination-address <ip-address> | destination-address-name
        <address-name>);
      destination-port (port-or-low <to high>);
      protocol [protocol-name-or-number];
      source-address [ip-address];
```


Copyright © 2016, Juniper Networks, Inc.

```
            source-address-name [address-name];
            source-port (port-or-low <to high>);
          }
          then source-nat;
            interface {
              persistent-nat {
                address-mapping;
                inactivity-timeout seconds;
                max-session-number value;
                permit (any-remote-host | target-host | target-host-port);
              }
            off;
            pool <pool-name>
              persistent-nat
                address-mapping;
                inactivity-timeout seconds;
                max-session-number number;
                permit (any-remote-host | target-host | target-host-port);
              }
            rule-session-count-alarm (clear-threshold value | raise-threshold value);
          }
        }
        to {
          interface [interface-name];
          routing-instance [routing-instance-name];
          zone [zone-name];
        }
      }
    }
    static rule-set rule-set-name;
        description text;
        from {
          interface [interface-name];
          routing-instance [routing-instance-name];
          zone [zone-name];
        }
        rule rule-name {
          description text;
          match {
            (destination-address <ip-address> | destination-address-name
              <address-name>);
            destination-port (port-or-low  | <to high>);
            source-address [ip-address];
            source-address-name [address-name];
            source-port (port-or-low  <to high>);
          }
          then static-nat;
            inet {
              routing-instance (routing-instance-name | default);
            }
            nptv6-prefix {
              address-prefix;
                routing-instance routing-instance-name;
            nptv6-prefix-name {
              address-prefix-name;
                routing-instance routing-instance-name;
```

```
                    [
                    prefix {
                      address-prefix;
                      mapped-port lower-port-range to upper-port-range;
                      routing-instance (routing-instance-name| default);
                    }
                    prefix-name {
                      address-prefix-name;
                      mapped-port lower-port-range to upper-port-range;
                      routing-instance (routing-instance-name | default);
                    }
                    rule-session-count-alarm (clear-threshold value | raise-threshold value);
                    }
                  }
                }
              }
            }
            traceoptions {
              file {
                filename;
                files number;
                match regular-expression;
                (world-readable | no-world-readable);
                size maximum-file-size;
              }
              flag flag;
              no-remote-trace;
            }
          }
        }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Introduction to NAT on page 4839

Copyright © 2016, Juniper Networks, Inc.

address (Security ARP Proxy)

Syntax address *ip-address* < to *ip-address*>;

Hierarchy Level [edit security nat proxy-arp interface *interface-name*],

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a single address or an address range of ARP proxy.

Options **to**—Specify the upper limit of the address range.

ip-address—IP address of an ARP proxy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

address (Security Destination NAT)

Syntax
```
address <ip-address> {
   (port port-number | to ip-address);
}
```

Hierarchy Level [edit security nat destination pool *pool-name*]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a single address or an address range of the destination NAT pool.

Options
- ***ip-address*** —IP address of a pool.
- **port** ***port-number***—Specify the port number.
- **to**—Specify the upper limit of the address range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

address (Security NDP Proxy)

Syntax

```
address ip-address {
   to ip-address;
}
```

Hierarchy Level [edit security nat proxy-ndp interface *interface-name*],

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a single address or an address range of NDP proxy.

Options

- ***ip-address***—IP address of an NDP proxy.
- **to**—Specify the upper limit of the address range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

address-mapping

Syntax address-mapping;

Hierarchy Level [edit security nat source rule-set *ruleset* rule *rule* then source-nat interface persistent-nat]
[edit security nat source rule-set *ruleset* rule *rule* then source-nat pool persistent-nat]

Release Information Statement introduced in Junos OS Release 10.2.

Description Allows requests from a specific internal IP address to be mapped to the same reflexive IP address (the public IP address created by the NAT device closest to the STUN server); internal and external ports can be any ports. An external host using any port can send a packet to the internal host by sending the packet to the reflexive IP address (with a configured incoming policy that allows external to internal traffic). If this option is not configured, the persistent NAT binding is for specific internal and reflexive transport addresses.

You can only specify this option when the persistent NAT type is **any-remote-host** and the source NAT rule action is one of the following:

- Source NAT pool with IP address shifting
- Source NAT pool with no port translation and no overflow pool

Required Privilege Level **security**—To view this statement in the configuration.
security-control—To add this statement to the configuration

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

address-persistent (Security Source NAT)

Syntax address-persistent;

Hierarchy Level [edit security nat source]

Release Information Statement modified in Junos OS Release 9.2.

Description Enable the device to assign the same, statically chosen, IP address from a source pool to a host for multiple sessions that require the same source IP address for each session. This option is a global configuration and is applied to all source pools. After a session is established from a host and NAT is performed, the subsequent session from the same host will always use the same translated address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

address-persistent (Security Source NAT Pool)

Syntax address-persistent subscriber ipv6-prefix-length *prefix-length*;

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Enable the device to translate an IPv6 address, with a consistent IPv6 prefix, to the same IPv4 address to ensure that IPv4 services can be used over IPv6-only networks.

Options ipv6-prefix-length *prefix-length*—Specify the subscriber IPv6 prefix length.

Range: 8 through 128.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation on page 4972
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

address-pooling (Security Source NAT)

Syntax address-pooling (paired | no-paired);

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Use the **address-pooling paired** option for applications that require all sessions associated with one internal IP address to be translated to the same external IP address for multiple sessions. Unlike the **address-persistent** option, this option is configured per pool, not globally. After a session is established from a host and NAT is performed, the subsequent sessions from the same host use the same translated address as long as the sessions are active.

Specify **address-pooling no-paired** for applications that can be assigned IP addresses in a round-robin fashion.

If either **address-pooling paired** or **address-pooling no-paired** is configured for a source NAT pool with port address translation, the persistent address option is disabled. If only **address-shared** is configured (no **address-pooling**) on a source NAT pool with no port translation, then the persistent address option is enabled. Both **address-shared** and **address-pooling paired** can be configured on the same source NAT pool with no port translation.

NOTE: The address-shared and address-pooling options are supported only on SRX Series devices.

Options **no-paired**—Allow address-pooling no-paired for a source pool without port translation.

paired—Allow address-pooling paired for a source pool with port translation.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Source NAT Pools with Address Pooling on page 4900
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

address-shared (Security Source NAT)

Syntax address-shared;

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specifies that multiple internal IP addresses can be mapped to the same external IP address. Use this option only when the source NAT pool is configured with no port translation.

When a source NAT pool configured with no port translation has few external IP addresses available, or only one external IP address, the **address-shared** option, with a many-to-one mapping, increases NAT resources and improves traffic.

Required Privilege Level security—To view this statement in the configuration
security-control—To add this statement to the configuration

Related Documentation
- Understanding Source NAT Pools with Shared Address on page 4921
- Example: Configuring a Single IP Address in a Source NAT Pool Without PAT on page 4913
- Security Configuration Statement Hierarchy on page 603

application (Security Destination NAT)

Syntax
```
application {
  [application];
  any;
}
```

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Specify an application name to match the rule. You can specify multiple application names, but the number of application terms must not exceed 3072.

Options ***application-name***—Name of the application.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- application (Security Policies) on page 334
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

application (Security Source NAT)

Syntax

```
application {
  [application];
  any;
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Specify an application name to match the rule. You can specify multiple application names, but the number of application terms must not exceed 3072.

Options ***application-name***—Name of the application.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- application (Security Policies) on page 334
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

application-services (Security Forwarding Process)

Syntax

```
application-services {
  enable-gtpu-distribution;
  maximize-alg-sessions;
  maximize-idp-sessions {
    inline-tap;
    weight (equal | firewall | idp);
  }
  packet-ordering-mode {
    (hardware | software);
  }
}
```

Hierarchy Level [edit security forwarding-process]

Release Information Statement introduced in Junos OS Release 9.6. Statement updated in Junos OS Release 10.4. Statement updated in Junos OS Release 15.1X49-D40 with the **enable-gtpu-distribution** option.

Description You can configure the device to switch from an integrated firewall mode to maximize intrusion detection and prevention (IDP) mode to increase the capacity of IDP processing with the **maximize-idp-sessions** option. When you maximize IDP, you are decoupling IDP processes from firewall processes, allowing the device to support the same number of firewall and IDP sessions.

You can configure maximum Application Layer Gateway (ALG) sessions by using the **maximize-alg-sessions** option. By default, the session capacity number for Real-Time Streaming Protocol (RTSP), FTP, and Trivial File Transfer Protocol (TFTP) ALG sessions is 10,000 per flow Services Processing Unit (SPU). You must reboot the device (and its peer in chassis cluster mode) for the configuration to take effect. The **maximize-alg-sessions** option now enables you to increase defaults as follows:

- RTSP, FTP, and TFTP ALG session capacity: 25,000 per flow SPU
- TCP proxy connection capacity: 40,000 per flow SPU

NOTE: Flow session capacity is reduced to half per flow SPU; therefore the aforementioned capacity numbers will not change on central point flow.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Juniper Networks Devices Processing Overview on page 1665
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

clear-threshold

Syntax clear-threshold *value*;

Hierarchy Level [edit security nat source pool-utilization-alarm]

Release Information Statement modified in Junos OS Release 9.2.

Description Configure the lower threshold at which an SNMP trap is triggered when pool utilization for a source pool without Port Address Translation (PAT) falls below the threshold.

Options **clear-threshold *value***—Threshold at which an SNMP trap is triggered.

Range: 40 through 100

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

description (Security NAT Pool)

Syntax description *text*;

Hierarchy Level [edit security nat destination pool *pool-name*]
[edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for a source or destination NAT pool.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options ***text***—Descriptive text about a source or destination NAT pool.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

description (Security NAT Rule)

Syntax description *text*;

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name*]
[edit security nat source rule-set *rule-set-name* rule *rule-name*]
[edit security nat static rule-set *rule-set-name* rule *rule-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for a source, destination, or static NAT rule.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options ***text***—Descriptive text about a source, destination, or static NAT rule.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

description (Security NAT Rule Set)

Syntax description *text*;

Hierarchy Level [edit security nat destination rule-set *rule-set-name*]
[edit security nat source rule-set *rule-set-name*]
[edit security nat static rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify descriptive text for a source, destination, or static NAT rule set.

NOTE: The descriptive text should not include characters, such as “<”, “>”, “&”, or “\n”.

Options *text*—Descriptive text about a source, destination, or static NAT rule set.

Range: 1 through 300 characters

NOTE: The upper limit of the description text range is related to character encoding, and is therefore dynamic. However, if you configure the descriptive text length beyond 300 characters, the configuration might fail to take effect.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination (Security Destination NAT)

Syntax

```
destination {
  pool pool-name {
    address <ip-address> {
      (port port-number | to ip-address);
    }
    description text;
    routing-instance (routing-instance-name | default);
  }
  rule-set rule-set-name {
    description text;
    from {
      interface [interface-name];
      routing-instance [routing-instance-name];
      zone [zone-name];
    }
    rule rule-name {
      description text;
      match {
        application {
          [application];
          any;
        }
        (destination-address ip-address| destination-address-name address-name);
        destination-port (port-or-low <to high>);
        protocol [protocol-name-or-number];
        source-address [ip-address];
        source-address-name [address-name];
      }
      then {
        destination-nat (off | pool pool-name |rule-session-count-alarm (clear-threshold
          value | raise-threshold value));
      }
    }
  }
}
```

Hierarchy Level [edit security nat]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Configure destination NAT, which allows you to configure the following:

- Translate destination IP address or addresses to a specific IP address.
- Translate destination IP address or addresses and port number(s) to a specific IP address and one port number.
- Translate a range of destination IP addresses to another range of IP addresses. This mapping is one-to-one, static, and without PAT.

Copyright © 2016, Juniper Networks, Inc.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

destination-address (Security Destination NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure one address or a subnet.

NOTE:

- If the destination address is IPv4 and the pool is an IPv6 prefix, the length of the IPv6 prefix must be 96.
- If the destination address is an IPv6 prefix and the pool is an IPv6 prefix, their length must be the same.

Options *ip-address*— Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-address (Security Source NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure multiple addresses or subnets.

Options ***ip-address***—Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

destination-address (Security Static NAT)

Syntax destination-address <*ip-address*>;

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address to match the rule. You can configure one address or a subnet.

Options ***ip-address***—Destination address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-address-name (Security Destination NAT)

Syntax destination-address-name <*address-name*>;

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address name to match the rule. You can configure multiple address names.

Options ***address-name***—Destination address name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

destination-address-name (Security Source NAT)

Syntax destination-address-name <*address-name*>;

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address name to match the rule. You can configure multiple address names.

Options ***address-name***—Destination address name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-address-name (Security Static NAT)

Syntax destination-address-name <*address-name*>;

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a destination address name to match the rule.

Options ***destination-address-name***—Name of the destination address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

destination-nat

Syntax destination-nat (off | pool *pool-name* | rule-session-count-alarm (clear-threshold *value* | raise-threshold *value*));

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. The **rule-session-count-alarm** option added in Junos OS Release 12.1X45-D10.

Description Specify the action of the destination NAT rule.

Options **off**—Do not perform destination NAT operation.

pool—Use user-defined destination NAT pool to perform destination NAT.

rule-session-count-alarm—Define session count alarm thresholds for a specific destination NAT rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

NOTE: If you enter a value for **raise-threshold** but not for **clear-threshold**, **clear-threshold** is automatically set to 80 percent of **raise-threshold**.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-port (Security Destination NAT)

Syntax destination-port (port-or-low <to *high*>);

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

Description Specify a destination port or port range to match the rule. Up to eight port or port ranges are supported.

Options ***port*** —Specify a destination port number.

low—Specify the lower limit of the destination port range.

<to *high*>—Specify the upper limit of the destination port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

destination-port (Security Source NAT)

Syntax destination-port (port-or-low <to *high*>);

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

Description Specify a destination port or port range to match the rule. Up to eight port or port ranges are supported.

Options ***port*** —Specify a destination port number.

low—Specify the lower limit of the destination port range.

<to *high*>—Specify the upper limit of the destination port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

destination-port (Security Static NAT)

Syntax destination-port (*port-or-low* | <to *high*>);

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify a destination port or port range to allow static NAT to map ports.

Options ***port-or-low***—Specify the port name or the lower limit of the port range.

to *high*—Specify the upper limit of the port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

from (Security NAT)

Syntax
```
from {
  interface [interface-name];
  routing-instance [routing-instance-name];
  zone [zone-name];
}
```

Hierarchy Level [edit security nat destination rule-set *rule-set-name*]
[edit security nat source rule-set *rule-set-name*]
[edit security nat static rule-set *rule-set-name*]

Release Information Statement modified in Junos OS Release 9.3.

Description Specify the source of the packet among the routing instance, interface, or zone.

Options
- **interface [*interface-name*]** —Name of the interface.
- **routing-instance [*routing-instance-name*]** —Name of the routing instance.
- **zone [*zone-name*]** —Name of the zone.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

host-address-base

Syntax host-address-base *ip-address*;

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the base address of the original source IP address range. This is used for IP shifting.

Options ***ip-address*** —IP address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

inactivity-timeout (Security Persistent NAT)

Syntax inactivity-timeout *seconds*;

Hierarchy Level [edit security nat source rule-set *ruleset* rule *rule* then source-nat interface persistent-nat]
[edit security nat source rule-set *ruleset* rule *rule* then source-nat pool persistent-nat]

Release Information Statement introduced in Junos OS Release 9.6.

Description The amount of time, in seconds, that the persistent NAT binding remains in the Juniper Networks device's memory when all the sessions of the binding entry are gone. When the configured timeout is reached, the binding is removed from memory.

Options *seconds*—Number of seconds.

Range: 60 through 7200 seconds

Default: 300 seconds (5 minutes)

Required Privilege Level **security**—To view this statement in the configuration.
security-control—To add this statement to the configuration

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

inet (Security Static NAT)

Syntax

```
inet {
  routing-instance (routing-instance-name| default);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify the automatic translation of IPv6 addresses to IPv4 addresses (and vice versa).

NOTE: If you use this option, you do not need to use the *prefix* option because the first 96 most significant bits are automatically stripped from the 128-bit IPv6 address.

Options

- **routing-instance** ***routing-instance-name*** —Use the user-defined static NAT routing-instance to perform static NAT.
- **default**—Use the default routing-instance to perform static NAT. When a **routing-instance-name** is not provided, the default routing-instance master is used, which refers to the main inet.0 (for IPv4 unicast routes) routing table or inet.6 (for IPv6 unicast routes) routing table.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

interface (Security NAT ARP Proxy)

Syntax

```
interface interface-name {
  address ip-address {
    to ip-address;
  }
}
```

Hierarchy Level [edit security nat proxy-arp]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify the interface on which the ARP proxy is to be configured. It should be a logical interface.

Options *interface-name*—Name of the logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

interface (Security NAT NDP Proxy)

Syntax

```
interface interface-name {
  address ip-address {
    to ip-address;
  }
}
```

Hierarchy Level [edit security nat proxy-ndp]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify the interface on which the NDP proxy is to be configured. It should be a logical interface.

Options *interface-name*—Name of the logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

interface (Security Source NAT)

Syntax interface (port-overloading off | port-overloading-factor *number*);

Hierarchy Level [edit security nat source]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Enable interface NAT with or without port overloading.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

interface (Security Source NAT Rule Set)

Syntax

```
interface {
  persistent-nat {
    address-mapping;
    inactivity-timeout seconds;
    max-session-number value;
    permit (any-remote-host | target-host | target-host-port);
  }
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* then source-nat]

Release Information Statement introduced in Junos OS Release 9.6.

Description Enable interface NAT with or without port overloading.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

mapped-port (Security Static NAT)

Syntax mapped-port *lower-port-range* to *upper-port-range*;

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat prefix]
[edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat prefix-name]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify a destination port or port range to allow static NAT to map ports.

Options
- **lower-port-range**—Specify the lower limit of the port range.
- **upper-port-range**—Specify the upper limit of the port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

match (Security Destination NAT)

Syntax

```
match {
    application {
      [application];
      any;
    }
  (destination-address ip-address | destination-address-name address-name);
  destination-port (port-or-low <to high>);
  protocol [protocol-name-or-number];
  source-address [ip-address];
  source-address-name [address-name];
}
```

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X47-D10.

Description Specify the destination rules to be used as match criteria.

NOTE: If the options **destination-port** and **protocol** are configured as match conditions, then you cannot also configure the **application** option as a match condition. The reverse is also true: if you configure the **application** option as a match condition for a rule, you cannot also configure the **destination-port** and **protocol** options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

match (Security Source NAT)

Syntax

```
match {
  application {
    [application];
    any;
  }
  (destination-address <ip-address> | destination-address-name <address-name>);
  destination-port (port-or-low <to high>);
  protocol [protocol-name | protocol-number];
  source-address [ip-address];
  source-address-name [address-name];
  source-port (port-or-low <to high>);
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Specify the source rules to be used as match criteria.

NOTE: If the options **source-port**, **destination-port**, and **protocol** are configured as match conditions, then you cannot also configure the **application** option as a match condition. The reverse is also true: if you configure the **application** option as a match condition for a rule, you cannot also configure the **source-port**, **destination-port**, and **protocol** options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

match (Security Static NAT)

Syntax

```
match {
  (destination-address <ip-address> | destination-address-name <address-name>);
  destination-port (port-or-low | <to high>);
  source-address [ip-address];
  source-address-name [ip-address-name];
  source-port (port-or-low <to high>);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Specify the static rules to be used as match criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

max-session-number

Syntax max-session-number *number*;

Hierarchy Level [edit security nat source rule-set *ruleset* rule *rule* then source-nat interface persistent-nat]
[edit security nat source rule-set *ruleset* rule *rule* then source-nat pool persistent-nat]

Release Information Statement introduced in Junos OS Release 9.6.

Description The maximum number of the sessions with which a persistent NAT binding can be associated. For example, if the **max-session-number** of the persistent NAT rule is 65,536, then a 65,537th session cannot be established if that session uses the persistent NAT binding created from the persistent NAT rule.

Options *number*—Maximum number of sessions.

Range: 8 through 65,536

Default: 30 sessions

Required Privilege Level **security**—To view this statement in the configuration.
security-control—To add this statement to the configuration

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

overflow-pool

Syntax overflow-pool (interface | *pool-name*);

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a source pool to use when the current address pool is exhausted. Currently the statement is applicable for IPv4 addresses only.

NOTE: The length of the IPv6 prefix must be 96 when the pool is used for NAT-PT.

Options

- **interface** — Allow the interface pool to support overflow.
- ***pool-name*** — Name of the source address pool.

NOTE: The source pool must have Port Address Translation (PAT) enabled. PAT is not supported when the address is an IPv6 prefix address.

Required Privilege Level security — To view this statement in the configuration.
security-control — To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

nptv6-prefix

Syntax

```
nptv6-prefix {
  address-prefix;
    routing-instance routing-instance-name;
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement introduced in Junos OS Release 12.3X48-D25.

Description Specify a static IPv6 address prefix. The longest prefix supported is /64.

Options

- ***address-prefix***—Specify the address prefix.
- **routing-instance**—Use the user-defined static NAT routing instance to perform static NAT.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

nptv6-prefix-name

Syntax

```
nptv6-prefix-name {
  address-prefix-name;
    routing-instance routing-instance-name;
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement introduced in Junos OS Release 12.3X48-D25.

Description Specify an address prefix name from an address book. The longest prefix name supported is /64.

Options

- ***address-prefix-name***—Specify an address prefix name from an address book.
- **routing-instance** —Use the user-defined static NAT routing instance to perform static NAT.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

permit (Security Persistent NAT)

Syntax permit (any-remote-host | target-host | target-host-port);

Hierarchy Level [edit security nat source rule-set *ruleset* rule *rule* then source-nat interface persistent-nat]
[edit security nat source rule-set *ruleset* rule *rule* then source-nat pool persistent-nat]

Release Information Statement introduced in Junos OS Release 9.6. Support for IPv6 addresses added in Junos OS Release 11.2.

Description Configure persistent NAT mappings.

Options

- **any-remote-host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. (The reflexive transport address is the public IP address and port created by the NAT device closest to the STUN server.) Any external host can send a packet to the internal host by sending the packet to the reflexive transport address.
- **target-host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address.
- **target-host-port**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.

NOTE: **The target-host-port configuration is not supported for NAT64 when configured with IPv6 address.**

Required Privilege Level **security**—To view this statement in the configuration.
security-control—To add this statement to the configuration

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

persistent-nat

Syntax

```
persistent-nat {
  address-mapping;
  inactivity-timeout seconds;
  max-session-number value;
  permit (any-remote-host | target-host | target-host-port);
}
```

Hierarchy Level

[edit security nat source rule-set *ruleset* rule *rule* then source-nat interface]
[edit security nat source rule-set *ruleset* rule *rule* then source-nat pool]

Release Information

Statement introduced in Junos OS Release 9.6. Support for address-mapping added in Junos OS Release 10.2.

Description

Use the **persistent-nat** feature to ensure that all requests from the same internal transport address are mapped to the same reflexive transport address (the public IP address and port created by the NAT device closest to the STUN server). The source NAT rule action can use a source NAT pool (with or without port translation) or an egress interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration

Related Documentation

- Understanding Persistent NAT and NAT64 on page 4969
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

pool (Security Destination NAT)

Syntax

```
pool pool-name {
  address <ip-address> {
    (port port-number | to ip-address);
  }
  description text;
  routing-instance (routing-instance-name | default);
}
```

Hierarchy Level [edit security nat destination]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1.

Description Define a destination NAT pool to identify the pool uniquely.

Options

- *pool-name*—Name of the pool.
- description—Description of the pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

pool (Security Source NAT)

Syntax

```
pool pool-name {
  address ip-address {
    to ip-address;
  }
  address-persistent subscriber ipv6-prefix-length prefix-length;
  address-pooling (paired | no-paired);
  address-shared;
  description text;
  host-address-base ip-address;
  overflow-pool (interface | pool-name);
  pool-utilization-alarm (clear-threshold value | raise-threshold value);
  port (no-translation | port-overloading-factor number | range port-low (to port-high));
  routing-instance routing-instance-name;
}
```

Hierarchy Level [edit security nat source]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D10.

Description Define a source NAT pool to identify the pool uniquely.

Options *pool-name*—Name of the pool.

description—Description of the pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

pool (Security Source NAT Rule Set)

Syntax

```
pool (pool-name) {
  persistent-nat {
    address-mapping;
    inactivity-timeout seconds;
    max-session-number number;
    permit (any-remote-host | target-host | target-host-port);
  }
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* then source-nat]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify to use source NAT pool.

Options *pool-name*—Name of the source NAT pool.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

pool-default-port-range

Syntax pool-default-port-range *lower-port-range* to *upper-port-range*;

Hierarchy Level [edit security nat source]

Release Information Statement introduced in Junos OS Release 11.4.

Description Set the global default single port range for source NAT pools with port translation. If the port range in source NAT pools is not specified, the configured default port range is used. If neither the port range in source NAT pools nor the default port range are configured, the default single port range is 1024 through 63,487.

To set the global twin port range for source NAT pools with port translation, use the **pool-default-twin-port-range** statement at the [edit security nat source] hierarchy. The twin port range is 63,488 through 65,535.

To set the single port range for a specific pool, use the **port range *port-low* (to *port-high*)** statement at the [edit security nat source pool] hierarchy level.

Options

- ***lower-port-range***—Specify the lower limit of the port range.
- ***upper-port-range***—Specify the upper limit of the port range.

Range: 1024 through 63,487. To view pool information, use the **show security nat source pool** command.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- pool (Security Source NAT) on page 5076
- pool-default-twin-port-range on page 5079

Copyright © 2016, Juniper Networks, Inc.

pool-default-twin-port-range

Syntax pool-default-twin-port-range *lower-port-range* to *upper-port-range*;

Hierarchy Level [edit security nat source]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Specify the global default twin port range for all source pools. Two ports within range (63488, 65535) are allocated at a time for RTP/RTCP applications such as SIP, H.323, and RTSP for source pools with PAT.

The default twin port range is 2048. If you have an SRX5400, SRX5600, or SRX5800 device that supports a maximum of 1 million IP addresses, use this option to limit the twin port range and avoid exceeding the port capacity of 384 million.

To set the twin port range for a specific pool, use the **port range twin-port *port-low* (to *port-high*)** statement at the [edit security nat source pool] hierarchy level.

Options

- ***lower-twin-port-range***—Specify the lower limit of the port range.
- ***upper-twin-port-range***—Specify the upper limit of the port range.

Range: 63,488 through 65,535.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- **pool (Security Source NAT) on page 5076**
- Understanding Source NAT Pool Capacities on page 4897

Copyright © 2016, Juniper Networks, Inc.

pool-utilization-alarm

Syntax pool-utilization-alarm (clear-threshold *value* | raise-threshold *value*);

Hierarchy Level [edit security nat source]

Release Information Statement modified in Junos OS Release 9.2.

Description Define the global pool utilization alarm thresholds for Network Address Translation (NAT) source IP address pools without Port Address Translation (PAT). When the pool utilization exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options

clear-threshold *value*—Lower threshold at which an SNMP trap is triggered.

Range: 40 through 100.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

Range: 50 through 100.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- pool-utilization-alarm (Security Source NAT Pool) on page 5081
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

pool-utilization-alarm (Security Source NAT Pool)

Syntax pool-utilization-alarm (clear-threshold *value* | raise-threshold *value*);

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Define utilization alarm thresholds for a specific Network Address Translation (NAT) source pool. When pool utilization exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered. Threshold settings that use this statement take precedence over thresholds that are set using the global **pool-utilization-alarm** statement in the **[security nat source]** hierarchy.

Options

clear-threshold *value*—Lower threshold at which an SNMP trap is triggered.

Range: 40 through 100.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

Range: 50 through 100.

NOTE: If you enter a value for raise-threshold but not for clear-threshold, clear-threshold is automatically set to 80 percent of raise-threshold.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- pool-utilization-alarm on page 5080
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

port (Security Source NAT)

Syntax

```
port {
  block-allocation {
    active-block-timeout timeout-interval;
    block-size block-size;
    log disable;
    maximum-blocks-per-host maximum-block-number;
  }
  deterministic {
    block-size block-size;
    host {
      address ip-address;
      address-name address-name ;
    }
  }
  no-translation;
  port-overloading-factor number;
  range {
    port-low <to port-high>;
    to port-high;
    twin-port port-low <to port-high>;
  }
}
```

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement introduced in Junos OS Release 9.2. Statement updated with **block-allocation**, **deterministic**, and **twin-port** options in Junos OS Release 12.1X47-D10.

Description Specify the Port Address Translation (PAT) for a source pool.

Options

- **block-allocation**—Allocates a block of ports for translation, instead of allocating individual ports.
- **deterministic**—An incoming (source) IP address and port always map to the specific destination address and port block, based on pre-defined deterministic NAT algorithm.
- **no-translation**—If set, no Port Address Translation is required. This option cannot be configured with the **port-overloading-factor** option or the **range** option.
- **port-overloading-factor** *number*—Configure the port overloading capacity in source NAT. This option cannot be configured with the **no-translation** option.
- **range** *port-low* **<to** *port-high***>**—Specify the port number range attached to each address in the pool. This option cannot be configured with the **no-translation** option.
- **twin port**—Configure the twin port range for source NAT pools to avoid port overloading.

The remaining statements are explained separately.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

port-overloading (Security Source NAT Interface)

Syntax port-overloading off

Hierarchy Level [edit security nat source interface]

Release Information Statement introduced in Junos OS Release 9.6.

Description Enable interface NAT with or without port overloading.

Options **off**—Specify off to disable interface port overloading.

NOTE: The **port-overloading** option should not be used in conjunction with the **port-overloading-factor** option because they can override each other. For example, if **port-overloading** has been set to **off** to disable interface port overloading, and subsequently the **port-overloading-factor** is configured with any value greater than 1, the **port-overloading-factor** setting will override the **port-overloading** setting. (Configuring **port-overloading-factor 1** is equivalent to configuring **port-overloading off.**)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- port-overloading-factor (Security Source NAT Interface) on page 5084

Copyright © 2016, Juniper Networks, Inc.

port-overloading-factor (Security Source NAT Interface)

Syntax port-overloading-factor *number*;

Hierarchy Level [edit security nat source interface]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Configure the port overloading capacity for the source NAT interface. If **port-overloading-factor** is set to *x*(1 up to the maximum port capacity), then x times the maximum port capacity is allocated for interface-based NAT.

NOTE: There is also a **port-overloading** option, but it is not supported for logical systems, and should not be used in conjunction with the **port-overloading-factor** option because the statements can overwrite each other. For example, if **port-overloading** has been set to **off** to disable interface port overloading, and subsequently **port-overloading-factor** is configured with any value greater than 1, the **port-overloading-factor** setting will override the **port-overloading** setting. (Configuring **port-overloading-factor 1** is equivalent to configuring **port-overloading off**.)

Options *number*—A number ranging from 1 through the maximum port capacity.

For example, if **port-overloading-factor** is set to 2, and it is multiplied by a maximum port capacity of 63,486, the port overloading threshold is 126,972. If the configured **port-overloading-factor** setting exceeds the maximum port capacity of the interface, an error message is generated during the configuration commit.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- port-overloading (Security Source NAT Interface) on page 5083
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

port-overloading-factor (Security Source NAT Pool)

Syntax port-overloading-factor

Hierarchy Level [edit security nat source pool *source-pool-name* port]

Release Information Statement introduced in Junos OS Release 11.2

Description Configures the port overloading capacity in source NAT. If the port-overloading-factor is set to *x*, each translated IP address will have *x* number of ports available.

NOTE: **The port-overloading-factor statement cannot be configured with port no-translation (source NAT pool without PAT).**

Options **Range:** 2 through 32

For example, If you set **port-overloading-factor** to 2 for a source pool with two IP addresses, each with the single port range of 1024 through 2047, the ports are multiplied by 2, increasing the port capacity for each from 1024 to 2048. If the configured port-overloading-factor setting exceeds the maximum port capacity of the pool, an error message is generated during the configuration commit.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

port-randomization

Syntax port-randomization disable;

Hierarchy Level [edit security nat source]

Release Information Statement introduced in Junos OS Release 9.6.

Description Disable random port allocation for pool-based and interface source NAT.

Options **disable**—Disables random port allocation for pool-based and interface source NAT. For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

port-round-robin

Syntax port-round-robin disable;

Hierarchy Level [edit security nat source]

Release Information Statement introduced in Junos OS Release 15.1X49-D30.

Description Disable round-robin port allocation for pool-based and interface source NAT.

Options **disable**—Disables round-robin port allocation for pool-based and interface source NAT.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

prefix (Security Static NAT)

Syntax

```
prefix {
  address-prefix;
  mapped-port lower-port-range to upper-port-range;
  routing-instance (routing-instance-name| default);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a static IP address prefix.

NOTE: If you use the *inet* option for translation of IPv6 addresses to IPv4 addresses (and vice versa), you do not need to specify a prefix because the *inet* option automatically strips the first 96 most significant bits from the 128-bit IPv6 address.

Options

- ***address-prefix***—Specify address prefix.
- **mapped-port *lower-port-range* to *upper-port-range***—Specify a destination port or port range to allow static NAT to map ports.
- **routing-instance** —Specify routing instance type:
  - ***routing-instance-name***—Use the user-defined static NAT routing instance to perform static NAT.
  - **default**—Use the default routing-instance.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

prefix-name (Security Static NAT)

Syntax

```
prefix-name {
  address-prefix-name;
  mapped-port lower-port-range to upper-port-range;
  routing-instance (routing-instance-name| default);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify an address from the address book.

Options

- ***address-prefix-name***—Specify address prefix name from address book.
- **mapped-port *lower-port-range* to *upper-port-range***—Specify a destination port or port range to allow static NAT to map ports.
- **routing-instance** —Specify routing instance type:
  - ***routing-instance-name***—Use the user-defined static NAT routing instance to perform static NAT.
  - **default**—Use the default routing-instance.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

protocol (Security Destination NAT)

Syntax protocol [*protocol-name-or-number*];

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify an IP protocol to match the rule. You can configure multiple protocol names or protocol numbers.

Options ***protocol-name-or-number***—Name or number of the specific protocol.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

protocol (Security Source NAT)

Syntax protocol [*protocol-name-or-number*];

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify an IP protocol to match the rule. You can configure multiple protocol names or protocol numbers.

Options ***protocol-name-or-number***—Name or number of the specific protocol.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

proxy-arp (Security NAT)

Syntax
```
proxy-arp {
  interface interface-name {
    address ip-address {
      to ip-address;
    }
  }
}
```

Hierarchy Level [edit security nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Configure Address Resolution Protocol (ARP) proxy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

proxy-ndp (Security NAT)

Syntax

```
proxy-ndp {
  interface interface-name {
    address ip-address {
      to ip-address;
    }
  }
}
```

Hierarchy Level [edit security nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Configure Neighbor Discovery Protocol (NDP) proxy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

raise-threshold

Syntax raise-threshold *value*;

Hierarchy Level [edit security nat source pool-utilization-alarm]

Release Information Statement modified in Junos OS Release 9.2.

Description Configure the upper threshold at which an SNMP trap is triggered when pool utilization for a source pool without Port Address Translation (PAT) rises above the threshold. This feature is disabled by default.

Options **raise-threshold *value***—Threshold at which an SNMP trap is triggered.

Range: 50 through 100

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

routing-instance (Security Destination NAT)

Syntax routing-instance (*routing-instance-name* | default);

Hierarchy Level [edit security nat destination pool *pool-name*]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify the routing instance on which to perform the route lookup for the address in the pool. It is not a mandatory flag.

A destination NAT pool that does not specify a specific routing instance will default to the routing instance of the ingress zone. You can configure a NAT pool to exist in the default routing instance. As a result, the NAT pool is reachable from zones in the default routing instance and from zones in other routing instances.

Options ***routing-instance-name***—Name of the routing instance.

default—Use the default routing instance.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

routing-instance (Security Source NAT)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit security nat source pool *pool-name*]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify the routing instance to which the pool is bound. It is not a mandatory flag. If the user does not configure the routing instance, by default the pool belongs to routing-instance **inet.0**.

Options ***routing-instance-name***—Name of the routing instance.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule (Security Destination NAT)

Syntax

```
rule rule-name {
  description text;
  match {
    application {
      [application];
      any;
    }
    (destination-address ip-address| destination-address-name address-name);
    destination-port (port-or-low <to high>);
    protocol [protocol-name-or-number];
    source-address [ip-address];
    source-address-name [address-name];
  }
  then {
    destination-nat (off | pool pool-name | rule-session-count-alarm (clear threshold value
      | raise-threshold value));
  }
}
```

Hierarchy Level [edit security nat destination rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 9.2. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Define a destination NAT rule.

Options

- ***rule-name***—Name of the destination NAT rule.
- **description**—Description of the destination NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule (Security Source NAT)

Syntax

```
rule rule-name {
  description text;
  match {
    application {
      [application];
      any;
    }
    (destination-address <ip-address> | destination-address-name <address-name>);
    destination-port (port-or-low <to high>);
    protocol [protocol-name-or-number];
    source-address [ip-address];
    source-address-name [address-name];
    source-port (port-or-low <to high>)
  }
  then {
    source-nat {
      interface {
        persistent-nat {
          address-mapping;
          inactivity-timeout seconds;
          max-session-number value;
          permit (any-remote-host | target-host | target-host-port);
        }
      }
      off;
      pool <pool-name>
        persistent-nat {
          address-mapping;
          inactivity-timeout seconds;
          max-session-number number;
          permit (any-remote-host | target-host | target-host-port);
        }
      }
    }
  }
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name*]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Define a source NAT rule.

Options

- ***rule-name***—Name of the source NAT rule.
- **description**—Description of the source NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule (Security Static NAT)

Syntax

```
rule rule-name {
  description text;
  match {
    (destination-address <ip-address> | destination-address-name <address-name>);
    destination-port (port-or-low | <to high>);
    source-address [ip-address];
    source-address-name [ip-address-name];
    source-port (port-or-low <to high>);
  }
  then {
    static-nat {
      inet {
        routing-instance (routing-instance-name| default);
      }
      nptv6-prefix {
        address-prefix;
          routing-instance routing-instance-name;
      nptv6-prefix-name {
        address-prefix-name;
          routing-instance routing-instance-name;
      [
      prefix {
        address-prefix;
        mapped-port lower-port-range to upper-port-range;
        routing-instance (routing-instance-name| default);
      }
      prefix-name {
        address-prefix-name;
        mapped-port lower-port-range to upper-port-range;
        routing-instance (routing-instance-name| default);
        }
      rule-session-count-alarm (clear-threshold value | raise-threshold value);
    }
  }
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 9.3. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

Description Define a static NAT rule.

Options

- ***rule-name***—Name of the static NAT rule.
- **Description**—Description of the static NAT rule.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

rule-session-count-alarm (Security Destination NAT Rule Set)

Syntax
rule-session-count-alarm (clear-threshold *value* | raise-threshold *value*):

Hierarchy Level
[edit security nat destination rule-set *rule-set-name* rule *rule-name* then destination-nat]

Release Information
Statement introduced in Junos OS Release 12.1X45-D10.

Description
Define session count alarm thresholds for a specific Network Address Translation (NAT) destination rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options
clear-threshold *value*—Lower threshold at which an SNMP trap is triggered.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

NOTE: If you enter a value for raise-threshold but not for clear-threshold, clear-threshold is automatically set to 80 percent of raise-threshold.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule-session-count-alarm (Security Source NAT Rule Set)

Syntax rule-session-count-alarm (clear-threshold *value* | raise-threshold *value*):

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* then source-nat]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Define session count alarm thresholds for a specific Network Address Translation (NAT) source rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options **clear-threshold *value***—Lower threshold at which an SNMP trap is triggered.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

NOTE: If you enter a value for raise-threshold but not for clear-threshold, clear-threshold is automatically set to 80 percent of raise-threshold.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule-session-count-alarm (Security Static NAT Rule Set)

Syntax rule-session-count-alarm (clear-threshold *value* | raise-threshold *value*);

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then static-nat]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Define session count alarm thresholds for a specific static Network Address Translation (NAT) rule. When the session count exceeds the upper (raise) threshold or falls below the lower (clear) threshold, an SNMP trap is triggered.

Options **clear-threshold *value***—Lower threshold at which an SNMP trap is triggered.

raise-threshold *value*—Upper threshold at which an SNMP trap is triggered.

NOTE: If you enter a value for raise-threshold but not for clear-threshold, clear-threshold is automatically set to 80 percent of raise-threshold.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

rule-set (Security Destination NAT)

Syntax

```
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      application {
        [application];
        any;
      }
      destination-address ip-address| destination-address-name address-name);
      destination-port (port-or-low <to high>);
      protocol [protocol-name-or-number];
      source-address [ip-address];
      source-address-name [address-name];
    }
    then {
      destination-nat (off | pool pool-name | rule-session-count-alarm (clear-threshold
        value | raise-threshold value));
      }
    {
  }
}
```

Hierarchy Level [edit security nat destination]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Configure a set of rules for destination NAT.

Options ***rule-set-name***—Name of the rule set.

description—Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule-set (Security Source NAT)

Syntax

```
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      application {
        [application];
        any;
      }
      (destination-address <ip-address> | destination-address-name <address-name>);
      destination-port (port-or-low <to high>);
      protocol [protocol-name-or-number];
      source-address [ip-address];
      source-address-name [address-name];
      source-port (port-or-low <to high>);
    }
    then {
      source-nat {
        interface {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number value;
            permit (any-remote-host | target-host | target-host-port);
          }
        }
        off;
        pool <pool-name>
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number number;
            permit (any-remote-host | target-host | target-host-port);
          }
        }
        rule-session-count-alarm (raise-threshold value | clear-threshold value);
      }
    }
  }
  to {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
}
```

Hierarchy Level [edit security nat source]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10.

Description Configure a set of rules for source NAT.

Options ***rule-set-name***—Name of the rule set.

description—Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

rule-set (Security Static NAT)

Syntax

```
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      (destination-address ip-address | destination-address-name address-name);
      destination-port (port | low to high);
      source-address ip-address;
      source-address-name address-name;
      source-port (port or low <to high>);
    }
    then {
      static-nat {
        inet {
          routing-instance (default | routing-instance-name);
        }
        prefix {
          address-prefix;
          mapped-port lower-port-range to upper-port-range;
          routing-instance (default | routing-instance-name);
        }
        prefix-name {
          address-prefix-name;
          mapped-port lower-port-range to upper-port-range;
          routing-instance (default | routing-instance-name);
          }
        rule-session-count-alarm (raise-threshold value | clear-threshold value);
      }
    }
  }
}
```

Hierarchy Level [edit security nat static]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. The **rule-session-count-alarm**, **source-address**, **source-address-name**, and **source-port** options added in Junos OS Release 12.1X45-D10.

Description Configure a set of rules for static NAT.

Options ***rule-set-name***—Name of the rule set.

description—Description of the rule set.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

source (Security Source NAT)

Syntax

```
source {
  address-persistent;
  interface (port-overloading off | port-overloading-factor number);
  pool pool-name {
    address ip-address {
      to ip-address;
    }
    address-persistent subscriber ipv6-prefix-length prefix-length;
    address-pooling (paired | no-paired);
    address-shared;
    description text;
    host-address-base ip-address;
    overflow-pool (interface | pool-name);
    pool-utilization-alarm (clear-threshold value | raise-threshold value);
    port {
      block-allocation {
        active-block-timeout timeout-interval;
        block-size block-size;
        log disable;
        maximum-blocks-per-host maximum-block-number
      }
      deterministic {
        block-size block-size;
        host {
          address ip-address;
          address-name address-name;
      }
      no-translation;
      port-overloading-factor number;
      range {
        port-low <to port-high>;
        to port-high;
        twin-port port-low <to port-high>;
      }
    }
    routing-instance routing-instance-name;
  }
  pool-default-port-range lower-port-range to upper-port-range;
  pool-default-twin-port-range lower-port-range to upper-port-range;
  pool-utilization-alarm (clear-threshold value | raise-threshold value);
  port-randomization disable;
  rule-set rule-set-name {
    description text;
    from {
      interface [interface-name];
      routing-instance [routing-instance-name];
      zone [zone-name];
    }
    rule rule-name {
      description text;
      match {
        application {
```


Copyright © 2016, Juniper Networks, Inc.

```
            [application];
            any;
          }
          (destination-address <ip-address> | destination-address-name <address-name>);
          destination-port (port-or-low <to high>);
          protocol [protocol-name-or-number];
          source-address [ip-address];
          source-address-name [address-name];
          source-port (port-or-low <to high>)
        }
        then source-nat
            interface {
              persistent-nat {
                address-mapping;
                inactivity-timeout seconds;
                max-session-number value;
                permit (any-remote-host | target-host | target-host-port);
              }
            }
            off;
            pool <pool-name>
              persistent-nat {
                address-mapping;
                inactivity-timeout seconds;
                max-session-number number;
                permit (any-remote-host | target-host | target-host-port);
              }
            rule-session-count-alarm (clear-threshold value | raise-threshold value);
          }
        }
      }
      to {
        interface [interface-name];
        routing-instance [routing-instance-name];
        zone [zone-name];
      }
    }
}
```

Hierarchy Level [edit security nat source pool pool-name port]

Release Information Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.1X47-D10. Statement modified in Junos OS Release 12.3X48-D10.

Copyright © 2016, Juniper Networks, Inc.

Description Configure source NAT, which allows you to configure the following:

- Translate source IP address or addresses to the egress interface's IP address.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is dynamic and without PAT.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is dynamic and with PAT.
- Translate a range of source IP addresses to another range of IP addresses. This mapping is one-to-one, static, and without PAT.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

source-address (Security Destination NAT)

Syntax source-address [*ip-address*];

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify source address to match the rule. You can configure multiple addresses or subnets.

Options ***ip-address*** —Source address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-address (Security Source NAT)

Syntax source-address [*ip-address*];

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify source address to match the rule. You can configure multiple addresses or subnets.

Options ***ip-address***—Source address or a subnet.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

source-address (Security Static NAT Rule Set)

Syntax source-address [*ip-address*];

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the source address to match the rule. Up to 8 addresses are supported.

Options ***ip-address*** —Source address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-address-name (Security Destination NAT)

Syntax source-address-name [*address-name*];

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a source address name to match the rule. You can configure multiple address names.

Options ***address-name***—Source address name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

source-address-name (Security Source NAT)

Syntax source-address-name [*address-name*];

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement modified in Junos OS Release 9.6.

Description Specify a source address name to match the rule. You can configure multiple address names.

Options ***address-name***—Source address name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-address-name (Security Static NAT Rule Set)

Syntax source-address-name [*address-name*];

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify a source address name to match the rule. Up to 8 address names are supported.

Options ***address-name***—Source address name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-nat

Syntax

```
source-nat {
  interface {
    persistent-nat {
      address-mapping;
      inactivity-timeout seconds;
      max-session-number value;
      permit (any-remote-host | target-host | target-host-port);
    }
  }
  off;
  pool <pool-name>;
    persistent-nat {
      address-mapping;
      inactivity-timeout seconds;
      max-session-number number;
      permit (any-remote-host | target-host | target-host-port);
  }
  rule-session-count-alarm (clear-threshold value | raise-threshold value);
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Specify the action of the source NAT rule.

Options

- **off**—Do not perform the source NAT operation.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security — To view this statement in the configuration.
security-control— To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

source-port (Security Source NAT Rule Set)

Syntax source-port (*port-or-low* <to *high*>);

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the port number or port range for a source rule. Up to 8 ports or port ranges are supported.

Options ***port***—Specify a port number.

low—Specify the lower limit of the port range.

<to *high*>—Specify the upper limit of the port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

source-port (Security Static NAT Rule Set)

Syntax source-port (*port or low* <to *high*>);

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* match]

Release Information Statement introduced in Junos OS Release 12.1X45-D10.

Description Specify the port or port range for a source rule. Up to 8 ports or port ranges are supported.

Options ***port***—Specify a port number.

low—Specify the lower limit of the port range.

<to *high*>—Specify the upper limit of the port range.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

static (Security NAT)

Syntax

```
static {
  rule-set rule-set-name {
    description text;
    from {
      interface [interface-name];
      routing-instance [routing-instance-name];
      zone [zone-name];
    }
    rule rule-name {
      description text;
      match {
        (destination-address <ip-address> | destination-address-name <address-name>);
        destination-port (port-or-low | <to high>);
        source-address [ip-address];
        source-address-name [ip-address-name];
        source-port (port-or-low <to high>);
      }
      then {
        static-nat {
          inet {
            routing-instance (routing-instance-name| default);
          }
          nptv6-prefix {
            address-prefix;
              routing-instance routing-instance-name;
          nptv6-prefix-name {
            address-prefix-name;
              routing-instance routing-instance-name;
          [
          prefix {
            address-prefix;
              mapped-port lower-port-range to upper-port-range;
            routing-instance (routing-instance-name| default);
          }
          prefix {
            address-prefix-name;
            mapped-port lower-port-range to upper-port-range;
            routing-instance (routing-instance-name| default);
          }
          rule-session-count-alarm (clear-threshold value | raise-threshold value);
        }
      }
    }
  }
}
```

Hierarchy Level [edit security nat]

Release Information Statement introduced in Junos OS Release 9.3. The **description** option added in Junos OS Release 12.1. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

Copyright © 2016, Juniper Networks, Inc.

Description Configure static NAT.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

static-nat

Syntax

```
static-nat {
  inet {
    routing-instance (default | routing-instance-name);
  }
  nptv6-prefix {
    address-prefix;
      routing-instance routing-instance-name;
  nptv6-prefix-name {
    address-prefix-name;
      routing-instance routing-instance-name;
  [
  prefix {
    address-prefix;
    routing-instance (default | routing-instance-name);
  }
  prefix-name {
    address-prefix-name;
    routing-instance (default | routing-instance-name);
  }
  rule-session-count-alarm (clear threshold value | raise threshold value);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name* then]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

Description Specify the translated address of the static NAT rule.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

to (Security Source NAT)

Syntax

```
to {
  interface [interface-name];
  routing-instance [routing-instance-name];
  zone [zone-name];
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the destination of the packet among the routing instance, interface, or zone.

Options

- **interface [*interface-name*]**—Name of the interface.
- **routing-instance [*routing-instance-name*]**—Name of the routing instance.
- **zone [*zone-name*]**—Name of the zone.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

then (Security Destination NAT)

Syntax

```
then {
  destination-nat (off | pool pool-name | rule-session-count-alarm (clear-threshold value |
    raise-threshold value));
}
```

Hierarchy Level [edit security nat destination rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Specify the action to be performed when traffic matches the destination NAT rule criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

then (Security Source NAT)

Syntax

```
then source-nat;
    interface {
      persistent-nat {
        address-mapping;
        inactivity-timeout seconds;
        max-session-number value;
        permit (any-remote-host | target-host | target-host-port);
      }
    }
    off;
    pool <pool-name>;
      persistent-nat {
        address-mapping;
        inactivity-timeout seconds;
        max-session-number number;
        permit (any-remote-host | target-host | target-host-port);
      }
    rule-session-count-alarm (clear-threshold value | raise-threshold value);
  }
}
```

Hierarchy Level [edit security nat source rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10.

Description Specify the action to be performed when traffic matches the source NAT rule criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

then (Security Static NAT)

Syntax

```
then {
  static-nat {
    inet {
      routing-instance (default | routing-instance-name);
    }
    nptv6-prefix {
      address-prefix;
        routing-instance routing-instance-name;
    nptv6-prefix-name {
      address-prefix-name;
        routing-instance routing-instance-name;
    [
    prefix {
      address-prefix;
      mapped-port lower-port-range to upper-port-range;
      routing-instance (default | routing-instance-name);
    }
    prefix-name {
      address-prefix-name;
      mapped-port lower-port-range to upper-port-range;
      routing-instance (default | routing-instance-name);
    }
  rule-session-count-alarm (clear-threshold value | raise-threshold value);
}
```

Hierarchy Level [edit security nat static rule-set *rule-set-name* rule *rule-name*]

Release Information Statement modified in Junos OS Release 9.6. Statement modified in Junos OS Release 12.1X45-D10. Statement modified in Junos OS Release 12.3X48-D25.

Description Specify the action to be performed when traffic matches the static NAT rule criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security NAT)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security nat]

Release Information Statement modified in Junos OS Release 9.6.

Description Configure NAT tracing options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
  - **destination-nat-pfe**—Trace destination NAT events on PFE-ukernel side
  - **destination-nat-re**—Trace destination NAT events on Routing Engine (RE) side
  - **destination-nat-rt**—Trace destination NAT events on Packet Forwarding Engine real-time (PFE-RT) side
  - **source-nat-pfe**—Trace source NAT events on PFE-ukernel side
  - **source-nat-re**—Trace source NAT events on RE side
  - **source-nat-rt**—Trace source NAT events on PFE-RT side
  - **static-nat-pfe**—Trace static NAT events on PFE-ukernel side
  - **static-nat-re**—Trace static NAT events on RE side
  - **static-nat-rt**—Trace static NAT events on PFE-RT side
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 210

Operational Commands

- clear security nat incoming-table
- clear security nat source persistent-nat-table
- clear security nat statistics destination pool
- clear security nat statistics destination rule
- clear security nat statistics source pool
- clear security nat statistics source rule
- clear security nat statistics static rule
- show security nat destination pool
- show security nat destination rule
- show security nat destination rule-application
- show security nat destination summary
- show security nat incoming-table
- show security nat interface-nat-ports
- show security nat resource-usage source-pool
- show security nat source deterministic
- show security nat source paired-address
- show security nat source persistent-nat-table
- show security nat source pool
- show security nat source port-block
- show security nat source rule
- show security nat source rule-application
- show security nat source summary
- show security nat static rule

Copyright © 2016, Juniper Networks, Inc.

clear security nat incoming-table

Syntax

clear security nat incoming-table
<node (*node-id* | all | local | primary)>

Release Information

Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description

Clear Network Address Translation (NAT) incoming table information.

Options

- none—Clear all information NAT incoming table.
- **node**—(Optional) For chassis cluster configurations, clear incoming table information on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level

clear

Related Documentation

- show security nat incoming-table on page 5136
- Security Configuration Statement Hierarchy on page 603

Output Fields

This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security nat source persistent-nat-table

Syntax clear security nat source persistent-nat-table
(all | interface | internal-ip *ip-address* <internal-port *port*> | pool *poolname*)

Release Information Command introduced in Junos OS Release 10.0.

Description Clear Network Address Translation (NAT) persistent NAT bindings that are in query mode, where all sessions of the binding are gone.

Options

- **all**—Clear all persistent NAT bindings that are in query mode.
- **interface**—Clear persistent NAT bindings that are in query mode for the specified interface.
- **internal-ip** ***ip-address***—Clear persistent NAT bindings for the specified internal IP address.
- **internal-ip** ***ip-address*** **internal-port** ***port***—Clear persistent NAT bindings that are in query mode for the specified internal IP address and port.
- **pool**—Clear persistent NAT bindings that are in query mode for the specified source NAT pool.

Required Privilege Level clear

Related Documentation

- show security nat source persistent-nat-table on page 5148
- Security Configuration Statement Hierarchy on page 603

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security nat statistics destination pool

Syntax

clear security nat statistics destination pool
<*pool-name*>
all

Release Information

Command introduced in Junos OS Release 11.1.

Description

Clear the destination NAT pool information.

Options

pool-name—Clear specified destination nat pool information.

all—Clear all destination nat pool information.

Required Privilege Level

clear

Related Documentation

- show security nat destination pool on page 5127
- show security nat destination summary on page 3701
- Security Configuration Statement Hierarchy on page 603

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

clear security nat statistics destination rule

Syntax

clear security nat statistics destination rule
<rule-name>
alll

Release Information

Command introduced in Junos OS Release 11.1.

Description

Clear the destination NAT rule information.

Options

rule-name—Clear specified destination nat rule-set information.

all—Clear all destination nat rule-set information.

Required Privilege Level

clear

Related Documentation

- show security nat destination rule on page 3698
- show security nat destination summary on page 3701
- Security Configuration Statement Hierarchy on page 603

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

clear security nat statistics source pool

Syntax

clear security nat statistics source pool
<*pool-name*>
all

Release Information

Command introduced in Junos OS Release 11.1.

Description

Clear the source NAT statistic pool information.

Options

pool-name—Clear the specified source nat pool information.

all—Clear all source pool information.

Required Privilege Level

clear

Related Documentation

- show security nat source pool on page 5151
- show security nat source summary on page 3707
- Security Configuration Statement Hierarchy on page 603

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

clear security nat statistics source rule

Syntax

clear security nat statistics source rule
<*rule-name*>
all

Release Information

Command introduced in Junos OS Release 11.1.

Description

Clear the source NAT statistic rule-set information.

Options

rule-name—Clear the specified source rule-set information.

all—Clear all source nat rule-set information.

Required Privilege Level

clear

Related Documentation

- show security nat source summary on page 3707
- show security nat source rule on page 3703
- Security Configuration Statement Hierarchy on page 603

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

clear security nat statistics static rule

Syntax

clear security nat statistics static rule
<*rule-name*>
all

Release Information

Command introduced in Junos OS Release 11.1.

Description

Clear the static NAT rule-set information.

Options

rule-name—Clear specified static nat rule-set information.

all—Clear all static nat rule-set information.

Required Privilege Level

clear

Related Documentation

- show security nat static rule on page 3709
- Security Configuration Statement Hierarchy on page 603

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2016, Juniper Networks, Inc.

show security nat destination pool

Syntax

show security nat destination pool
pool-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

Description

Display information about the specified Network Address Translation (NAT) destination address pool.

Options

pool-name—Name of the destination address pool.

all—Display information about all the destination NAT address pools.

logical-system (*logical-system-name* | all)—Display information about the destination NAT pools for the specified logical system or for all logical systems.

root-logical-system—Display information about the destination NAT pools for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Destination NAT) on page 5075
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat destination pool dst-nat-pool1 on page 5128
show security nat destination pool all on page 5128

Output Fields

Table 448 lists the output fields for the **show security nat destination pool** command. Output fields are listed in the approximate order in which they appear.

Table 448: show security nat destination pool Output Fields

Field Name	Field Description
Pool name	Name of the destination pool.
Description	Description of the destination pool.
Pool id	Pool identification number.
Routing instance	Name of the routing instance.
Total address	Number of IP addresses that are in use.

Copyright © 2016, Juniper Networks, Inc.

Table 448: show security nat destination pool Output Fields *(continued)*

Field Name	Field Description
Translation hits	Number of translation hits.
Address range	IP address or IP address range for the pool.

Sample Output

show security nat destination pool dst-nat-pool1

```
user@host> show security nat destination pool dst-p1

Pool name       : dst-p1
Description     : The destination pool dst-p1 is for the sales team
Pool id         : 1
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                        Port
     203.0.113.1 -203.0.113.1               0
```

Sample Output

show security nat destination pool all

```
user@host> show security nat destination pool all

Total destination-nat pools: 2

Pool name       : dst-p1
Description     : The destination pool dst-p1 is for the sales team
Pool id         : 1
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                        Port
     203.0.113.1 -203.0.113.1               0

Pool name       : dst-p2
Description     : The destination pool dst-p2 is for the sales team
Pool id         : 2
Routing instance: default
Total address   : 1
Translation hits: 0
Address range                        Port
     2001:db8::1 - 2001:db8::1               0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat destination rule

Syntax

show security nat destination rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the **application** option field added in Junos OS Release 12.1X47-D10.

Description

Display information about the specified destination Network Address Translation (NAT) rule.

Options

rule-name—Display information about the specified destination NAT rule.

all—Display information about all the destination NAT rules.

logical-system (*logical-system-name* | all)—Display information about the destination NAT rules for the specified logical system or for all logical systems.

root-logical-system—Display information about the destination NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Destination NAT) on page 5092
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat destination rule dst2-rule on page 5130
show security nat destination rule all on page 5131

Output Fields

Table 365 lists the output fields for the **show security nat destination rule** command. Output fields are listed in the approximate order in which they appear.

Table 449: show security nat destination rule Output Fields

Field Name	Field Description
Total destination-nat rules	Number of destination NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly as address names and as address set names in the rule.
Destination NAT rule	Name of the destination NAT rule.

Copyright © 2016, Juniper Networks, Inc.

Table 449: show security nat destination rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the destination NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the destination NAT rule.
From routing instance	Name of the routing instance from which the packets flow.
From interface	Name of the interface from which the packets flow.
From zone	Name of the zone from which the packets flow.
Source addresses	Name of the source addresses that match the rule. The default value is any.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: • **destination NAT pool**—Use user-defined destination NAT pool to perform destination NAT. • **off**—Do not perform destination NAT.
Destination ports	Destination ports number that match the rule. The default value is any.
Application	Indicates whether the application option is configured.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat destination rule dst2-rule

```
user@host>show security nat destination rule dst2-rule

 Destination NAT rule: dst2-rule           Rule-set: dst2
  Description               : The destination rule dst2-rule is for the sales
team
  Rule-Id                   : 1
  Rule position             : 1
  From routing instance     : ri1
                            : ri2
  Match
    Source addresses        : add1
```


Copyright © 2016, Juniper Networks, Inc.

```
                                add2
    Destination addresses     : add9
  Action                      : off

  Destination port            : 0
 Translation hits             : 68
   Successful sessions        : 25
   Failed sessions            : 43
 Number of sessions           : 2
```

Sample Output

show security nat destination rule all

```
user@host> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0

Destination NAT rule: r4                    Rule-set: rs4
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : untrust
  Match
    Source addresses        : 192.0.2.0 - 192.0.2.255
    Destination addresses   : 198.51.100.0 - 198.51.100.255
    Application             : configured
  Action                    : off
  Translation hits          : 0
    Successful sessions     : 0
    Failed sessions         : 0
  Number of sessions        : 0
```

Copyright © 2016, Juniper Networks, Inc.

show security nat destination rule-application

Syntax

show security nat destination rule-application
rule-name
all
logical-system *logical-system-name*
root-logical-system

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Display information about the specified destination Network Address Translation (NAT) rule application.

Options

rule-name—Display information about the specified destination NAT rule application.

all—Display information about all the destination NAT rule applications.

logical-system *logical-system-name* —Display information about the destination NAT rule applications for the specified logical system.

root-logical-system—Display information about the destination NAT rule applications for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- *Logical Systems Feature Guide for Security Devices*

List of Sample Output

show security nat destination rule-application for port application on page 5133
show security nat destination rule-application for ICMP application on page 5133

Output Fields

Table 450 lists the output fields for the **show security nat destination rule-application** command. Output fields are listed in the approximate order in which they appear.

Table 450: show security nat destination rule-application Output Fields

Field Name	Field Description
Destination NAT rule	Name of the destination NAT rule.
Rule-set	Rule set identification number.
Rule-Id	Rule identification number.
Application	Name of the application or application set.
IP protocol	IP protocol identifier.
Source port range	Source port range identifier.

Copyright © 2016, Juniper Networks, Inc.

Table 450: show security nat destination rule-application Output Fields *(continued)*

Field Name	Field Description
Destination port range	Destination port identifier.
ICMP information	• type—ICMP message type. • code—Code corresponding to the ICMP message type.

Sample Output

show security nat destination rule-application for port application

```
user@host>show security nat destination rule-application all

 Destination NAT rule: r4                   Rule-set: rs4
  Rule-Id                   : 2
    Application: app-set1
      IP protocol: 17
      Source port range: [40-50]
      Destination port range: [50-60]
      IP protocol: 17
      Source port range: [100-200]
      Destination port range: [300-500]
```

Sample Output

show security nat destination rule-application for ICMP application

```
user@host>show security nat destination rule-application all

Destination NAT rule: r1                   Rule-set: rs1
  Rule-Id                   : 1
    Application: junos-icmp-all
      IP protocol: icmp
      ICMP Information: type=255, code=0
    Application: icmp1
      IP protocol: icmp
      ICMP Information: type=1, code=1
    Application: junos-icmp6-all
      IP protocol: 58
      ICMP Information: type=255, code=0
```

Copyright © 2016, Juniper Networks, Inc.

show security nat destination summary

Syntax

show security nat destination summary
<logical-system (*logical-system-name* | all)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

Description

Display a summary of Network Address Translation (NAT) destination pool information.

Options

none—Display summary information about the destination NAT pool.

logical-system (*logical-system-name* | all)—Display summary information about the destination NAT for the specified logical system or for all logical systems.

root-logical-system—Display summary information about the destination NAT for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Destination NAT) on page 5075
- rule (Security Destination NAT) on page 5092
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat destination summary on page 5135

Output Fields

Table 366 lists the output fields for the **show security nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 451: show security nat destination summary Output Fields

Field Name	Field Description
Total destination nat pool number	Number of destination NAT pools.
Pool name	Name of the destination address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.
Total	Number of IP addresses that are in use.
Available	Number of IP addresses that are free for use.
Total destination nat rule number	Number of destination NAT rules.

Copyright © 2016, Juniper Networks, Inc.

Table 451: show security nat destination summary Output Fields *(continued)*

Field Name	Field Description
Total hit times	Number of times a translation in the translation table is used for all the destination NAT rules.
Total fail times	Number of times a translation in the translation table failed to translate for all the destination NAT rules.

Sample Output

show security nat destination summary

```
user@host> show security nat destination summary

Total pools: 2
Pool name           Address                     Routing        Port  Total
                    Range                       Instance             Address
dst-p1              203.0.113.1 -203.0.113.1        default        0     1
dst-p2              2001:db8::1 - 2001:db8::1       default        0     1

Total rules: 171
Rule name           Rule set       From                          Action
dst2-rule           dst2           ri1
                                   ri2
                                   ri3
                                   ri4
                                   ri5
                                   ri6
                                   ri7
dst3-rule           dst3           ri9                           off
                                   ri1
                                   ri2
                                   ri3
                                   ri4
                                   ri5

...
```

show security nat incoming-table

Syntax **show security nat incoming-table**
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display Network Address Translation (NAT) table information.

NOTE: The incoming dip NAT table is replaced with ALG cone NAT binding table and the **show security nat incoming-table** command is obsolete from Junos OS Release 11.2 onward. The **show security nat incoming-table** command works as is in the previous releases.

Options

- none—Display all information NAT incoming table.
- **node**—(Optional) For chassis cluster configurations, display incoming table information on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- clear security nat incoming-table on page 5120
- Security Configuration Statement Hierarchy on page 603

List of Sample Output show security nat incoming-table on page 5137

Output Fields Table 452 lists the output fields for the **show security nat incoming-table** command. Output fields are listed in the approximate order in which they appear.

Table 452: show security nat incoming-table Output Fields

Field Name	Field Description
In use	Number of entries in the NAT table.
Maximum	Maximum number of entries possible in the NAT table.
Entry allocation failed	Number of entries failed for allocation.

Copyright © 2016, Juniper Networks, Inc.

Table 452: show security nat incoming-table Output Fields *(continued)*

Field Name	Field Description
Destination	Destination IP address and port number.
Host	Host IP address and port number that the destination IP address is mapped.
References	Number of sessions referencing the entry.
Timeout	Timeout, in seconds, of the entry in the NAT table.
Source-pool	Name of source pool where translation is allocated.

Sample Output

show security nat incoming-table

```
user@host> show security nat incoming-table
In use: 1, Maximum: 1024, Entry allocation failed: 0
Destination           Host                  References Timeout Source-pool
10.1.1.26:1028        203.0.113.10:5060              1    3600 p1
```

show security nat interface-nat-ports

Syntax

show security nat interface-nat-ports
<node (*node-id* | all | local | primary)>
<logical-system (*logical-system-name* | all)>

Release Information

Command modified in Junos OS Release 9.2. The **node** options added in Junos OS Release 9.0. Logical system support added in Junos OS Release 12.1X45-D10.

Description

Display port usage for an interface source pool for Network Address Translation (NAT).

Options

none—Display all port usage information for an interface source pool.

node—For chassis cluster configurations, display interface NAT ports information on a specific node.

node-id—Identification number of the node. It can be 0 or 1.

all—Display information about all nodes.

local—Display information about the local node.

primary—Display information about the primary node.

logical-system (*logical-system-name* | all)—Display port usage information for the specified logical system or for all logical systems.

Required Privilege Level

view

Related Documentation

- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat interface-nat-ports on page 5139
show security nat interface-nat-ports logical-system all on page 5139

Output Fields

Table 453 lists the output fields for the **show security nat interface-nat-ports** command. Output fields are listed in the approximate order in which they appear.

Table 453: show security nat interface-nat-ports Output Fields

Field Name	Field Description
Pool Index	Port pool index.
Total Ports	Total number of ports in a port pool. In SRX Series devices, 10 interface NAT ports are supported.
Single Ports Allocated	Number of ports allocated one at a time that are in use.
Single Ports Available	Number of ports allocated one at a time that are free for use.
Twin Ports Allocated	Number of ports allocated two at a time that are in use.

Copyright © 2016, Juniper Networks, Inc.

Table 453: show security nat interface-nat-ports Output Fields *(continued)*

Field Name	Field Description
Twin Ports Available	Number of ports allocated two at a time that are free for use.

Sample Output

show security nat interface-nat-ports

```
user@host> show security nat interface-nat-ports
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
    index  ports     allocated     available   allocated   available
     0  64510             0         63486           0        1024
     1  64510             0         63486           0        1024
     2  64510             0         63486           0        1024
     3  64510             0         63486           0        1024
     4  64510             0         63486           0        1024
     5  64510             0         63486           0        1024
     6  64510             0         63486           0        1024
     7  64510             0         63486           0        1024
     8  64510             0         63486           0        1024
     9  64510             0         63486           0        1024
```

Sample Output

show security nat interface-nat-ports logical-system all

```
user@host> show security nat interface-nat-ports logical-system all
Logical system: root-logical-system
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
index  ports     allocated     available   allocated   available
    0  64510             0         63486           0        1024
Logical system: LSYS1
Pool   Total  Single ports  Single ports  Twin ports  Twin ports
index  ports     allocated     available   allocated   available
    0  64510             0         63486           0        1024
    1  64510             0         63486           0        1024
    2  64510             0         63486           0        1024
    3  64510             0         63486           0        1024
    4  64510             0         63486           0        1024
    5  64510             0         63486           0        1024
    6  64510             0         63486           0        1024
    7  64510             0         63486           0        1024
    8  64510             0         63486           0        1024
    9  64510             0         63486           0        1024
   10  64510             0         63486           0        1024
   11  64510             0         63486           0        1024
   12  64510             0         63486           0        1024
   13  64510             0         63486           0        1024
   14  64510             0         63486           0        1024
   15  64510             0         63486           0        1024
   16  64510             0         63486           0        1024
   17  64510             0         63486           0        1024
   18  64510             0         63486           0        1024
   19  64510             0         63486           0        1024
   20  64510             0         63486           0        1024
   21  64510             0         63486           0        1024
   22  64510             0         63486           0        1024
```

Copyright © 2016, Juniper Networks, Inc.

23	64510	0	63486	0	1024
24	64510	0	63486	0	1024
25	64510	0	63486	0	1024
26	64510	0	63486	0	1024
27	64510	0	63486	0	1024
28	64510	0	63486	0	1024
29	64510	0	63486	0	1024
30	64510	0	63486	0	1024
31	64510	0	63486	0	1024
32	64510	0	63486	0	1024
33	64510	0	63486	0	1024
34	64510	0	63486	0	1024
35	64510	0	63486	0	1024
36	64510	0	63486	0	1024
37	64510	0	63486	0	1024
38	64510	0	63486	0	1024
39	64510	0	63486	0	1024
40	64510	0	63486	0	1024
41	64510	0	63486	0	1024
42	64510	0	63486	0	1024
43	64510	0	63486	0	1024
44	64510	0	63486	0	1024
45	64510	0	63486	0	1024
45	64510	0	63486	0	1024
46	64510	0	63486	0	1024
47	64510	0	63486	0	1024
48	64510	0	63486	0	1024
49	64510	0	63486	0	1024
50	64510	0	63486	0	1024
51	64510	0	63486	0	1024
52	64510	0	63486	0	1024
53	64510	0	63486	0	1024
54	64510	0	63486	0	1024
55	64510	0	63486	0	1024
56	64510	0	63486	0	1024
57	64510	0	63486	0	1024
58	64510	0	63486	0	1024
59	64510	0	63486	0	1024

Copyright © 2016, Juniper Networks, Inc.

show security nat resource-usage source-pool

Syntax
show security nat resource-usage source-pool
all | *source-pool-name*
logical-system *logical-system-name* | root logical system

Release Information
Command introduced in Junos OS Release 12.1X45-D10.

Description
Display source NAT pool usage information. In pools without Port Address Translation (PAT), information about IP addresses is displayed. In pools with PAT, information about ports is displayed.

Options
all—Display resource use information for all source NAT pools.

source-pool-name—Display resource use information for the specified source NAT pool.

logical-system *logical-system-name*—Display resource use information for the source NAT pools in the specified logical system.

root-logical-system—Display resource use information for the source NAT pools in the root logical system.

Required Privilege Level
view

Related Documentation
- clear security nat statistics source pool on page 5124
- Security Configuration Statement Hierarchy on page 603

List of Sample Output
show security nat resource-usage resource-pool all on page 5142
show security nat resource-usage resource-pool pool-name (Without PAT) on page 5142
show security nat resource-usage resource-pool pool-name (with PAT) on page 5142

Output Fields
Table 454 lists the output fields for the **show security nat resource-usage source-pool** command. Output fields are listed in the approximate order in which they appear. You can use the **clear security nat statistics** command to reset the peak usage statistics.

Table 454: show security nat resource-usage source-pool Output Fields

Field Name	Field Description
Pool	Name of the pool.
Address	Address of the pool.
Factor-index	Port pool index.
Total address	Number of addresses in the pool.
Port-range	Number of ports allocated at a time.

Copyright © 2016, Juniper Networks, Inc.

Table 454: show security nat resource-usage source-pool Output Fields *(continued)*

Field Name	Field Description
Used	Number of used resources in the pool.
Avail	Number of available resources in the pool.
Usage	Percent of resources used. In a PAT pool, use includes single and twin ports.
Current usage	Percent of current resources used.
Peak usage	Percent of resources used during the peak date and time.
Total	Number of used and available resources.
Total ports	Number of used and available ports.
Port-overloading-factor	Port overloading capacity for the pool.

Sample Output

show security nat resource-usage resource-pool all

```
user@host> show security nat resource-usage source-pool all

PAT pools(including address-shared pool) port utilization:
Pool      Address          Used       Avail      Total      Usage
SpoolA        512       2387968    29593600   31981568         7%
SpoolB        128        393216      655360    1048576        38%

Non-PAT pools address utilization:
Pool         Used    Avail    Total   Usage
Spool1        300     3796     4096      7%
Spool2        512      512     1024     50%
```

show security nat resource-usage resource-pool pool-name (Without PAT)

```
user@host> show security nat resource-usage source-pool Spool1
Logical system: root
Peak usage: 60% @ 2012-08-26 20:16:20 UTC

Pool         Used     Avail     Total    Usage
Spool1        300      3796      4096       7%
```

show security nat resource-usage resource-pool pool-name (with PAT)

```
user@host> show security nat resource-usage source-pool sp3
Logical system: root
Pool name: sp3
Total address: 2
Port-overloading-factor: 2
Total ports: 258048 Used: 60563  Avail: 197485
```


Copyright © 2016, Juniper Networks, Inc.

```
Current usage: 23%   Peak usage: 35% at 2012-11-12 20:15:26 CST
  Address    Factor-index Port-range          Used       Avail       Total Usage

  192.0.2.113
             0            Single Ports       30001       32463       62464   48%

             -            Alg Ports            462        1586        2048   22%

             1            Single Ports           0       62464       62464    0%

             -            Alg Ports              0        2048        2048    0%

             Sum          Single Ports       30001       94927      124928   24%

             -            Alg Ports            462        3634        4096   11%

  192.0.2.114
             0            Single Ports       29600       32864       62464   47%

             -            Alg Ports            500        1548        2048   24%

             1            Single Ports           0       62464       62464    0%

             -            Alg Ports              0        2048        2048    0%

             Sum          Single Ports       29600       95328      124928   23%

             -            Alg Ports            500        3596        4096   12%
```

Copyright © 2016, Juniper Networks, Inc.

show security nat source deterministic

Syntax

show security nat source deterministic
pool-name
host-ip *host ip address*
host-address-range
xlated-ip *xlated-ip-address*
xlated-port *xlated-port*
node
root-logical-system | logical-system {*lsys-name* | all }

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Verify the mapping relation when Deterministic-Nat is on.

Options

host-address-range—Display deterministic host address range without overlap.

pool-name—Display Deterministic NAT port block table for the specified source pool name.

node—Display source NAT deterministic port block table on specific node.

host ip address—Display deterministic NAT port block table based on internal host ip address.

xlated ip address—Display deterministic NAT port block table based on translated IP address.

xlated-port—Display deterministic NAT port block table based on translated IP and port; *xlated-port* can be used only with *xlated-ip* together for display.

root-logical-system—Display information about the source NAT pools for the master (root) logical system.

logical-system (*lsys-name* | all)—Display information about the specified logical system source NAT pools or all logical system source NAT pools.

Required Privilege Level

view

Related Documentation

- show security nat source pool on page 5151
- show security nat source port-block on page 5155
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source deterministic on page 5145

Output Fields

Table 455 lists the output fields for the **show security nat source deterministic** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 455: show security nat source deterministic Output Fields

Field Name	Field Description
Pool name	Name of pool.
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.
Used/total port blocks	Port block used number and port block total number for this source NAT pool.
Host IP	IP address of host.
External IP	IP address of external router.
Port_Range	The range of ports in a block, ranging from lowest to highest.
Ports_Used/Total	Number of ports used and total ports.

Sample Output

show security nat source deterministic

```
user@host> show security nat source deterministic
Pool name: SRC_P_3
Port-overloading-factor:     1     Port block size: 10000
Used/total port blocks: 0/12
Host_IP                      External_IP                 Port_Range
Ports_Used/Total
10.2.0.1                     203.0.113.1                      1024-11023
     0/10000*1
10.2.0.2                     203.0.113.1                     11024-21023
     0/10000*1
```

Copyright © 2016, Juniper Networks, Inc.

show security nat source paired-address

Syntax

show security nat source paired-address
<internal-ip *internal-ip-address*>
<logical-system *logical-system-name*>
<pool-name *pool-name*>
<root-logical-system>
<xlated-ip *x-lated-ip-address*>

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

Description

Display information about the Network Address Translation (NAT) source paired addresses.

Options

none—Display all paired IP address information.

internal ip *internal-ip-address*—Display information about the specified internal IP address.

logical-system *logical-system-name*—Display information about the source NAT pools for the specified logical system.

pool-name *pool-name*—Display paired address information for the specified pool.

root-logical-system—Display information about the source NAT pools for the master (root) logical system.

x-lated-ip *x-lated-ip-address*—Display information about the specified translated external IP address.

Additional Information

Required Privilege Level

view

Related Documentation

- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source paired-address on page 5147
show security nat source paired-address pool-name on page 5147
show security nat source paired-address pool-name internal-ip on page 5147
show security nat source paired-address pool-name xlated-ip on page 5147

Output Fields

Table 456 lists the output fields for the **show security nat source paired-address** command. Output fields are listed in the approximate order in which they appear.

Table 456: show security nat source paired-address Output Fields

Field Name	Field Description
Pool name	Name of the source pool.
Internal address	Internal IP address.

Copyright © 2016, Juniper Networks, Inc.

Table 456: show security nat source paired-address Output Fields *(continued)*

Field Name	Field Description
External address	External IP address.

Sample Output

show security nat source paired-address

```
user@host> show security nat source paired-address
Pool name: sp1
Internal address          External address
198.51.100.240                203.0.113.105

Pool name: sp2
Internal address          External address
198.51.100.126                192.0.2.1
198.51.100.127                192.0.2.1
198.51.100.125                192.0.2.1
198.51.100.130                192.0.2.1
198.51.100.128                192.0.2.1
198.51.100.129                192.0.2.1
```

show security nat source paired-address pool-name

```
user@host> show security nat source paired-address pool-name sp1
Pool name: sp1
Internal address        External address
192.168.1.1             192.0.2.1
192.168.1.2             192.0.2.2
192.168.1.3             192.0.2.3
```

show security nat source paired-address pool-name internal-ip

```
user@host> show security nat source paired-address pool-name sp1 internal-ip 192.168.1.1
Pool name: sp1
Internal address        External address
192.168.1.1             192.0.2.1
```

show security nat source paired-address pool-name xlated-ip

```
user@host> show security nat source paired-address pool-name sp1 xlated-ip 192.0.2.2
Pool name: sp1
Internal address        External address
192.168.1.2             192.0.2.2
```

Copyright © 2016, Juniper Networks, Inc.

show security nat source persistent-nat-table

Syntax show security nat source persistent-nat-table (all | interface | internal-ip *ip-address* <internal-port *port*> | pool *poolname*)

Release Information Command introduced in Junos OS Release 9.6. Support for IPv6 addresses added in Junos OS Release 11.2.

Description Display a summary of persistent Network Address Translation (NAT) information.

Options

- **all**—Display all persistent NAT bindings.
- **interface**—Display persistent NAT bindings for the interface.
- **internal-ip *ip-address***—Display persistent NAT bindings for the specified internal IP address.
- **internal-ip *ip-address* internal-port *port***—Display persistent NAT bindings for the specified internal IP address and port.
- **pool**—Display persistent NAT bindings for the specified source NAT pool.
- **summary**—Display persistent NAT bindings summary.

Required Privilege Level view

Related Documentation

- clear security nat source persistent-nat-table on page 5121
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source persistent–nat–table internal-ip internal-port on page 5149
show security nat source persistent–nat–table all on page 5149
show security nat source persistent-nat-table summary on page 5149

Output Fields Table 457 lists the output fields for the **show security nat source persistent–nat–table** command. Output fields are listed in the approximate order in which they appear.

Table 457: show security nat source persistent–nat–table Output Fields

Field Name	Field Description
Internal IP/Port	Internal transport IP address and port number of the outgoing session from internal to external.
Reflexive IP/Port	Translated IP address and port number of the source IP address and port.
Source NAT Pool	The name of the source pool where persistent NAT is used.
Type	Persistent NAT type.
Left_time/Conf_time	The inactivity timeout period that remains and the configured timeout value.

Copyright © 2016, Juniper Networks, Inc.

Table 457: show security nat source persistent–nat–table Output Fields *(continued)*

Field Name	Field Description
Current_Sess_Num/Max_Sess_Num	The number of current sessions associated with the persistent NAT binding.
Source NAT Rule	Name of the source NAT rule to which this persistent NAT binding applies.

Sample Output

show security nat source persistent–nat–table internal-ip internal-port

```
user@host> show security nat source persistent–nat–table internal-ip 192.0.2.1 internal-port
60784


Internal                          Reflective        Source     Type
Left_time/  Curr_Sess_Num/ Source
 In_IP  In_Port I_Proto Ref_IP    Ref_Port R_Proto NAT Pool
Conf_time   Max_Sess_Num  NAT Rule
192.0.2.1  60784   udp  198.51.100.68  60784     udp   dynamic-customer-source
any-remote-host  254/300  0/30 105
```

Sample Output

show security nat source persistent–nat–table all

```
user@host> show security nat source persistent–nat–table all
 Internal             Reflective                   Source     Type
Left_time/  Curr_Sess_Num/  Source
 In_IP     In_Port I_Proto Ref_IP       Ref_Port R_Proto NAT Pool
     Conf_time   Max_Sess_Num    NAT Rule
192.0.2.1    63893   tcp    198.51.100.68   63893     tcp
dynamic-customer-source any-remote-host  192/300   0/30 105
192.0.2.1    64014   udp    198.51.100.68   64014     udp
dynamic-customer-source any-remote-host  244/300   0/30 105
192.0.2.1    60784   udp    198.51.100.68   60784     udp
dynamic-customer-source any-remote-host  254/300   0/30 105
192.0.2.1    57022   udp    198.51.100.68   57022     udp
dynamic-customer-source any-remote-host  264/300   0/30 105
192.0.2.1    53009   udp    198.51.100.68   53009     udp
dynamic-customer-source any-remote-host  268/300   0/30 105
192.0.2.1    49225   udp    198.51.100.68   49225     udp
dynamic-customer-source any-remote-host  272/300   0/30 105
192.0.2.1    52150   udp    198.51.100.68   52150     udp
dynamic-customer-source any-remote-host  274/300   0/30 105
192.0.2.1    59770   udp    198.51.100.68   59770     udp
dynamic-customer-source any-remote-host  278/300   0/30 105
192.0.2.1    61497   udp    198.51.100.68   61497     udp
dynamic-customer-source any-remote-host  282/300   0/30 105
192.0.2.1    56843   udp    198.51.100.68   56843     udp
dynamic-customer-source any-remote-host    -/300   1/30 105
```

Sample Output

show security nat source persistent-nat-table summary

```
user@host> show security nat source persistent-nat-table summary
```

Copyright © 2016, Juniper Networks, Inc.

```
Persistent NAT Table Statistics on FPC5 PIC0:
binding total : 65536
binding in use : 0
enode total : 524288
enode in use : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source pool

Syntax

show security nat source pool
pool-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. **Description** output field added in Junos OS Release 12.1. The **Address assignment** output field and IPv6 logical system support added in Junos OS Release 12.1X45-D10. The **twin-port** output field added in Junos OS Release 12.1X47-D10. The **Address-persistent** output field added in Junos OS Release 12.3X48-D10.

Description

Display information about the specified Network Address Translation (NAT) source address pool and the configured twin port range per pool.

Options

pool-name—Display source NAT information for the specified address pool.

all—Display information about all source NAT address pools.

logical-system (*logical-system-name* | all)—Display information about the specified logical system source NAT pools or all logical system source NAT pools.

root-logical-system—Display information about the source NAT pools for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Source NAT) on page 5076
- clear security nat statistics source pool on page 5124
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source pool src-p1 on page 5152
show security nat source pool all on page 5153
show security nat source pool sp1 on page 5154
show security nat source pool P_1 on page 5154
show security nat source pool src-nat-v4-with-pat on page 5154

Output Fields

Table 458 lists the output fields for the **show security nat source pool** command. Output fields are listed in the approximate order in which they appear.

Table 458: show security nat source pool Output Fields

Field Name	Field Description
Pool name	Name of the source pool.
Description	Description of the source pool.

Copyright © 2016, Juniper Networks, Inc.

Table 458: show security nat source pool Output Fields *(continued)*

Field Name	Field Description
Pool id	Pool identification number.
Routing Instance	Name of the routing instance.
Host address base	Base address of the original source IP address range.
Port	Port numbers used for the source pool.
Twin port	Upper and lower limits of the twin port.
port overloading	Number of port overloading for the source pool.
Address assignment	Type of address assignment.
Total addresses	Number of IP addresses that are in use.
Translation hits	Number of translation hits.
Port block size	Block size for the deterministic pool.
Determ host range num	Host range for the deterministic pool.
Address range	IP address or IP address range for the source pool.
Address-Persistent	Address persistent information for IPv4 source pools: • IPv6 prefix length–Configured IPv6 prefix length. • IPv6 subscriber out of port–Number of port allocation failures.
Single Ports	Number of allocated single ports.
Twin Ports	Number of allocated twin ports.

Sample Output

show security nat source pool src-p1

```
user@host> show security nat source pool src-p1

Pool name          : src-p1
Description        : The source pool src-p1 is for the sales team
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : paired
port overloading   : 1
Total addresses    : 4
Translation hits   : 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Address range                          Single Ports   Twin Ports
            203.0.113.0 - 203.0.113.0           0             0
```

Sample Output

show security nat source pool all

```
user@host> show security nat source pool all

Total pools: 4

Pool name          : src-p1
Description        : The source pool src-p1 is for the sales team
Pool id            : 4
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : paired
port overloading   : 1
Total addresses    : 4
Translation hits   : 0
Address range                          Single Ports   Twin Ports
            203.0.113.0 - 203.0.113.0           0             0

Pool name          : src-p2
Description        : The source pool src-p2 is for the sales team
Pool id            : 5
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 4
Translation hits   : 0
Address range                        Single Ports   Twin Ports
            192.0.2.0 - 192.0.2.3           0             0

Pool name          : src-p3
Description        : The source pool src-p3 is for the sales team
Pool id            : 6
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 1
Translation hits   : 0
Address range                          Single Ports   Twin Ports
            2001:db8::1 - 2001:db8::1           0             0

Pool name          : src-p4
Description        : The source pool src-p4 is for the sales team
Pool id            : 7
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Address assignment : no-paired
port overloading   : 1
Total addresses    : 1
Translation hits   : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
Address range                          Single Ports   Twin Ports
            2001:db8::2 - 2001:db8::2           0             0
```

show security nat source pool sp1

```
user@host>show security nat source pool sp1
Pool name          : sp1
Description        : The source pool src-p1 is for the sales team
Pool id            : 12
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Twin port          : [63488, 64515]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 1
Translation hits   : 0
Address range                        Single Ports   Twin Ports
            192.0.2.1 - 192.0.2.1           0             0
```

show security nat source pool P_1

```
user@host>show security nat source pool P_1
Pool name          : P_1
Pool id            : 4
Routing instance   : default
Port               : [12345, 17890]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 256
Translation hits   : 0
Port block size      : 1000
Determ host range num: 3
Address range                       Single Ports   Twin Ports
         203.0.113.0 - 203.0.113.255            0            0
```

show security nat source pool src-nat-v4-with-pat

```
user@host>how security nat source pool src-nat-v4-with-pat

Pool name          : src-nat-v4-with-pat
Pool id            : 5
Routing instance   : default
Host address base  : 0.0.0.0
Port               : [1024, 63487]
Port overloading   : 1
Address assignment : no-paired
Total addresses    : 10
Translation hits   : 0
Address-persistent
        IPv6 prefix length: 64
        IPv6 subscriber out of port: 0
Address range                       Single Ports   Twin Ports
            203.0.113.1 - 203.0.113.10          0             0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source port-block

Syntax

show security nat source port-block
pool-name
host-ip *host ip address*
xlated-ip *xlated-ip-address*
xlated-port *xlated-port*
root-logical-system | logical-system {*lsys-name* | all}

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Display the port blocks allocated by the host.

Options

pool-name—Display the PBA port block table for the specified source pool.

host ip address—Display the PBA port block table based on the host-ip address.

xlated ip address—Display the PBA port block table based on the translated IPaddress.

xlated-port—Display the PBA port block table based on the translated IP address and the translated port information.

root-logical-system—Display the PBA port block table for the master (root) logical system.

logical-system (*lsys-name* | all)—Display information about the specified logical system source NAT pools or all logical system source NAT pools.

Required Privilege Level

view

Related Documentation

- show security nat source pool on page 5151
- show security nat source deterministic on page 5144
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source port-block on page 5156

Output Fields

Table 459 lists the output fields for the **show security nat source port-block** command. Output fields are listed in the approximate order in which they appear.

Table 459: show security nat source port-block Output Fields

Field Name	Field Description
Pool name	Name of pool.
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.
Max port blocks per host	Maximum number of blocks that one host can use for translation.

Copyright © 2016, Juniper Networks, Inc.

Table 459: show security nat source port-block Output Fields *(continued)*

Field Name	Field Description
Port block active timeout	Longest duration that a block remains active for port allocation.
Used/total port blocks	Current number of used ports and the total number of ports in this source pool.
Host IP	Address of the host IP.
External IP	Address of an external IP.
Port_Block Range	Port range of one PBA port block entry from the lowest to the highest port number that can be allowed to allocate ports for this block.
Ports_Used/Ports_Total	Current number of used ports and total number ports in this source pool.
Block_State/Left_Time(s)	PBA port block entry state for NAT port allocation, including Active, Inactive, Query, and the time left for a port block that is in the Active state or Query state. • Active—When an internal subscriber initiates a NAT request, a port block is allocated from the pool, and the status is set to Active. When there is a subsequent request from the same subscriber, a port is allocated from the existing Active block. • Inactive—When there is a request from an internal subscriber who has previously had a port allocated from this port block, but the time on the Active port block has expired or the ports are used up, the port block status changes from Active to Inactive. • InactiveB—When a chassis cluster is in active/passive mode, and a port block is created on the active node, the status for the synced port block on the backup node is InactiveB. • Query—When no ports are used in an Active port block, the status changes from Active to Query.

Sample Output

show security nat source port-block

```
user@host> show security nat source port-block
Pool name: p1
Port-overloading-factor:     1     Port block size:  128
Max port blocks per host:    4     Port block active timeout:    0
Used/total port blocks: 1/118944
Host_IP   External_IP    Port_Block   Ports_Used/    Block_State/
                          Range       Ports_Total    Left_Time(s)
203.0.113.1   198.51.100.20    51328-51455       2/128*1   Active/-
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source rule

Syntax

show security nat source rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 11.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Source port**, **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the **application** output field added in Junos OS Release 12.1X47-D10.

Description

Display information about the specified source Network Address Translation (NAT) rule.

Options

rule-name—Name of the rule.

all—Display information about all the source NAT rules.

logical-system (*logical-system-name* | all)—Display information about the source NAT rules for the specified logical system or for all logical systems source NAT rules.

root-logical-system—Display information about the source NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Source NAT) on page 5093
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source rule r2 on page 5159
show security nat source rule all on page 5159

Output Fields

Table 367 lists the output fields for the **show security nat source rule** command. Output fields are listed in the approximate order in which they appear

Table 460: show security nat source rule Output Fields

Field Name	Field Description
Source NAT rule	Name of the source NAT rule.
Total rules	Number of source NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.

Copyright © 2016, Juniper Networks, Inc.

Table 460: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the source NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the source NAT rule.
From zone	Name of the zone from which the packets flow.
To zone	Name of the zone to which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
To routing instance	Name of the routing instance to which the packets flow.
From interface	Name of the interface from which the packets flow.
To interface	Name of the interface to which the packets flow.
Source addresses	Name of the source addresses that match the rule.
Source port	Source port numbers that match the rule.
Destination address	Name of the destination addresses that match the rule.
Destination ports	Destination port numbers that match the rule.
Application	Indicates whether the application option is configured.
Action	The action taken in regard to a packet that matches the rule's tuples. Actions include the following: • **off**—Do not perform source NAT. • **source NAT pool**—Use user-defined source NAT pool to perform source NAT • **interface**—Use egress interface's IP address to perform source NAT.
Persistent NAT type	Persistent NAT type.
Persistent NAT mapping type	Persistent NAT mapping type.
Inactivity timeout	Inactivity timeout for persistent NAT binding.
Max session number	Maximum number of sessions.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.

Copyright © 2016, Juniper Networks, Inc.

Table 460: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat source rule r2

```
user@host> show security nat source rule r2

 source NAT rule: r2        Rule-set: src-nat
  Description               : The source rule r2 is for the sales team
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : zone1
  To zone                   : zone9
  Match
   Source addresses         : add1
                              add2
   Destination addresses    : add9
                              add10
   Destination port         : 1002            - 1002
  Action                       : off
    Persistent NAT type        : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout         : 0
    Max session number         : 0
  Translation hits             : 4719
   Successful sessions      : 2000
   Failed sessions          : 2719
 Number of sessions         : 5
```

Sample Output

show security nat source rule all

```
user@host> show security nat source rule all
Logical system: root
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 3/0

source NAT rule: r2                  Rule-set: rs2
  Rule-Id                  : 2
  Rule position            : 1
  From zone                : trust
  To zone                  : untrust
  Match
    Source addresses       : 192.0.2.0 - 192.0.2.255
    Destination addresses  : 203.0.113.0 - 203.0.113.255
                             198.51.100.0 - 198.51.100.2555
    Application            : configured
  Action                      : off
    Persistent NAT type       : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout        : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
  Max session number          : 0
Translation hits              : 0
  Successful sessions         : 0
  Failed sessions             : 0
Number of sessions            : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source rule-application

Syntax

show security nat source rule-application
rule-name
all
logical-system *logical-system-name*
root-logical-system

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Display information about the specified source Network Address Translation (NAT) rule application.

Options

rule-name—Display information about the specified source NAT rule application.

all—Display information about all the source NAT rule applications.

logical-system *logical-system-name* —Display information about the source NAT rule applications for the specified logical system.

root-logical-system—Display information about the source NAT rule applications for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- *Logical Systems Feature Guide for Security Devices*

List of Sample Output

show security nat source rule-application for port application on page 5162
show security nat source rule-application for ICMP application on page 5162

Output Fields

Table 461 lists the output fields for the **show security nat source rule-application** command. Output fields are listed in the approximate order in which they appear.

Table 461: show security nat source rule-application Output Fields

Field Name	Field Description
Destination NAT rule	Name of the source NAT rule.
Rule-set	Rule set identification number.
Rule-Id	Rule identification number.
Application	Name of the application or application set.
IP protocol	IP protocol identifier.
Source port range	Source port range identifier.

Copyright © 2016, Juniper Networks, Inc.

Table 461: show security nat source rule-application Output Fields *(continued)*

Field Name	Field Description
Destination port range	Destination port identifier.
ICMP information	• type—ICMP message type. • code—Code corresponding to the ICMP message type.

Sample Output

show security nat source rule-application for port application

```
user@host>show security nat source rule-application all

source NAT rule: r2                   Rule-set: rs2
Rule-Id                    : 2
  Application: app1
    IP protocol: 3
    Source port range: [90-90]
    Destination port range: [0-0]
    IP protocol: 4
    Source port range: [100-100]
    Destination port range: [200-200]
  Application: app2
    IP protocol: 7
    Source port range: [400-500]
    Destination port range: [80-80]
```

Sample Output

show security nat source rule-application for ICMP application

```
user@host>show security nat source rule-application all

source NAT rule: r1                   Rule-set: rs1
  Rule-Id                    : 1
    Application: junos-icmp-all
      IP protocol: icmp
      ICMP Information: type=255, code=0
    Application: icmp1
      IP protocol: icmp
      ICMP Information: type=1, code=1
    Application: junos-icmp6-all
      IP protocol: 58
      ICMP Information: type=255, code=0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat source summary

Syntax

show security nat source summary
<logical-system (*logical-system-name* | all)>
<root-logical-system>

Release Information

Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.

Description

Display a summary of Network Address Translation (NAT) source information.

Options

none—Display summary source NAT information.

logical-system (*logical-system-name* | all)—Display summary information about the source NAT for the specified logical system or for all logical systems.

root-logical-system—Display summary information about the source NAT for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- pool (Security Source NAT) on page 5076
- rule (Security Source NAT) on page 5093
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat source summary on page 5164

Output Fields

Table 368 lists the output fields for the **show security nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 462: show security nat source summary Output Fields

Field Name	Field Description
Total source nat pool number	Number of source NAT pools.
Pool name	Name of the source address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).
Total Address	Number of IP addresses that are in use.
Total source nat rule number	Number of source NAT rules.

Copyright © 2016, Juniper Networks, Inc.

Table 462: show security nat source summary Output Fields *(continued)*

Field Name	Field Description
Total port number usage for port translation pool	Number of ports assigned to the pool.
Maximum port number for port translation pool	Maximum number of NAT or PAT transactions done at any given time.

Sample Output

show security nat source summary

```
user@host> show security nat source summary logical-system all


Logical system: root-logical-system
Total port number usage for port translation pool: 67108864
Maximum port number for port translation pool: 134217728

Logical system: lsys1
Total port number usage for port translation pool: 193536
Maximum port number for port translation pool: 134217728
Total pools: 2

Logical system: root-logical-system
Pool              Address                   Routing   PAT   Total
Name              Range                     Instance        Address
pool1             10.1.1.0-10.1.4.255-        default   yes   2048
                  10.1.5.0-10.1.8.255


Logical system: lsys1
Pool              Address                    Routing   PAT   Total
Name              Range                      Instance        Address
pool2             203.0.113.1-203.0.113.3     default   yes   3


Total rules: 1

Logical system: root-logical-system
Rule name         Rule set      From         To             Action
rule 1            ruleset1      ge-2/2/2.0   ge-2/2/3.0     pool1
rule 1                          ge-2/2/4.0   ge-2/2/5.0
```


Copyright © 2016, Juniper Networks, Inc.

show security nat static rule

Syntax

show security nat static rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information

Command introduced in Junos OS Release 9.3. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Successful sessions**, **Failed sessions**, **Number of sessions**, **Source addresses,** and **Source ports** output fields added in Junos OS Release 12.1X45-D10. The **Destination NPTv6 addr** and **Destination NPTv6 Netmask** output fields added in Junos OS Release 12.3X48-D25.

Description

Display information about the specified static Network Address Translation (NAT) rule.

Options

rule-name—Name of the rule.

all—Display information about all the static NAT rules.

logical-system (*logical-system-name* | all)—Display information about the static NAT rules for the specified logical system or for all logical systems.

root-logical-system—Display information about the static NAT rules for the master (root) logical system.

Required Privilege Level

view

Related Documentation

- rule (Security Static NAT) on page 5095
- Security Configuration Statement Hierarchy on page 603

List of Sample Output

show security nat static rule on page 5166
show security nat static rule (IPv6) on page 5167
show security nat static rule all on page 5167

Output Fields

Table 369 lists the output fields for the **show security nat static rule** command. Output fields are listed in the approximate order in which they appear.

Table 463: show security nat static rule Output Fields

Field Name	Field Description
Static NAT rule	Name of the static NAT rule.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Rule-set	Name of the rule set. Currently, you can configure 8 rules within the same rule set.

Copyright © 2016, Juniper Networks, Inc.

Table 463: show security nat static rule Output Fields *(continued)*

Field Name	Field Description
Description	Description of the static NAT rule.
Rule-Id	Rule identification number.
Rule position	Position of the rule that indicates the order in which it applies to traffic.
From interface	Name of the interface from which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
From zone	Name of the zone from which the packets flow.
Destination addresses	Name of the destination addresses that match the rule.
Destination NPTv6 addr	Destination address that matches the rule.
Source addresses	Name of the source addresses that match the rule.
Host addresses	Name of the host addresses that match the rule.
Netmask	Subnet IP address.
Destination NPTv6 Netmask	Subnet IPv6 address.
Host routing-instance	Name of the host routing instance.
Destination port	Destination port numbers that match the rule. The default value is any.
Source port	Source port numbers that match the rule.
Total static-nat rules	Number of static NAT rules.
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat static rule

```
user@host> show security nat static rule sta-r2

 Static NAT rule: sta-r2                  Rule-set: sta-nat
```


Copyright © 2016, Juniper Networks, Inc.

```
  Description               : The static rule sta-r2 is for the sales team
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : zone9
  Destination addresses     : add3
  Host addresses            : add4
  Netmask                   : 24
  Host routing-instance     : N/A
  Translation hits          : 2
   Successful sessions     : 2
   Failed sessions         : 0
 Number of sessions        : 2
```

Sample Output

show security nat static rule (IPv6)

```
user@host> show security nat static rule r1

Static NAT rule: r1                  Rule-set: rs1
   Rule-Id                   : 1
   Rule position             : 1
   From zone                 : trust
   Destination NPTv6 addr    : 2001:db8::
   Destination NPTv6 Netmask : 48
   Host addresses            : 2001:db8::3000:
   Netmask                   : 48
   Host routing-instance     : N/A
   Translation hits          : 0
      Successful sessions     : 0
      Failed sessions         : 0
   Number of sessions        : 0
```

Sample Output

show security nat static rule all

```
user@host> show security nat static rule all

Static NAT rule: r1                  Rule-set: rs1
  Rule-Id                   : 1
  Rule position             : 1
  From zone                 : trust
  Source addresses          : 192.0.2.0 -192.0.2.3
: addr1
  Source ports              : 200 - 300
  Destination addresses     : 198.51.100.0
  Host addresses            : 203.0.113.0
  Netmask                   : 24
  Host routing-instance     : N/A
  Translation hits          : 4
   Successful sessions     : 4
   Failed sessions         : 0
 Number of sessions        : 4
Static NAT rule: r2                  Rule-set: rs1
  Rule-Id                   : 2
  Rule position             : 2
  From zone                 : trust
  Source addresses          : 192.0.2.0 -192.0.2.255  Destination addresses
  : 203.0.113.1
  Destination ports         : 100 - 200
```

Copyright © 2016, Juniper Networks, Inc.

```
  Host addresses             : 192.0.2.1
  Host ports                 : 300 - 400
  Netmask                    : 32
  Host routing-instance      : N/A
  Translation hits           : 4
   Successful sessions      : 4
   Failed sessions          : 0
Number of sessions          : 4
```


Copyright © 2016, Juniper Networks, Inc.

Authentication and Integrated User Firewalls Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 66

Overview

- Introduction to User Authentication on page 5173

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 211

Introduction to User Authentication

- Understanding User Authentication for Security Devices on page 5173
- Understanding the Three-Tiered User Firewall Features on page 5173

Understanding User Authentication for Security Devices

Firewall user authentication lets you define firewall users and create policies that require the users to authenticate themselves through one of two authentication schemes: pass-through authentication or web authentication.

User role firewall policies can be integrated with firewall authentication both to authenticate users and to retrieve username and role information. The information is mapped to the IP address of the traffic, stored in the firewall authentication table, and used for user role firewall policy enforcement.

Infranet authentication occurs when an SRX Series device acts as an Infranet Enforcer for an IC Series device. You deploy the Infranet Enforcer in front of the servers and resources that you want to protect. Authentication occurs on the IC Series device and provides policies to the Enforcer to determine whether or not to allow an endpoint access to protected resources.

Understanding the Three-Tiered User Firewall Features

Juniper Networks offers three tiers of user firewall. The three features have different characteristics that are appropriate in different environments. Figure 222 illustrates the relative security level of the three tiers. Table 464 compares them to help you decide which best suits your implementation.

Copyright © 2016, Juniper Networks, Inc.

Figure 222: Three-Tiered User Firewall Features

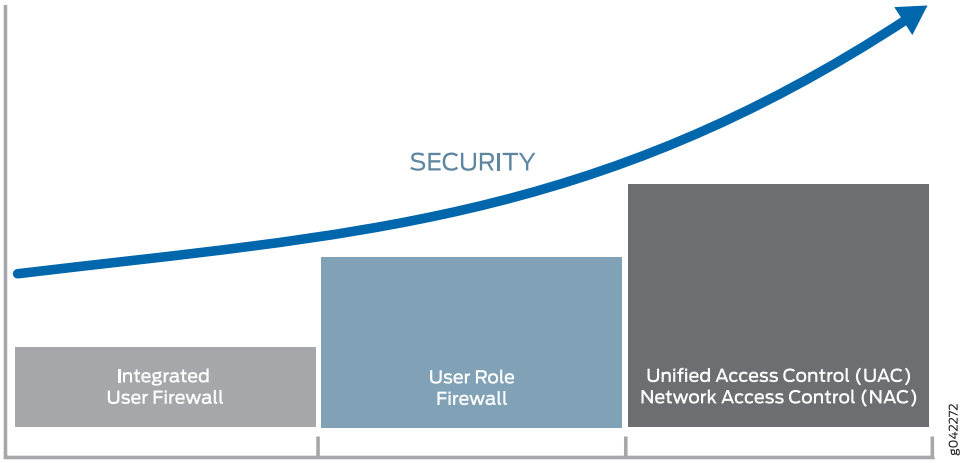

Table 464 describes the basic differences among the three features.

Table 464: Comparison of User Firewall Features

	Integrated User Firewall	User Role Firewall	Unified Access Control (UAC) Network Access Control (NAC)
Authentication	Passive authentication–Does not interact with client directly; polls the Active Directory for login information.	Active authentication–Queries the client.	End-to-end–Authenticates the user down to the access level where user connects, whether wired or wireless.
Extent of Authentication	Best effort.	Deterministic–User is identified.	Deterministic–User is identified.
Where Enforced	Enforced at firewall.	Enforced at firewall.	Enforced at access (switch or WiFi) and firewall.
Devices Needed	SRX Series	SRX Series and MAG Series	SRX Series and MAG Series
Ideal Environments	• Needs visibility into who is accessing the SRX Series • Small-to-medium business • Low-scale deployment	• Security-conscious environments • Scales up to 50,000 users	• Large-scale deployment • Interface for Metadata Access Points (IF-MAP) federation

- You can upgrade to a higher tier if you choose. From integrated user firewall, simply add the MAG Series to get user role firewall. From there, add licenses to get full UAC NAC.
- The three offerings provide maximum flexibility; they are supported on all SRX Series hardware platforms.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 67

Configuring Firewall User Authentication

- Understanding Firewall Authentication on page 5179
- Configuring Encrypted Files Using SSL Proxy on page 5183
- Configuring Pass-Through Authentication on page 5205
- Configuring Web Authentication on page 5221
- Configuring External Authentication Servers on page 5235
- Configuring Client Groups on page 5247
- Customizing the Firewall Authentication Banner on page 5251

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 212

Understanding Firewall Authentication

- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 5180

Firewall User Authentication Overview

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

Junos OS also supports the administrator and Point-to-Point Protocol (PPP) user types.

NOTE: **HTTPS-based authentication is introduced on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.**

After you define firewall users, you can create a policy that requires the users to authenticate themselves through one of three authentication schemes:

- **Pass-through authentication**—A host or a user from one zone tries to access resources on another zone. You must use an FTP client, a Telnet client, an HTTP client, or an HTTPS client to access the IP address of the protected resource and to get authenticated by the firewall. The device uses FTP, Telnet, HTTP, or HTTPS to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication. When the device is using an HTTPS server, and after the authentication is done, the subsequent traffic from the user is always terminated whether the authentication is successful or not.
- **Pass-through with web-redirect authentication**—This authentication scheme is for HTTP and HTTPS. when firewall authentication is performed using the pass-through authentication scheme for HTTP and HTTPs clients, the web-redirect feature optionally, redirects HTTP or HTTPS requests to the device's internal webserver by sending a redirect HTTP or HTTPS response to the client system to reconnect to the webserver for user authentication. The interface on which the client's request arrived is the interface to which the request is redirected.

Copyright © 2016, Juniper Networks, Inc.

Using this feature allows for a richer user login experience. For example. instead of a popup prompt asking for username and password, users can get the login page in a browser. Enabling **web-redirect** has the same effect as users typing the Web authentication IP address in a client browser. Using **web-redirect** provides a more seamless authentication experience because users do not need to know the Web authentication IP address but only the IP address of the resource they are trying to access. After the user has been authenticated this way, traffic from user's IP address is authenticated to go through the **web-redirect** method.

When you web-redirect, you http or https the destination url, the authentication page prompts you after successful authentication.

A message is displayed to inform you about the successful authentication. After successful authentication, the browser launches your original destination URL without your needing to retype the URL.

The following message is displayed:

```
Redirecting to the original url, please wait
```

- Web authentication—Users try to connect, using HTTP or HTTPS, to an IP address on the device that is enabled for Web authentication; in this scenario, you do not use HTTP or HTTPS to get to the IP address of the protected resource. You are prompted for the username and password that are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

Related Documentation

- Understanding Pass-Through Authentication on page 5205
- Understanding Web Authentication on page 5221
- Understanding External Authentication Servers on page 5235
- Configuring SSL Proxy on page 540

Obtaining Username and Role Information Through Firewall Authentication

User role firewall policies can be integrated with firewall authentication both to authenticate users and to retrieve username and role information. The information is mapped to the IP address of the traffic, stored in the firewall authentication table, and used for user role firewall policy enforcement.

The following CLI statements configure firewall authentication for user role firewall enforcement.

1. If not already established, define the access profile to be used for firewall authentication. You can skip this step if an existing access profile provides the client data needed for your implementation.

   The access profile is configured in the **[edit access profile]** hierarchy as with other firewall authentication types. It defines clients as firewall users and the passwords that provide them access. Use the following command to define a profile and add client names and passwords for firewall authentication.

Copyright © 2016, Juniper Networks, Inc.

```
set access profile profile-name client client-name firewall-user password pwd
```

2. If HTTPS traffic is expected, define the access profile to be used for SSL termination services. You can skip this step if an existing SSL termination profile provides the services needed for your implementation.

   The SSL termination profile is configured in the **[edit services ssl]** hierarchy.

   ```
   set services ssl termination profile ssl-profile-name server-certificate certificate-type
   ```

3. Enable the firewall authentication table as an authentication source.

   ```
   set security user-identification authentication-source firewall-authentication priority
     priority
   ```

   The priority value determines the sequence in which authentication sources are checked. The default value is 150 for the firewall authentication table. (It is 100 for the local authentication table and 200 for the Unified Access Control (UAC) authentication table.) By default, the local authentication table is checked first, the firewall authentication table is next, and the UAC authentication table is third if it is enabled. You can change this sequence by changing the priority value of one or more of the tables.

4. Configure policies that permit traffic for user firewall authentication.

   ```
   edit security policies from-zone zone to-zone zone policy policy-name
   set match source-identity unauthenticated-user
   set then permit firewall-authentication user-firewall access-profile profile-name
     ssl-termination-profile profile-name
   ```

   When unauthenticated traffic is permitted for firewall authentication, the user is authenticated based on the access profile configured in this statement. The **ssl-termination-profile** option is needed only for HTTPS traffic.

   By specifying the authentication type **user-firewall**, the firewall authentication table is propagated with the IP address, the username, and any group names associated with the authenticated user. (Group names from firewall authentication are interpreted as roles by the user role firewall.) Any further traffic from this IP address will match the IP address in the firewall authentication table, and not require authentication. The associated username and roles are retrieved from the table for use as potential match criteria in subsequent security policies.

Related Documentation

- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 213

Configuring Encrypted Files Using SSL Proxy

- SSL Proxy Overview on page 5183
- Configuring SSL Proxy on page 5191
- Enabling Debugging and Tracing for SSL Proxy on page 5202

SSL Proxy Overview

Secure Sockets Layer (SSL) is an application-level protocol that provides encryption technology for the Internet. SSL, also called Transport Layer Security (TLS), ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private-public key exchange pairs for this level of security.

Server authentication guards against fraudulent transmissions by enabling a Web browser to validate the identity of a webserver. Confidentiality mechanisms ensure that communications are private. SSL enforces confidentiality by encrypting data to prevent unauthorized users from eavesdropping on electronic communications. Finally, message integrity ensures that the contents of a communication have not been tampered with.

SSL proxy is transparent; that is, it performs SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence. Existing features like SSL offload and SSL inspection require the servers to share their secret keys to be able to decrypt the SSL traffic. However, sharing server keys is sometimes not feasible or might not be available in certain circumstances, in which case the SSL traffic cannot be decrypted.

SSL proxy addresses this problem by ensuring that it has the keys to encrypt and decrypt the payload:

- For the server, SSL proxy acts as a client—Because SSL proxy generates the shared pre-master key, it determines the keys to encrypt and decrypt.
- For the client, SSL proxy acts as a server—SSL proxy first authenticates the original server and replaces the public key in the original server certificate with a key that is known to it. It then generates a new certificate by replacing the original issuer of the certificate with its own identity and signs this new certificate with its own public key (provided as a part of the proxy profile configuration). When the client accepts such

Copyright © 2016, Juniper Networks, Inc.

a certificate, it sends a shared pre-master key encrypted with the public key on the certificate. Because SSL proxy replaced the original key with its own key, it is able to receive the shared pre-master key. Decryption and encryption take place in each direction (client and server), and the keys are different for both encryption and decryption.

Figure 29 depicts how SSL inspection (on an existing SRX Series IDP module) is typically used to protect servers. SSL inspection requires access to the private keys used by the servers so that the SRX Series device can decrypt the encrypted traffic.

Figure 223: SSL Inspection on an Existing SRX Series IDP Module

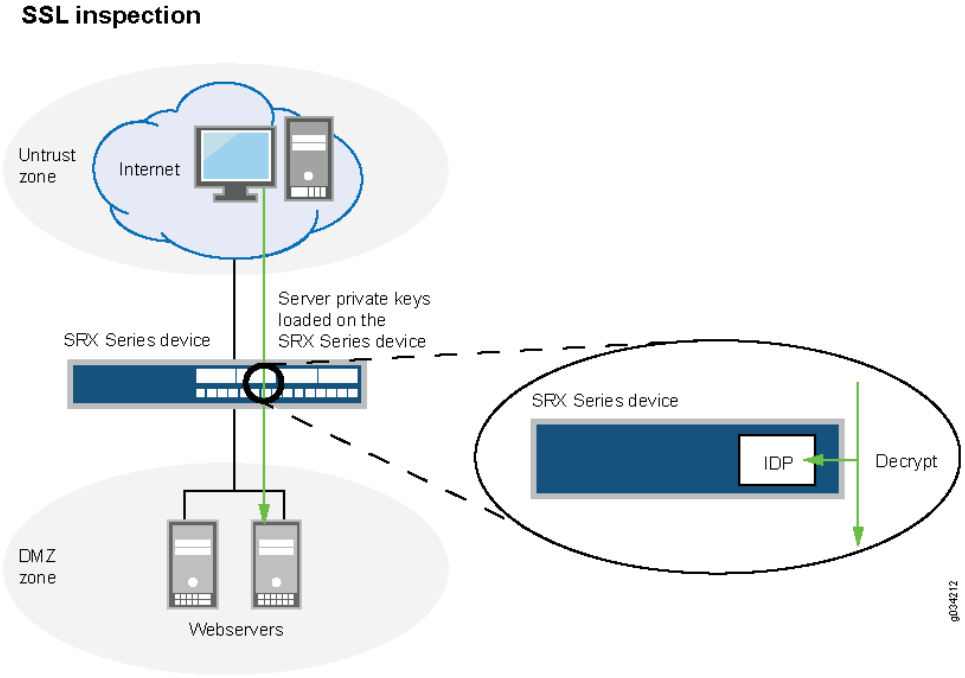

Figure 30 shows how SSL proxy works on an encrypted payload. When application firewall (AppFW), Intrusion Detection and Prevention (IDP), or application tracking (AppTrack) is configured, the SSL proxy acts as an SSL server by terminating the SSL session from the client and establishing a new SSL session to the server, the SRX Series device decrypts and then reencrypts all SSL proxy traffic. SSL proxy uses the following:

- SSL-T-SSL terminator on the client side.
- SSL-I-SSL initiator on the server side.
- Configured AppFW, IDP, or AppTrack services use the decrypted SSL sessions.

NOTE: If none of the services (AppFW, IDP, or AppTrack) are configured, then SSL proxy services are bypassed even if an SSL proxy profile is attached to a firewall policy.

Copyright © 2016, Juniper Networks, Inc.

Figure 224: SSL Proxy on an Encrypted Payload

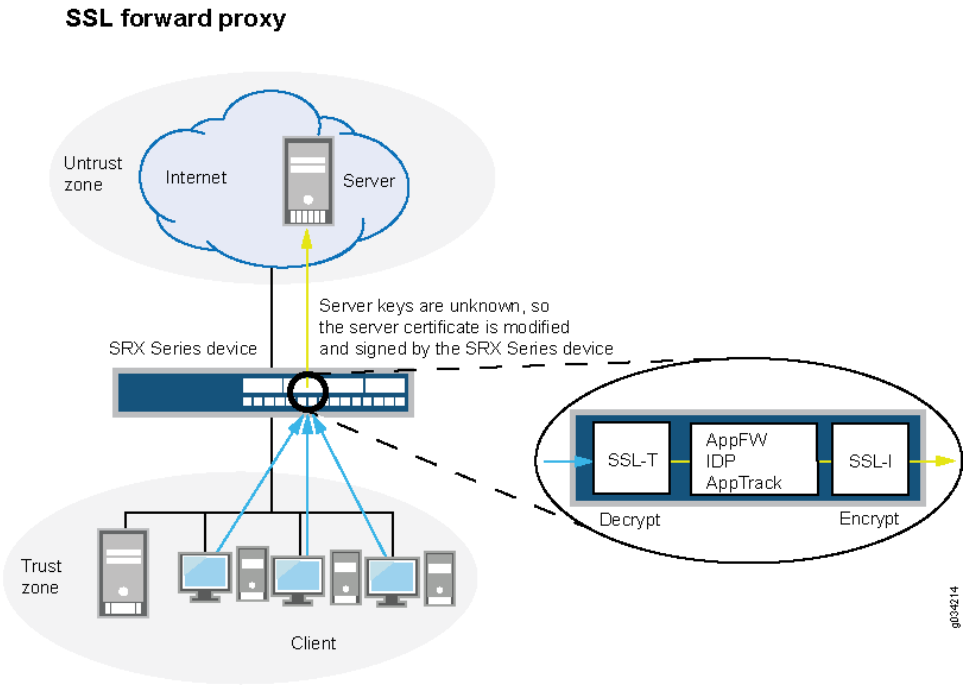

Supported Ciphers in Proxy Mode

An SSL cipher comprises encryption ciphers, authentication method, and compression. Table 34 displays a list of supported ciphers. NULL ciphers are excluded.

The following SSL protocols are supported on SRX Series devices:

- TLS version 1.0—Provides secure communication over networks by providing privacy and data integrity between communicating applications
- TLS version 1.1—This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.
- TLS version 1.2 — This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.

Table 465: Supported SSL Cipher List

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_RC4_128_MD5	RSA key exchange	128-bit RC4	Message Digest 5 (MD5) hash
RSA_WITH_RC4_128_SHA	RSA key exchange	128-bit RC4	Secure Hash Algorithm (SHA) hash
RSA_WITH_DES_CBC_SHA	RSA key exchange	DES CBC	SHA hash
RSA_WITH_3DES_EDE_CBC_SHA	RSA key exchange	3DES EDE/CBC	SHA hash

Copyright © 2016, Juniper Networks, Inc.

Table 465: Supported SSL Cipher List *(continued)*

SSL Cipher	Key Exchange Algorithm	Data Encryption	Message Integrity
RSA_WITH_AES_128_CBC_SHA	RSA key exchange	128-bit AES/CBC	SHA hash
RSA_WITH_AES_256_CBC_SHA	RSA key exchange	256-bit AES/CBC	SHA hash
RSA_EXPORT_WITH_RC4_40_MD5	RSA-export	40-bit RC4	MD5 hash
RSA_EXPORT_WITH_DES40_CBC_SHA	RSA-export	40-bit DES/CBC	SHA hash
RSA_EXPORT1024_WITH_DES_CBC_SHA	RSA 1024 bit export	DES/CBC	SHA hash
RSA_EXPORT1024_WITH_RC4_56_MD5	RSA 1024 bit export	56-bit RC4	MD5 hash
RSA_EXPORT1024_WITH_RC4_56_SHA	RSA 1024 bit export	56-bit RC4	SHA hash

NOTE: Cipher suites that have “export” in the title are intended for use outside of the United States and might have encryption algorithms with limited key sizes.

Export ciphers are not enabled by default. You need to either configure the export ciphers to enable or install a domestic package.

NOTE: Supported SSL ciphers for HTTPS firewall authentication are RSA_WITH_3DES_EDE_CBC_SHA , RSA_WITH_AES_128_CBC_SHA, and RSA_WITH_AES_256_CBC_SHA.

Server Authentication

Implicit trust between the client and the device (because the client accepts the certificate generated by the device) is an important aspect of SSL proxy. It is extremely important that server authentication is not compromised; however, in reality, self-signed certificates and certificates with anomalies are in abundance. Anomalies can include expired certificates, instances of common name not matching a domain name, and so forth. Server authentication is governed by setting the **ignore-server-auth-failure** option in the SSL proxy.

- By default, the **ignore-server-auth-failure** option is not defined as an action in the SSL proxy profile, and the following occurs:
  - If authentication succeeds, a new certificate is generated by replacing the keys and changing the issuer name to the issuer name that is configured in the root CA certificate in the proxy profile.
  - If authentication fails, the connection is dropped.

Copyright © 2016, Juniper Networks, Inc.

- If the **ignore-server-auth-failure** option is defined as an action in the SSL proxy profile, the following occurs:
  - If the certificate is self-signed, a new certificate is generated by replacing the keys only. The issuer name is not changed. This ensures that the client browser displays a warning that the certificate is not valid.
  - If the certificate has expired or if the common name does not match the domain name, a new certificate is generated by replacing the keys and changing the issuer name to SSL-PROXY: DUMMY_CERT:GENERATED DUE TO SRVR AUTH FAILURE. This ensures that the client browser displays a warning that the certificate is not valid.

Trusted CA List

SSL proxy ensures secure transmission of data between a client and a server. Before establishing a secure connection, SSL proxy checks CA certificates to verify signatures on server certificates. For this reason, a reasonable list of trusted CA certificates is required to effectively authenticate servers.

Junos OS provides the following options for trusted CA certificates:

- Loading the default trusted CA list—Junos OS provides a default list of certificates that contains well-known trusted CA certificates similar to the default certificates used by most common browsers. Without these default certificates, browsers would not be able to validate the identity of most websites and would mark them as untrusted sites.

  The Junos OS package contains the default CA certificates as a PEM file (for example, trusted_CA.pem). After you download the package and reboot your device, you can easily load the default certificates on your system using a CLI command.

  We recommend you load the default trusted CA list if you want to trust the same CA certificates as common browsers and avoid importing CA certificates manually.

- Importing the trusted CA list manually—You can import your own trusted CA certificates using the Public Key Infrastructure (PKI). The PKI helps verify and authenticate the validity of the trusted CA certificates. You create CA profile groups that include trusted CA certificates, then import the group on your device for server authentication.

- Ignoring server authentication—You can use the **ignore-server-auth-failure** option to ignore server authentication completely. In this case, SSL proxy ignores errors encountered during the server certificate verification process (such as CA signature verification failure, self-signed certificates, and certificate expiry).

  We do not recommend this option for authentication, because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause for dropped SSL sessions. See "Enabling Debugging and Tracing for SSL Proxy" on page 552.

Root CA

In a public key infrastructure (PKI) hierarchy, the root CA is at the top of the trust path. The root CA identifies the server certificate as a trusted certificate.

Copyright © 2016, Juniper Networks, Inc.

Client Authentication

Currently, client authentication is not supported in SSL proxy. If a server requests client authentication, a warning is issued that a certificate is not available. The warning lets the server determine whether to continue or to exit.

Whitelists

Because SSL encryption and decryption are complicated and expensive procedures, network administrators can selectively bypass SSL proxy processing for some sessions. Such sessions mostly include connections and transactions with trusted servers or domains with which network administrators are very familiar. There are also legal requirements to exempt financial and banking sites. Such exemptions are achieved by configuring the IP addresses or domain names of the servers under whitelists.

Dynamic Resolution of Domain Names

The IP addresses associated with domain names are dynamic and can change at any time. Whenever a domain IP address changes, it is propagated to the SSL proxy configuration (similar to what is done in the firewall policy configuration).

Session Resumption

An SSL session refers to the set of parameters and encryption keys created by performing a full handshake. A connection is the conversation or active data transfer that occurs within the session. The computational overhead of a complete SSL handshake and generation of master keys is considerable. In short-lived sessions, the time taken for the SSL handshake can be more than the time for data transfer.

To improve throughput and still maintain an appropriate level of security, SSL session resumption provides a session caching mechanism so that session information, such as the pre-master secret key and agreed-upon ciphers, can be cached for both the client and server. The cached information is identified by a session ID. In subsequent connections both parties agree to use the session ID to retrieve the information rather than create a new pre-master secret key. Session resumption shortens the handshake process and accelerates SSL transactions.

Session Renegotiation

After a session is created and SSL tunnel transport has been established, a change in SSL parameters requires renegotiation. SSL proxy supports both secure (RFC 5746) and nonsecure (TLS v1.0, TLS v1.1, and TLS v1.2) renegotiation. When session resumption is enabled, session renegotiation is useful in the following situations:

- Cipher keys need to be refreshed after a prolonged SSL session.
- Stronger ciphers need to be applied for a more secure connection.

A change in an SSL proxy profile that modifies a certificate, cipher strength, or trusted CA list flushes cache entries when the modified policy is committed. When a session is resumed, the SSL parameters associated with its session ID are retrieved from the cache. If the SSL proxy profile is not altered, cache entries corresponding to that profile are not

Copyright © 2016, Juniper Networks, Inc.

flushed and the session continues. If the cache has been flushed, however, a full handshake must be performed to establish the new SSL parameters. (There is no impact to non-SSL sessions.)

SSL Proxy Logs

When logging is enabled in an SSL proxy profile, SSL proxy can generate the messages shown in Table 35.

Table 466: SSL Proxy Logs

Syslog Type	Description
SSL_PROXY_SSL_SESSION_DROP	Logs generated when a session is dropped by SSL proxy.
SSL_PROXY_SSL_SESSION_ALLOW	Logs generated when a session is processed by SSL proxy even after encountering some minor errors.
SSL_PROXY_SESSION_IGNORE	Logs generated if non-SSL sessions are initially mistaken as SSL sessions.
SSL_PROXY_SESSION_WHITELIST	Logs generated when a session is whitelisted.
SSL_PROXY_ERROR	Logs used for reporting errors.
SSL_PROXY_WARNING	Logs used for reporting warnings.
SSL_PROXY_INFO	Logs used for reporting general information.

All logs contain similar information as shown in the following example (actual order of appearance):

```
logical-system-name, session-id, source-ip-address, source-port,
destination-ip-address,destination-port,
nat-source-ip-address, nat-source-port, nat-destination-ip-address,
nat-destination-port, proxy profile name, source-zone-name,
source-interface-name, destination-zone-name,destination-interface-name, message
```

The **message** field contains the reason for the log generation. One of three prefixes shown in Table 36 identifies the source of the message. Other fields are descriptively labeled.

Table 467: SSL Proxy Log Prefixes

Prefix	Description
system	Logs generated due to errors related to the device or an action taken as part of the SSL proxy profile. Most logs fall into this category.
openssl error	Logs generated during the handshaking process if an error is detected by the openssl library.
certificate error	Logs generated during the handshaking process if an error is detected in the certificate (x509 related errors).

Copyright © 2016, Juniper Networks, Inc.

Sample logs:

```
Jun  1 05:11:13 4.0.0.254 junos-ssl-proxy: SSL_PROXY_SSL_SESSION_DROP: lsys:root
  23 <4.0.0.1/35090->5.0.0.1/443>  NAT:< 4.0.0.1/35090->5.0.0.1/443 >
ssl-inspect-profile  <untrust:ge-0/0/0.0->trust:ge-0/0/1.0> message:certificate
error: self signed certificate
```

NOTE: These logs capture sessions that are dropped by SSL proxy, not sessions that are marked by other modules that also use SSL proxy services.

For SSL_PROXY_SESSION_WHITELIST messages, an additional **host** field is included after the **session-id** and contains the IP address of the server or domain that has been whitelisted.

```
Jun  1 05:25:36 4.0.0.254 junos-ssl-proxy: SSL_PROXY_SESSION_WHITELIST: lsys:root
 24 host:5.0.0.1 <4.0.0.1/35090->5.0.0.1/443>  NAT:< 4.0.0.1/35090->5.0.0.1/443
> ssl-inspect-profile  <untrust:ge-0/0/0.0->trust:ge-0/0/1.0> message:system:
session whitelisted
```

Leveraging Dynamic Application Identification

SSL proxy uses application identification services to dynamically detect if a particular session is SSL encrypted. SSL proxies are allowed only if a session is SSL encrypted. The following rules apply for a session:

- Session is marked **Encrypted=Yes** in the application system cache. If the session is marked **Encrypted=Yes**, it indicates that the final match from application identification for that session is SSL encrypted, and SSL proxy transitions to a state where proxy functionality can be initiated.
- Session is marked **Encrypted=No** in the application system cache. If a non-SSL entry is found in the application system cache, it indicates that the final match from application identification for that session is non-SSL and SSL proxy ignores the session.
- An entry is not found in the application system cache. This can happen on the first session, or when the application system cache has been cleaned or has expired. In such a scenario, SSL proxy cannot wait for the final match (requires traffic in both directions). In SSL proxy, traffic in reverse direction happens only if SSL proxy has initiated an SSL handshake. Initially, for such a scenario SSL proxy tries to leverage prematch or aggressive match results from application identification , and if the results indicate SSL, SSL proxy will go ahead with the handshake.
- Application identification fails due to resource constraints and other errors. Whenever the result from application identification is not available, SSL proxy will assume static port binding and will try to initiate SSL handshake on the session. This will succeed for actual SSL sessions, but it will result in dropped sessions for non SSL sessions.

Logical Systems Support

It is possible to enable SSL proxy on firewall policies that are configured using logical systems; however, note the following limitations:

Copyright © 2016, Juniper Networks, Inc.

- The “services” category is currently not supported in logical systems configuration. Because SSL proxy is under “services,” you cannot configure SSL proxy profiles on a per-logical-system basis.
- Because proxy profiles configured at a global level (within “services ssl proxy”) are visible across logical system configurations, it is possible to configure proxy profiles at a global level and then attach them to the firewall policies of one or more logical systems.

Limitations

The following limitations currently apply to SSL proxy:

- Only TLSv1.0, TLSv1.1, and TLSv1.2 are supported.
- SSL v3.0 is not supported.
- Only X.509v3 certificate is supported
- Client authentication of SSL handshake is not supported.

NOTE: On all high-end SRX Series devices, for a particular session, the SSL proxy is only enabled if a relevant feature related to SSL traffic is also enabled. Features that are related to SSL traffic are IDP, application identification, application firewall, and application tracking. If none of the above listed features are active on a session, the SSL proxy bypasses the session and logs are not generated in this scenario

NOTE: On all SRX Series devices, the current SSL proxy implementation has the following connectivity limitations:

- The SSLv2 protocol is not supported. SSL sessions using SSLv2 are dropped
- SSL sessions where client certificate authentication is mandatory are dropped
- SSL sessions where renegotiation is requested are dropped

Related Documentation

- Understanding Address Books on page 1069
- Understanding Global Address Books on page 1071
- Understanding Self-Signed Certificates on page 6433
- Configuring SSL Proxy on page 540

Configuring SSL Proxy

SSL proxy works transparently between the client and the server. All requests from a client first go to the proxy server; the proxy server evaluates the request, and if the request

Copyright © 2016, Juniper Networks, Inc.

is valid, forwards the request to the outbound side. Similarly, inbound requests are also evaluated by the proxy server. Both client and server interpret that they are communicating with each other; however, it is the SSL proxy that functions between the two.

SSL proxies provide encryption and decryption by residing between the server and the client. Because SSL proxies are hidden from both the server and the client, secret keys are shared between the two to decrypt the SSL traffic. Proxies are known as *forward proxies* because proxy servers are used to hide any detailed information from the servers.

Integrity, confidentiality, and authenticity of traffic are validated through PKI, which includes digital certificates issued by the CA, certificate validity and expiration dates, details about the certificate owner and issuer, and security policies.

- SSL Proxy Configuration Overview on page 5192
- Configuring a Root CA Certificate on page 5193
- Configuring a CA Profile Group on page 5195
- Configuring a Trusted CA Profile on page 5196
- Importing a Root CA Certificate into a Browser on page 5197
- Applying an SSL Proxy Profile to a Security Policy on page 5198
- Creating a Whitelist of Exempted Destinations on page 5199
- Configuring SSL Proxy Logging on page 5200
- Configuring Ciphers on page 5200
- Exporting Certificates to a Specified Location on page 5201
- Ignoring Server Authentication on page 5201

SSL Proxy Configuration Overview

Figure 31 displays an overview of how SSL proxy is configured. It includes some required steps, such as configuring the root CA certificate, loading a CA profile group, and applying an SSL proxy profile to a security policy, and some optional steps, such as creating whitelists and SSL proxy logging.

Copyright © 2016, Juniper Networks, Inc.

Figure 225: SSL Proxy Configuration Overview

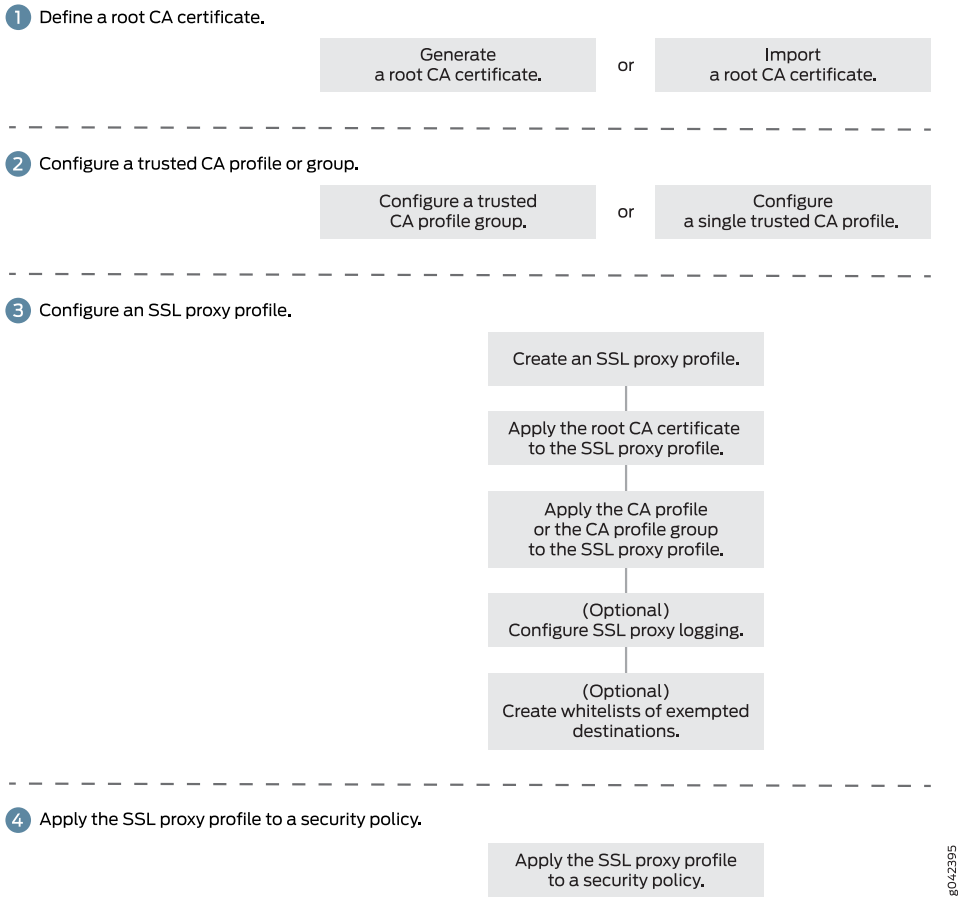

Configuring a Root CA Certificate

A CA can issue multiple certificates in the form of a tree structure. A root certificate is the topmost certificate of the tree, the private key of which is used to *sign* other certificates. All certificates immediately below the root certificate inherit the signature or trustworthiness of the root certificate. This is somewhat like the *notarizing* of an identity.

You can configure a root CA certificate by first obtaining a root CA certificate (by either generating a self-signed one or importing one) and then applying it to an SSL proxy profile. There are two ways you can obtain a root CA certificate—by using the Junos OS CLI on an SRX Series device or by using OpenSSL on a UNIX device.

To generate a root CA certificate using the Junos OS CLI, follow these steps on an SRX Series device:

1. From operational mode, generate a PKI public/private key pair for a local digital certificate.

   user@host>**request security pki generate-key-pair certificate-id** ***certificate-id*** **size** ***size*** **type** ***type***

Copyright © 2016, Juniper Networks, Inc.

2. From operational mode, define a self-signed certificate. Specify certificate details such as the certificate identifier (generated in the previous step), a fully qualified domain name (FQDN) for the certificate, and an e-mail address of the entity owning the certificate. You can also specify other information such as the common name and the organization involved. By configuring the **add-ca-constraint** option, you make sure that the certificate can be used for signing other certificates.

   user@host>**request security pki local-certificate generate-self-signed certificate-id** ***certificate-id*** **domain-name** ***domain-name*** **subject** ***subject*** **email** ***email-id*** **add-ca-constraint**

3. From configuration mode, apply the loaded certificate as root-ca in the SSL proxy profile.

   [edit]
   user@host# **set services ssl proxy profile** ***profile-name*** **root-ca** ***certificate-id***

4. Import the root CA as a trusted CA into client browsers. This is required for the client browsers to trust the certificates signed by the SRX Series device. See "Importing a Root CA Certificate into a Browser" on page 545.

To generate a root CA certificate using OpenSSL, follow these steps on a UNIX device:

1. Create folders **keys** and **certs**.

   mkdir /etc/pki/tls/keys
   mkdir /etc/pki/tls/certs

2. Change to the **openssl** directory.

   cd /etc/pki/tls

3. Create a CA certificate key. The following command creates an RSA key using the 3DES encryption named **ca.key** that is 2048 in length. You also need to enter a password that is used to encrypt the private key. This is critical to security if the key is lost because it will still be encrypted.

   % openssl genrsa -des3 -out keys/ssl-proxy-ca.key 2048

4. Create a CA certificate based on the CA private key (created in the previous step). The expiration date for this certificate is 3 years or 1095 days. However, you can set it to a different value. When creating the certificate, you need to enter the password and the certificate information that includes distinguished name (DN), country name, and so forth.

   % openssl req -new -x509 -days 1095 –key keys/ssl-proxy-ca.key -out certs/ssl-inspect-ca.cer

5. Import the CA private and public keys into the SRX Series device. Copy the **ca.key** and **ca.cer** keys to the **/var/tmp** directory on the SRX Series device. You can copy using SCP, or open the files and copy them into "vi" on the SRX Series device to create new files.

   user@host> request security pki local-certificate load certificate-id ssl-inspect-ca key /var/tmp/ssl-inspect-ca.key filename /var/tmp/ssl-inspect-ca.cer passphrase password

Copyright © 2016, Juniper Networks, Inc.

6. From configuration mode, apply the loaded certificate as root-ca in the SSL proxy profile.

   ```
   [edit]
   user@host# set services ssl proxy profile ssl-inspect-profile root-ca ssl-inspect-ca
   ```

7. Import the root CA as a trusted CA into client browsers. This is required for the client browsers to trust the certificates signed by the SRX Series device. See "Importing a Root CA Certificate into a Browser" on page 545.

Configuring a CA Profile Group

The CA profile defines the certificate information to be used for authentication. It includes the public key that SSL proxy uses when generating a new certificate. Junos OS allows you to create a group of CA profiles and load multiple certificates in one action, view information about all certificates in a group, and delete unwanted CA groups.

You can load a group of CA profiles by obtaining a list of trusted CA certificates, defining a CA group, and attaching the CA group to the SSL proxy profile.

1. Obtain a list of trusted CA certificates by following one of these methods:

   - Junos OS provides a default list of trusted CA certificates that you can load on your system using the **default** command option. The Junos OS package contains the default CA certificates as a PEM file (for example, **trusted_CA.pem**). After you download the Junos OS package and reboot your device, the default certificates are available on your system.

     From operational mode, load the default trusted CA certificates (the group name identifies the CA profile group):

     ```
     user@host> request security pki ca-certificate ca-profile-group load ca-group-name
       group-name filename default
     ```

   - Alternatively, you can define your own list of trusted CA certificates and import them on your system. You get the list of trusted CAs in a single PEM file (for example **IE-all.pem**) and save the PEM file in a specific location (for example, **/var/tmp**). See Knowledge Base Article KB23144.

     From operational mode, load the trusted list to the device (the group name identifies the CA profile group):

     ```
     user@host> request security pki ca-certificate ca-profile-group load ca-group-name
       group-name filename /var/tmp/IE-all.pem
     ```

2. From configuration mode, attach the CA profile group to the SSL proxy profile. You can attach one or all CA profile groups at a time:

   - To attach one CA profile group (the group name identifies the CA profile group):

     ```
     [edit]
     user@host# set services ssl proxy profile profile-name trusted-ca group-name
     ```

   - To attach all CA profile groups:

     ```
     [edit]
     user@host# set services ssl proxy profile profile-name trusted-ca all
     ```

Copyright © 2016, Juniper Networks, Inc.

You can easily display information about all certificates in a CA profile group:

```
user@host> show security pki ca-certificates ca-profile-group group-name
```

You can delete a CA profile group. Remember that deleting a CA profile group deletes all certificates that belong to that group:

```
user@host> clear security pki ca-certificates ca-profile-group group-name
```

Configuring a Trusted CA Profile

Typically, you import a list of trusted CA certificates by creating a group of CA profiles. However, you can also configure a single CA profile (containing one or multiple certificates) and import it using PKI commands. This section shows you how to import a trusted CA certificate from your browser's certificate store into your SRX Series device. The certificate that is configured under the trusted CA is loaded using the PKI commands and is used for validating the server certificate chain.

1. From configuration mode, configure the CA profile used for loading the certificate.

   ```
   [edit]
   user@host# set security pki ca-profile profile-name ca-identity ca-identity
   ```

2. From operational mode, load the certificate using PKI commands.

   ```
   user@host> request security pki ca-certificate load ca-profile profile-name filename
     filename
   ```

3. From configuration mode, disable the revocation check.

   NOTE: CRL checks are not supported; we recommend that you disable revocation checks.

   ```
   [edit]
   user@host# set security pki ca-profile profile-name ca-identity ca-identity
     revocation-check disable
   ```

4. From configuration mode, configure the loaded certificate as a trusted CA in the SSL proxy profile.

   ```
   [edit]
   user@host# set services ssl proxy profile ssl-proxy-profile-name trusted-ca
     ca-profile-name
   ```

   NOTE: More than one trusted CA can be configured for a profile.

5. (Optional) If you have multiple trusted CA certificates, you do not have to specify each trusted CA separately. You can load *all* the trusted CA certificates using the following command from configuration mode.

   ```
   [edit]
   user@host# set services ssl proxy profile ssl-proxy-profile-name trusted-ca all
   ```


Copyright © 2016, Juniper Networks, Inc.

NOTE: Alternatively, you can import a set of trusted CAs from your browser into the SRX Series device. See Knowledge Base article KB23144.

Importing a Root CA Certificate into a Browser

In order to have your browser or system automatically trust all certificates signed by the root CA configured in the SSL proxy profile, you must instruct your platform or browser to trust the CA root certificate.

To import a root CA certificate:

1. Generate a PEM format file for the configured root CA.

   request security pki local-certificate export certificate-id *root-ca* type pem filename *path/file-name.pem*

2. Import a root CA certificate into a browser.

   From Internet Explorer (version 8.0):

   a. From the Tools menu, choose **Internet Options**.

   b. On the Content tab, click **Certificates**.

   c. Select the **Trusted Root Certification Authorities** tab and click **Import**.

   d. In the Certificate Import Wizard, navigate to the required root CA certificate and select it.

   From Firefox (version 39.0):

   a. From the Tools menu, choose **Options**.

   b. From the Advanced menu, select the **Certificates** tab and click **View Certificate**.

   c. In the Certificate Manager window, select the **Authorities** tab and click **Import**.

   d. Navigate to the required root CA certificate and select it.

   From Google Chrome (45.0):

   a. From the Settings menu, choose **Show Advanced Settings**.

   b. From the Advanced menu, select the **Certificates** tab and click **View Certificate**.

   c. Under HTTPS/SSL, click **Manage Certificates**.

   d. In the Certificate window, select **Trusted Root Certification Authorities** and click **Import**.

   e. In the Certificate Import Wizard, navigate to the required root CA certificate and select it.

Copyright © 2016, Juniper Networks, Inc.

Applying an SSL Proxy Profile to a Security Policy

SSL proxy is enabled as an application service within a security policy. In a security policy, you specify the traffic that you want the SSL proxy enabled on as match criteria and then specify the SSL proxy CA profile to be applied to the traffic. Figure 32 displays a graphical view of SSL proxy profile and security policy configuration.

Figure 226: Applying an SSL Proxy Profile to a Security Policy

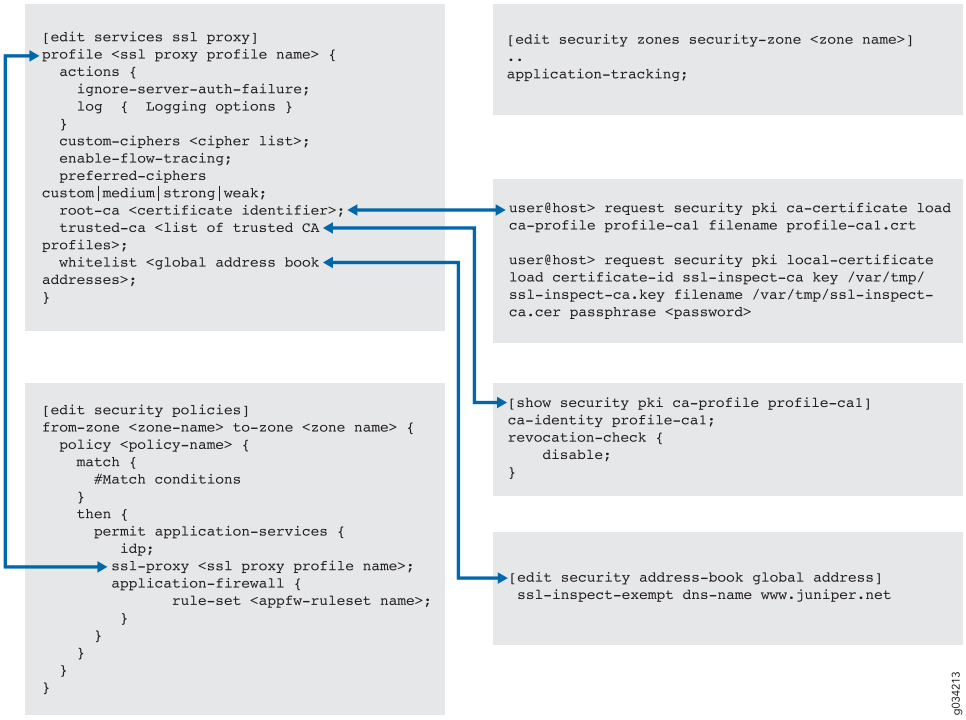

To enable SSL proxy in a security policy:

1. Create a security policy and specify the match criteria for the policy. As match criteria, specify the traffic for which you want to enable SSL proxy.

   [edit]
   user@host# set security policies from-zone trust to-zone untrust policy *policy-name* match source-address *source-address*
   user@host# set security policies from-zone trust to-zone untrust policy *policy-name* match destination-address *destination-address*
   user@host# set security policies from-zone trust to-zone untrust policy *policy-name* match application *application*

2. Apply the SSL proxy profile to the security policy.

   [edit]
   user@host# set security policies from-zone trust to-zone untrust policy *policy-name* then permit application-services ssl-proxy profile-name *profile-name*

Copyright © 2016, Juniper Networks, Inc.

Creating a Whitelist of Exempted Destinations

SSL encryption and decryption are complicated and expensive procedures. You can selectively bypass SSL proxy processing for some sessions by configuring a whitelist. Typically, you would configure the whitelist to include trusted servers or domains with which you are very familiar. You might also include financial and banking sites that you are legally required to include.

Whitelists include addresses that you want to exempt from undergoing SSL proxy processing. For example, if you want to exempt all sessions to www.example.net then you would include it in the whitelist. To configure the whiltelist, you specify the domain that you want to exempt in an address book and then configure the address in the SSL proxy profile.

1. Configure the domain in the address book.

   ```
   [edit]
   user@host# set security address-book global address address dns-name
     www.example.net
   ```

2. Specify the global address book address in the SSL proxy profile.

   ```
   [edit]
   user@host# set services ssl proxy profile profile-name whitelist address
   ```

Whitelist addresses and address sets are created under the global address book. The following type of addresses (from the global address book) are supported:

- IPv4 addresses (plain text). For example:

  ```
  [edit]
  user@host# set security address-book global address address-name ipv4-prefix
  ```

- IPv4 address range. For example:

  ```
  [edit]
  user@host# set security address-book global address address-name range-address
    range-low to range-high
  ```

- IPv4 wildcard. For example:

  ```
  [edit]
  user@host# set security address-book global address address-name wildcard-address
    a.d.d.r/netmask
  ```

  Noncontiguous netmasks are not supported. For example:

  - 5.0.0.1/255.255.0.0 is supported.
  - 5.0.0.1/255.255.0.255 is NOT supported.

- IPv6 address (plain text). For example:

  ```
  [edit]
  user@host# set security address-book global address address-name ipv6-prefix
  ```

- DNS name. For example:

  ```
  [edit]
  ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security address-book global address address-name dns-name
  domain-name
```

- Translated IP addresses. Sessions are whitelisted based on the actual IP address and not on the translated IP address. Because of this, in the whitelist configuration of the SSL proxy profile, the actual IP address should be provided and not the translated IP addresses.

  For example, consider a destination NAT rule that translates destination IP address **192.0.2.10/24** to **198.51.100.8/24** using the following commands:

```
[edit]
user@host# set security nat destination pool d1 address 198.51.100.8/24
user@host# set security nat destination rule-set dst-nat rule r1 match
  destination-address 192.0.2.10/24
user@host# set security nat destination rule-set dst-nat rule r1 then destination-nat
  pool d1
```

In this scenario, to exempt a session from SSL proxy inspection, the following IP address should be added to the whitelist:

```
[edit]
user@host# set security address-book global address ssl-proxy-exempted-addr
  192.0.2.10/24
user@host# set services ssl proxy profile ssl-inspect-profile whitelist
  ssl-proxy-exempted-addr
```

Configuring SSL Proxy Logging

When configuring SSL proxy, you can choose to set the option to receive some or all of the logs. SSL proxy logs contain the logical system name, SSL proxy whitelists, policy information, SSL proxy information, and other information that helps you troubleshoot when there is an error.

You can configure logging of *all* or specific events, such as error, warning, and information events. You can also configure logging of sessions that are whitelisted, dropped, ignored, or allowed after an error occurs.

```
[edit]
user@host# set services ssl proxy profile profile-name actions log all
user@host# set services ssl proxy profile profile-name actions log sessions-whitelisted
user@host# set services ssl proxy profile profile-name actions log sessions-allowed
user@host# set services ssl proxy profile profile-name actions log errors
```

You can use **enable-flow-tracing** option to enable debug tracing.

Configuring Ciphers

You can configure the following ciphers for an SSL proxy profile:

- **preferred-ciphers**–Preferred ciphers allow you to define an SSL cipher that can be used with acceptable key strength. Ciphers are divided in three categories depending on their key strength: strong, medium, or weak.
- **custom-ciphers**–Custom ciphers allow you to define your own cipher list. If you do not want to use one of the three categories, you can select ciphers from each of the

Copyright © 2016, Juniper Networks, Inc.

categories to form a custom cipher set. To configure custom ciphers, you must set **preferred-ciphers** to custom.

The following example shows how to create a custom cipher. In this example, you set **prefered-cipher** to custom and add the cipher list (rsa-with-3des-ede-cbc-sha and rsa-with-aes-256-cbc-sha):

```
set services ssl proxy profile profile-name preferred-ciphers custom
set services ssl proxy profile profile-name custom-ciphers rsa-with-3des-ede-cbc-sha
set services ssl proxy profile profile-name custom-ciphers rsa-with-aes-256-cbc-sha
```

Exporting Certificates to a Specified Location

When a self-signed certificate is generated using a PKI command, the newly generated certificate is stored in a predefined location (**var/db/certs/common/local**).

Use the following command to export the certificate to a specific location (within the device). You can specify the certificate ID, the filename, and the type of file format (DER/PEM):

```
user@host> request security pki local-certificate export certificate-id certificate-id
user@host> request security pki local-certificate export filename filename
user@host> request security pki local-certificate export type der
```

Ignoring Server Authentication

Junos OS allows you to configure an option to ignore server authentication completely. If you configure your system to ignore authentication, then any errors encountered during server certificate verification at the time of the SSL handshake are ignored. Commonly ignored errors include the inability to verify CA signature, incorrect certificate expiration dates, and so forth. If this option is not set, all the sessions where the server sends self-signed certificates are dropped when errors are encountered.

We do not recommend using this option for authentication because configuring it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause of dropped SSL sessions.

From configuration mode, specify to ignore server authentication:

```
[edit]
user@host# set services ssl proxy profile profile-name actions ignore-server-auth-failure
```

Related Documentation

- SSL Proxy Overview on page 531
- Enabling Debugging and Tracing for SSL Proxy on page 552
- Understanding Self-Signed Certificates on page 6433
- show services ssl proxy statistics on page 829
- clear services ssl proxy statistics on page 752

Copyright © 2016, Juniper Networks, Inc.

Enabling Debugging and Tracing for SSL Proxy

Debug tracing on both Routing Engine and the Packet Forwarding Engine can be enabled for SSL proxy by setting the following configuration:

```
user@host# set services ssl traceoptions
```

Table 37 shows the supported levels for trace options.

Table 468: Trace Levels

Cause Type	Description
Brief	Only error traces on both the Routing Engine and the Packet Forwarding Engine.
Detail	Packet Forwarding Engine–Only event details up to the handshake should be traced. Routing Engine–Traces related to commit. No periodic traces on the Routing Engine will be available
Extensive	Packet Forwarding Engine–Data transfer summary available. Routing Engine–Traces related to commit (more extensive). No periodic traces on the Routing Engine will be available.
Verbose	All traces are available.

Table 38 shows the flags that are supported.

Table 469: Supported Flags in Trace

Cause Type	Description
cli-configuration	Configuration-related traces only.
initiation	Enable tracing on the SSL-I plug-in.
proxy	Enable tracing on the SSL-Proxy-Policy plug-in.
termination	Enable tracing on the SSL-T plug-in.
selected-profile	Enable tracing only for profiles that have **enable-flow-tracing** set.

You can enable logs in the SSL proxy profile to get to the root cause for the drop. The following errors are some of the most common:

- Server certification validation error. Check the trusted CA configuration to verify your configuration.
- System failures such as memory allocation failures.
- Ciphers do not match.
- SSL versions do not match.

Copyright © 2016, Juniper Networks, Inc.

- SSL options are not supported.
- Root CA has expired. You need to load a new root CA.

You can enable the **ignore-server-auth-failure** option in the SSL proxy profile to ensure that certificate validation, root CA expiration dates, and other such issues are ignored. If sessions are inspected after the **ignore-server-auth-failure** option is enabled, the problem is localized.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 214

Configuring Pass-Through Authentication

- Understanding Pass-Through Authentication on page 5205
- Example: Configuring Pass-Through Authentication on page 5207
- Example: Configuring HTTPS Traffic to Trigger Pass-Through Authentication on page 5213

Understanding Pass-Through Authentication

Pass-through user authentication is a form of active authentication; the user is prompted to enter a username and password when pass-through authentication is invoked. If the user's identity is validated, the user is allowed to pass through the firewall and gain access to the requested resources.

When a user attempts to initiate an HTTP, an HTTPS, an FTP, or a Telnet connection request that has a policy requiring authentication, the device intercepts the request and prompts the user to enter a username and password. Depending on the configuration, the device validates the username and password by checking them against those stored in the local database or on an external authentication server.

If an external authentication server is used, after the user's credentials are collected, they are processed through firewall user authentication. The following external authentication servers are supported:

- RADIUS authentication and authorization (compatible with Juniper Steel-Belted Radius servers)

  You can use an external RADIUS server if, in addition to authentication, you want to obtain authorization information about the user's access right (what the user can do on the network).
- LDAP authentication only (supports LDAP version 3, compatible with Windows AD)
- SecurID authentication only (uses an RSA SecurID external authentication server)

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall.

user@host# set access profile profile1 client joan firewall-user password talk

You can put several user accounts together to form a user group, which you can store on the local database or on a RADIUS, an LDAP, or a SecurID server. When you reference

Copyright © 2016, Juniper Networks, Inc.

an authentication user group and an external authentication server in a policy, the traffic matching the policy triggers an authentication check.

NOTE: You use family inet to assign an IPv4 address. You use family inet6 to assign an IPv6 address. An interface can be configured with both an IPv4 and an IPv6 address. For the sake of brevity, these examples use IPv4 addresses only.

Figure 227: Policy Lookup for a User

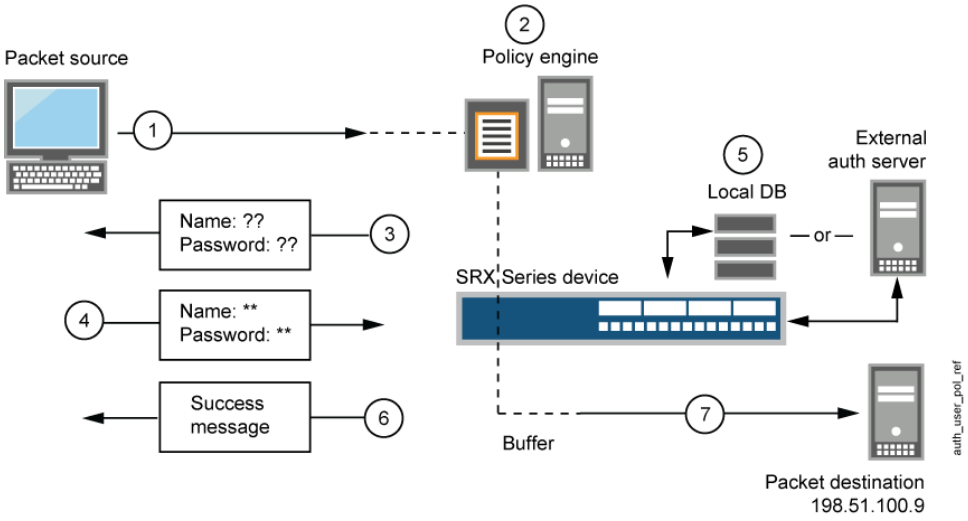

The steps in Figure 227 are as follows:

1. A client user sends an FTP, an HTTP, an HTTPS, or a Telnet packet to 198.51.100.9.
2. The device intercepts the packet, notes that its policy requires authentication from either the local database or an external authentication server, and buffers the packet.
3. The device prompts the user for login information through FTP, HTTP, HTTPS, or Telnet.
4. The user replies with a username and password.
5. The device either checks for an authentication user account on its local database or sends the login information to the external authentication server as specified in the policy.
6. Finding a valid match (or receiving notice of such a match from the external authentication server), the device informs the user that the login has been successful.
7. For HTTP, HTTPS, or Telnet traffic, the device forwards the packet from its buffer to its destination IP address, 198.51.100.9. However, for FTP traffic, after successful authentication the device closes the session and the user must reconnect to the FTP server at IP address 198.51.100.9.

After the device authenticates a user at a particular source IP address, it subsequently permits traffic—as specified in the policy requiring authentication through pass

Copyright © 2016, Juniper Networks, Inc.

through—from any other user at that same address. This might be the case if the user originates traffic from behind a NAT device that changes all original source addresses to a single translated address.

NOTE: The pass-through user authentication method is recommended in situations when security has a higher priority than convenience. This authentication method applies only to the session and child sessions matching the policy that triggered it. You can apply this method on Internet-facing links, if used with caution.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Understanding Web Authentication on page 5221
- Example: Configuring Pass-Through Authentication on page 5207

Example: Configuring Pass-Through Authentication

This example shows how to configure pass-through authentication to authenticate firewall users. A firewall user is a network user who must provide a username and password when initiating a connection across the firewall.

Pass-through authentication allows SRX Series administrators to restrict users who attempt to access a resource in another zone using FTP, Telnet, HTTP, or HTTPS. If the traffic matches a security policy whose action is pass-through authentication, the user is required to provide login information.

For HTTPS, to ensure security the HTTPS default certificate key size is 2048 bits. If you do not specify a certificate size, the default size is assumed.

- Requirements on page 5207
- Overview on page 5208
- Configuration on page 5208
- Verification on page 5212

Requirements

Before you begin, define firewall users. See Firewall User Authentication Overview.

This example uses the following hardware and software components:

- SRX Series device
- Firewall user's system
- Packet destination system

Copyright © 2016, Juniper Networks, Inc.

Overview

The pass-through authentication process is triggered when a client, referred to as a firewall user, attempts to initiate an FTP, a Telnet, or an HTTP session to access a resource in another zone. The SRX Series firewall acts as a proxy for an FTP, a Telnet, an HTTP, or an HTTPS server so that it can authenticate the firewall user before allowing the user access to the actual FTP, Telnet, or HTTP server behind the firewall.

If traffic generated from a connection request sent by a firewall user matches a security policy rule bidirectionally and that rule specifies pass-through firewall authentication as the action of its **then** clause, the SRX Series device requires the firewall user to authenticate to a Junos OS proxy server.

If the authentication is successful, subsequent traffic from the same source IP address is automatically allowed to pass through the SRX Series device if the traffic matches the security policy tuples.

Figure 228 shows the topology used in this example.

Figure 228: Configuring Pass-Through Firewall Authentication

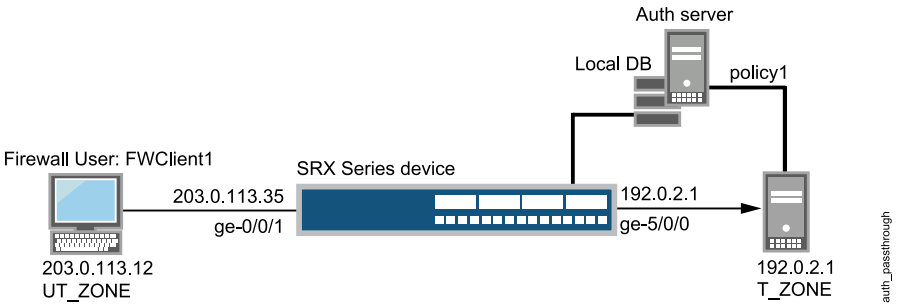

NOTE: Although the topology shows use of an external server, it is not covered in the configuration. It is outside the scope of this example.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 203.0.113.35/24
set interfaces ge-5/0/0 unit 0 family inet address 192.0.2.1/24
set access profile FWAUTH client FWClient1 firewall-user password pwd
set access firewall-authentication pass-through default-profile FWAUTH
set access firewall-authentication pass-through telnet banner success "WELCOME TO JUNIPER TELNET SESSION"
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone UT-ZONE host-inbound-traffic system-services all
set security zones security-zone UT-ZONE interfaces ge-0/0/1.0 host-inbound-traffic protocols
all
set security zones security-zone T-ZONE host-inbound-traffic system-services all
set security zones security-zone T-ZONE interfaces ge-5/0/0.0 host-inbound-traffic protocols
all
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match source-address any
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match destination-address
any
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match application junos-telnet
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 then permit
firewall-authentication pass-through client-match FWClient1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see Using the CLI Editor in Configuration Mode.

To configure pass-through authentication:

1. Configure two interfaces and assign IP addresses to them.

   NOTE: For this example, it is optional to assign two addresses to the interfaces.

   ```
   [edit]
   user@host# set interfaces ge-0/0/1 unit 0 family inet address 203.0.113.35/24
   user@host# set interfaces ge-5/0/0 unit 0 family inet address 203.0.113.35/24
   ```

2. Create the FWAUTH access profile for the FWClient1 user, specify the user's password, and define a success banner for Telnet sessions.

   ```
   [edit access]
   user@host# set profile FWAUTH client FWClient1 firewall-user password pwd
   user@host# set firewall-authentication pass-through default-profile FWAUTH
   user@host# set firewall-authentication pass-through telnet banner success
     "WELCOME TO JUNIPER TELNET SESSION"
   ```

3. Configure security zones.

   NOTE: For this example, it is optional to configure a second interface for a security zone.

   ```
   [edit security zones]
   user@host# set security-zone UT-ZONE host-inbound-traffic system-services all
   user@host# set security-zone UT-ZONE interfaces ge-0/0/1.0 host-inbound-traffic
     protocols all
   user@host# set security-zone T-ZONE host-inbound-traffic system-services all
   user@host# set security-zone T-ZONE interfaces ge-5/0/0.0 host-inbound-traffic
     protocols all
   ```

4. Assign security policy P1 to the security zones.

   ```
   [edit security policies]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match
  source-address any
user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match
  destination-address any
user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match application
  junos-telnet
user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 then permit
  firewall-authentication pass-through client-match FWClient1
```

5. Use Telnet to authenticate the FWClient1 firewall user to host2.

```
user@FWClient1# run telnet 192.0.2.1/24
Trying 192.0.2.1/24...
Connected to 192.0.2.1/24
Escape character is '^]'.
Firewall User Authentication
Username: FWClient1
Password:***
WELCOME TO JUNIPER TELNET SESSION
Host1 (ttyp0)
login: user
Password: ***
--- JUNOS 10.1R1.1 built 2009-10-12 13:30:18 UTC
%
```

Results From configuration mode, confirm your configuration by entering these commands.

- show interfaces
- show access
- show security zones
- show security policies

If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, the output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@host# show interfaces
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 203.0.113.35;
      }
    }
  }
  ge-5/0/0 {
    unit 0 {
      family inet {
        address 192.0.2.1/24;
      }
    }
  }
...
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# show access
  profile FWAUTH {
    authentication-order password;
    client FWClient1 {
      firewall-user {
        password "$ABC123"; ## SECRET-DATA
      }
    }
  }
  firewall-authentication {
    pass-through {
      default-profile FWAUTH;
      telnet {
        banner {
          success "WELCOME TO JUNIPER TELNET SESSION";
        }
      }
    }
  }

user@host# show security zones
  security-zone UT-ZONE {
    host-inbound-traffic {
      system-services {
        all;
      }
    }
      interfaces {
        ge-0/0/1.0 {
          host-inbound-traffic {
            protocols {
              all;
            }
          }
        }
      }
      security-zone T-ZONE {
        host-inbound-traffic {
          system-services {
            all;
          }
        }
        interfaces {
          ge-5/0/0.0 {
            host-inbound-traffic {
              protocols {
                all;
              }
            }
          }
        }
      }

user@host# show security policies
...
from-zone UT-ZONE to-zone T-ZONE {
  policy P1 {
```

```
        match {
          source-address any;
          destination-address any;
          application junos-telnet;
        }
        then {
          permit {
            firewall-authentication {
              pass-through {
                client-match FWClient1;
              }
            }
          }
        }
      }
    }
```

If you are done configuring the device, enter commit from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table on page 5212

Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table

Purpose Display firewall authentication user history and verify the number of firewall users who successfully authenticated and the number of firewall users who failed to log in.

Action From operational mode, enter these **show** commands:

```
user@host> show security firewall-authentication history
History of firewall authentication data:
Authentications: 2
Id Source Ip Date Time Duration Status User
1 203.0.113.12 2010-10-12 21:24:02 0:00:24 Failed FWClient1
2 203.0.113.12 2010-10-12 21:24:48 0:00:22 Success FWClient1

user@host> show security firewall-authentication history identifier 1
Username: FWClient1
Source IP: 203.0.113.12
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2010-10-12
Access start time: 21:24:02
Duration of user access: 0:00:24
Source zone: UT-ZONE
Destination zone: T-ZONE
Access profile: FWAUTH
Bytes sent by this user: 0
Bytes received by this user: 2660

user@host> show security firewall-authentication users
```


Copyright © 2016, Juniper Networks, Inc.

```
Firewall authentication data:
Total users in table: 1
Id Source Ip Src zone Dst zone Profile Age Status User
4 203.0.113.12     UT-ZONE  T-ZONE   FWAUTH     1 Success   FWClient1

user@host> show security firewall-authentication users identifier 3
Username: FWClient1
Source IP: 203.0.113.12
Authentication state: Success
Authentication method: Pass-through using Telnet
Age: 3
Access time remaining: 9
Source zone: UT-ZONE
Destination zone: T-ZONE
Access profile: FWAUTH
Interface Name: ge-0/0/1.0
Bytes sent by this user: 0
Bytes received by this user: 1521
```

Related Documentation

- Firewall User Authentication Overview on page 5179
- Understanding Pass-Through Authentication on page 5205

Example: Configuring HTTPS Traffic to Trigger Pass-Through Authentication

This example shows how to configure HTTPS traffic to trigger pass-through authentication. HTTPS is more secure than HTTP, so it has become more popular and is more widely used.

- Requirements on page 5213
- Overview on page 5214
- Configuration on page 5215
- Verification on page 5219

Requirements

This example uses the following hardware and software components:

- SRX Series device
- Two PCs running Linux and Open SSL. One PC acts as a client and another as an HTTPS server. The two PCs are used to create key files and to send traffic.
- Junos OS Release 12.1X44-D10 or later for high-end SRX Series Services Gateways and Junos OS Release 15.1X49-D40 or later for vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways.

Before you begin:

An SRX Series device has to decode HTTPS traffic to trigger pass-through authentication. Then, SSL termination proxy creates and installs a private key file and a certification file. The following list describes the steps to create and install a private key file and a certification key file.

Copyright © 2016, Juniper Networks, Inc.

NOTE: If you have an official .crt file and .key file, then you can directly upload and install the files on the SRX Series device. If you do not have a .crt file and .key file, follow the procedure to create and install the files. Instructions specified in Step 1 and Step 2 must be run on a PC with Linux and OpenSSL installed. Instructions specified in Step 3 and Step 4 must be run in operational mode.

To create and install a private key file and a certification file:

1. On a PC create the **.key** file.

   openssl genrsa -out /tmp/device.key 1024

2. On a PC, create the **.crt** file.

   openssl req -new -x509 -days 365 -key /tmp/server.key -out /tmp/device.crt -subj "/C=CN/ST=BJ/L=BJ/O=JNPR/OU=CNRD/CN=2.2.2.10/emailAddress=device@juniper.net"

3. Upload the **.key** and **.crt** files to an SRX Series device, and install the files on the device using the following command from operational mode:

   user@host> **request security pki local-certificate load filename /var/tmp/device.crt key /var/tmp/device.key certificate-id device**

Overview

Firewall authentication initiates a secure connection to be established across two devices. A network user must provide a username and password for authentication when initiating a connection across the firewall. Firewall authentication supports HTTPS traffic for pass-through authentication. HTTPS can secure HTTP firewall authentication traffic between users and the SRX Series device.

HTTPS is the secure version of HTTP, the protocol over which data is sent between the user and the device that the user is connected to. All communications between the user and the connected devices are encrypted. HTTPS is often used to protect highly confidential online transactions like online banking and online shopping order forms.

In this example, HTTPS traffic is used to trigger pass-through authentication because HTTPS is more secure than HTTP. For HTTPS traffic to trigger pass-through authentication you must first configure the SSL termination profile.

Figure 229 shows an example of pass-through authentication using HTTPS traffic. In this example, a host or a user from an untrust zone tries to access resources on the trust zone. The SRX Series device uses HTTPS to collect the username and password information. Subsequent traffic from the host or user is allowed or denied based on the result of this authentication.

Copyright © 2016, Juniper Networks, Inc.

Figure 229: Pass-Through Authentication Using HTTPS Traffic

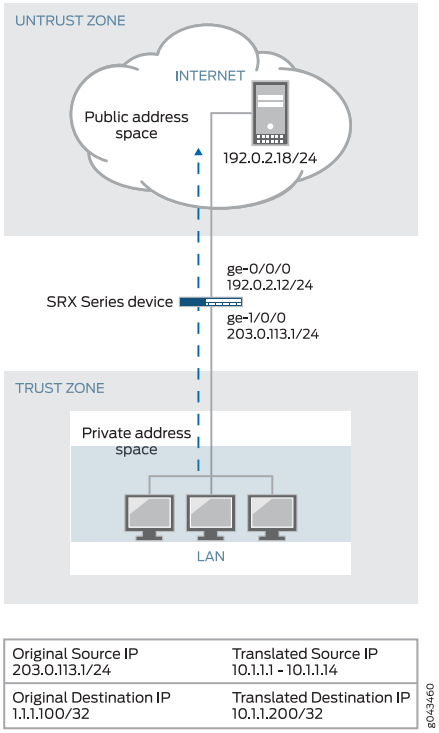

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.12/24
set interfaces ge-1/0/0 unit 0 family inet address 203.0.113.1/24
set security policies from-zone trust to-zone untrust policy p1 match source-address any
set security policies from-zone trust to-zone untrust policy p1 match destination-address
    any
set security policies from-zone trust to-zone untrust policy p1 match application any
set security policies from-zone trust to-zone untrust policy p1 then permit
    firewall-authentication pass-through access-profile local_pf
set security policies from-zone trust to-zone untrust policy p1 then permit
    firewall-authentication pass-through ssl-termination-profile ssl_pf
set security policies from-zone trust to-zone untrust policy p1 then log session-init
set security policies from-zone trust to-zone untrust policy p1 then log session-close
set security zones security-zone trust interfaces ge-0/0/0.0 host-inbound-traffic
    system-services all
set security zones security-zone trust interfaces ge-0/0/0.0 host-inbound-traffic protocols
    all
set security zones security-zone untrust interfaces ge-1/0/0.0 host-inbound-traffic
    system-services all
set security zones security-zone untrust interfaces ge-1/0/0.0 host-inbound-traffic protocols
    all
```

Copyright © 2016, Juniper Networks, Inc.

```
set access profile local_pf client user1 firewall-user password user1
set access firewall-authentication pass-through default-profile local_pf
set services ssl termination profile ssl_pf server-certificate device
```

Step-by-Step Procedure

To configure HTTPS traffic to trigger pass-through authentication:

1. Configure interfaces and assign IP addresses.

   ```
   [edit interfaces]
   user@host# set ge-0/0/0 unit 0 family inet address 192.0.2.12/24
   user@host# set ge-1/0/0 unit 0 family inet address 203.0.113.1/24
   ```

2. Configure security policies to permit firewall authenticated traffic from zone trust to zone untrust.

   ```
   [edit security policies]
   user@host# set from-zone trust to-zone untrust policy p1 then permit
     firewall-authentication pass-through access-profile local_pf
   user@host# set from-zone trust to-zone untrust policy p1 then permit
     firewall-authentication pass-through ssl-termination-profile ssl_pf
   ```

3. Specify a policy action to take when a packet matches the criteria.

   ```
   [edit security policies]
   user@host# set from-zone trust to-zone untrust policy p1 match source-address
     any
   user@host# set from-zone trust to-zone untrust policy p1 match destination-address
     any
   user@host# set from-zone trust to-zone untrust policy p1 match application any
   user@host# set from-zone trust to-zone untrust policy p1 then log session-init
   user@host# set from-zone trust to-zone untrust policy p1 then log session-close
   ```

4. Configure security zones and assign interfaces.

   ```
   [edit security zones]
   user@host#  set security-zone trust interfaces ge-0/0/0.0 host-inbound-traffic
     protocols all
   user@host# set security-zone trust interfaces ge-0/0/0.0 host-inbound-traffic
     system-services all
   ```

5. Configure application services for zones.

   ```
   [edit security zones]
   user@host# set security-zone trust host-inbound-traffic system-services all
     protocols all
   user@host# set security-zone untrust host-inbound-traffic system-services all
     protocols all
   ```

6. Create an access profile and configure the client as a firewall user and set the password.

   ```
   [edit access]
   user@host# set profile local_pf client user1 firewall-user password user1
   ```

7. Configure the type of firewall and the default profile name where the authentication settings are defined.

   ```
   [edit access]
   user@host# set firewall-authentication pass-through default-profile local_pf
   ```


Copyright © 2016, Juniper Networks, Inc.

8. Configure the SSL termination profile and enter a local certificate identifier name.

```
[edit services]
user@host# set ssl termination profile ssl_pf server-certificate device
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security policies**, **show security zones**, **show access**, and **show services ssl termination** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show interfaces
...
interfaces
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.0.2.12/24;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      family inet {
        address 203.0.113.1/24;
      }
    }
  }
```

```
user@host# show security policies
...
policies
  from-zone trust to-zone untrust {
    policy p1 {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          firewall-authentication {
            pass-through {
              access-profile local_pf;
              ssl-termination-profile ssl_pf;
            }
          }
        }
        log {
          session-init;
          session-close;
        }
      }
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security zones
...
zones {
    security-zone trust {
        interfaces {
            ge-0/0/0.0 {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
            }
        }
    }
    security-zone untrust {
        interfaces {
            ge-1/0/0.0 {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
            }
        }
    }
```

```
user@host# show access
...
access {
    profile local_pf {
        client user1 {
            firewall-user {
                password "$ABC123”;
            }
        }
    }
    firewall-authentication {
        pass-through {
            default-profile local_pf;
        }
    }
```

```
user@host# show services ssl termination
...
services {
    ssl {
        termination {
            profile ssl_pf {
                server-certificate device;
            }
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration

Purpose Verify that the configuration is correct.

Action From operational mode, enter the **show security firewall-authentication users** command for identifier 1.

```
user@host> show security firewall-authentication users identifier 1
  Username: user1
  Source IP: 203.0.113.1/24
  Authentication state: Success
  Authentication method: Pass-through using HTTPS
  Age: 0
  Access time remaining: 10
  Lsys: root-logical-system
  Source zone: trust
  Destination zone: untrust
  Access profile: local_pf
  Interface Name: ge-0/0/0.0
  Bytes sent by this user: 946
Bytes received by this user: 0
```

Meaning The **show security firewall-authentication users** command displays the firewall authentication user information for the specified identifier. If the output displays Pass-through using HTTPS in the Authentication method field and Success in the Authentication state field, then your configuration is correct.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Understanding Pass-Through Authentication on page 5205
- Example: Configuring Pass-Through Authentication on page 5207

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 215

Configuring Web Authentication

- Understanding Web Authentication on page 5221
- Example: Configuring Web Authentication on page 5223
- Example: Configuring HTTPS Traffic to Trigger Web Authentication on page 5230

Understanding Web Authentication

Web authentication is an alternative to pass-through user authentication. Instead of pointing to the resource that you want to connect to from your client browser, you point the browser to an IP address on the device that is enabled for Web authentication. This initiates an HTTP session to the IP address hosting the Web authentication feature on the device. The device then prompts you for your username and password and caches the result in the device. Later, when traffic encounters a Web authentication policy, you are allowed or denied access based on the prior Web authentication results, as shown in Figure 230.

NOTE: You use family inet to assign an IPv4 address. You use family inet6 to assign an IPv6 address. An interface can be configured with both an IPv4 and an IPv6 address. For the sake of brevity, these examples use IPv4 addresses only.

Copyright © 2016, Juniper Networks, Inc.

Figure 230: Web Authentication Example

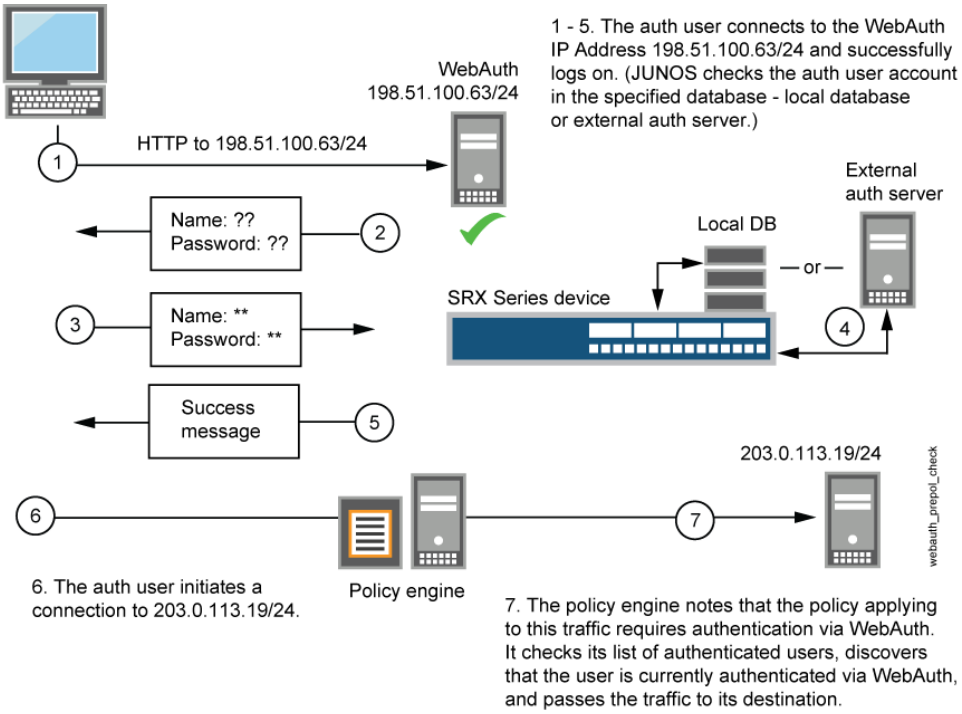

Follow these Web authentication guidelines:

- You can leave the default Web authentication server as the local database or you can choose an external authentication server for the role. The default Web authentication profile determines if the user authenticates using the local database or the external authentication server. An access profile stores usernames and passwords of users or points to external authentication servers where such information is stored.
- The Web authentication address must be in the same subnet as the interface that you want to use to host it. For example, if you want authentication users to connect using Web authentication through ethernet3, which has IP address 203.0.113.1/24, then you can assign Web authentication an IP address in the 203.0.113.0/24 subnet.
- You can put a Web authentication address in the same subnet as the IP address of any physical interface or virtual security interface (VSI). (For information about different types of interfaces, see "Security Zones and Interfaces Overview" on page 1049.)
- You can put Web authentication addresses on multiple interfaces.
- After a device authenticates a user at a particular source IP address, it subsequently permits traffic—as specified in the policy requiring authentication through Web authentication—from any other user at that same address. This might be the case if the user originates traffic from behind a NAT device that changes all original source addresses to a single translated address.
- With Web authentication enabled, any HTTP traffic to the IP address will get the Web authentication login page instead of the administrator login page. Disabling this option

Copyright © 2016, Juniper Networks, Inc.

will show the administrator login page (assuming that **[system services web-management HTTP]** is enabled.

- We recommend that you have a separate primary or preferred IP address, if an address is used for Web authentication.

NOTE: The Web authentication method is recommended in situations when the client devices are immediately adjacent to the security gateway and there is high assurance that the client devices are not multiuser hosts. This authentication method is best applied to wireless links and DMZ, or conference room links.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Understanding Pass-Through Authentication on page 5205
- Example: Configuring Web Authentication on page 5223

Example: Configuring Web Authentication

This example shows how to enable Web authentication and set up a policy that allows access to a user when traffic encounters a policy that has Web authentication enabled.

- Requirements on page 5223
- Overview on page 5223
- Configuration on page 5224
- Verification on page 5228

Requirements

Before you begin:

- Define firewall users. See "Firewall User Authentication Overview" on page 5179.
- Add the Web authentication HTTP flag under the interface's address hierarchy to enable Web authentication.

Overview

To enable Web authentication, you must specify the IP address of the device hosting the HTTP session. These settings are used if the firewall user accessing a protected resource wants to be authenticated by directly accessing the webserver or by Web authentication. The following instructions show how to set up a policy that allows access to the FWClient1 user when traffic encounters a policy that has Web authentication enabled (Policy-W). (See Figure 231.) In this example, FWClient1 has already authenticated through the Web authentication login page.

The FWClient1 firewall user does the following to get authenticated:

Copyright © 2016, Juniper Networks, Inc.

a. Points the browser to the Web authentication IP (198.51.100.63/24) to get authenticated first

b. Starts traffic to access resources specified by the policy-W policy

Figure 231: Web Authentication Example

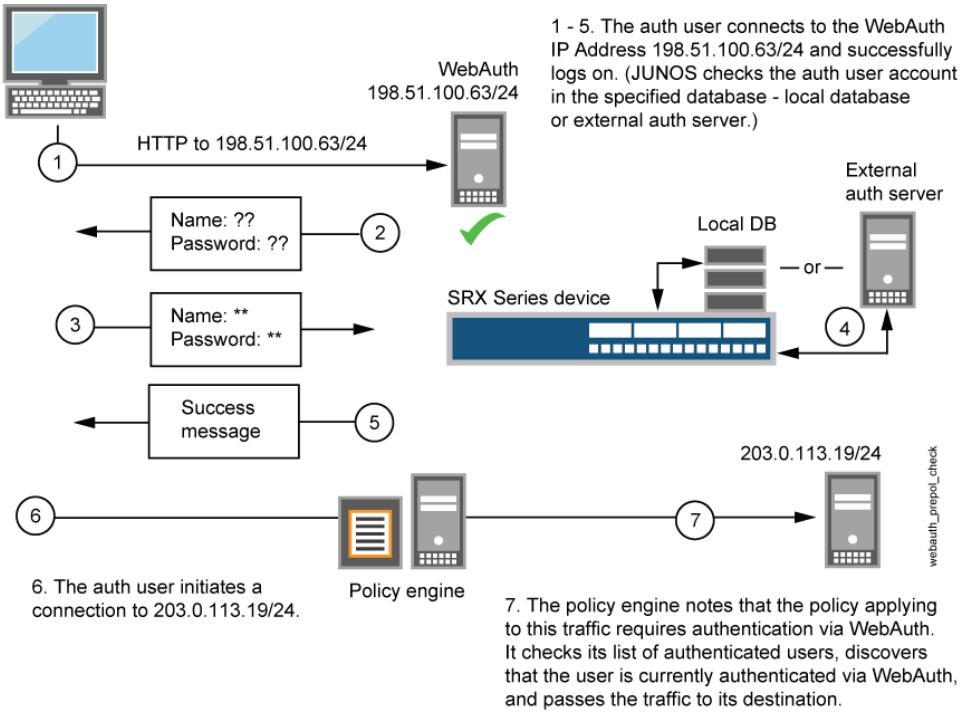

When you configure the device as described in these instructions and the user successfully authenticates, the screen illustrated in Figure 232 appears.

Figure 232: Web Authentication Success Banner

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.23/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.23/24 web-authentication http
set interfaces fe-5/0/0 unit 0 family inet address 203.0.113.15/24
set access profile WEBAUTH client FWClient1 firewall-user password pwd
set access firewall-authentication web-authentication default-profile WEBAUTH
set access firewall-authentication web-authentication banner success "WEB AUTH LOGIN SUCCESS"
set security zones security-zone UT-ZONE host-inbound-traffic system-services all
set security zones security-zone UT-ZONE interfaces ge-0/0/1.0 host-inbound-traffic protocols all
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone T-ZONE host-inbound-traffic system-services all
set security zones security-zone T-ZONE interfaces ge-5/0/0.0 host-inbound-traffic protocols
all
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match source-address any
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match destination-address
any
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 match application any
set security policies from-zone UT-ZONE to-zone T-ZONE policy P1 then permit
firewall-authentication web-authentication client-match FWClient1
set system services web-management http interface ge-0/0/1.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure Web authentication:

1. Configure two interfaces and assign IP addresses to them.

   NOTE: For this example, it is optional to assign two addresses to the interfaces.

   ```
   [edit]

   user@host# set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.23/24

   user@host# set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.63/24
     web-authentication http

   user@host# set interfaces fe-5/0/0 unit 0 family inet address 203.0.113.15/24
   ```

2. Create the WEBAUTH access profile for the FWClient1 user, specify the user's password, and define a success banner.

   ```
   [edit access]

   user@host# set profile WEBAUTH client FWClient1 firewall-user password pwd

   user@host# set firewall-authentication web-authentication default-profile
     WEBAUTH

   user@host# set firewall-authentication web-authentication banner success "WEB
     AUTH LOGIN SUCCESS"
   ```

3. Configure security zones.

   NOTE: For this example, it is optional to configure a second interface for a security zone.

   ```
   [edit security zones]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security-zone UT-ZONE host-inbound-traffic system-services all

user@host# set security-zone UT-ZONE interfaces ge-0/0/1.0 host-inbound-traffic
  protocols all

user@host# set security-zone T-ZONE host-inbound-traffic system-services all

user@host# set security-zone T-ZONE interfaces ge-5/0/0.0 host-inbound-traffic
  protocols all
```

4. Assign security policy P1 to the security zones.

```
[edit security policies]

user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match
  source-address any

user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match
  destination-address any

user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 match application
  any

user@host# set from-zone UT-ZONE to-zone T-ZONE policy P1 then permit
  firewall-authentication web-authentication client-match FWClient1
```

5. Activate the HTTP process (daemon) on your device.

```
[edit]

user@host# set system services web-management http interface ge-0/0/1.0
```

Results From configuration mode, confirm your configuration by entering these commands:

- **show interfaces**
- **show access**
- **show security zones**
- **show security policies**
- **show system services**

If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@host# show interfaces
...
}
ge-0/0/1{
    unit 0 {
        family inet {
```

 Copyright © 2016, Juniper Networks, Inc.

```
            address 198.51.100.23/24 {
            address 198.51.100.63/24 {
                web-authentication http;
            }
        }
    }
}
fe-5/0/0 {
    unit 0 {
        family inet {
            address 198.51.100.14/24;
        }
    }
}
...

user@host# show access
profile WEBAUTH {
    client FWClient1 {
        firewall-user {
            password "$ABC123"; ## SECRET-DATA
        }
    }
}
firewall-authentication {
    web-authentication {
        default-profile WEBAUTH;
        banner {
            success "WEB AUTH LOGIN SUCCESS";
        }
    }
}

user@host# show security zones
...
}
security-zone UT-ZONE {
    host-inbound-traffic {
       system-services {
            all;
        }
    }
    interfaces {
        ge-0/0/1.0 {
            host-inbound-traffic {
                protocols {
                    all;
                }
            }
        }
    }
}
security-zone T-ZONE {
    host-inbound-traffic {
       system-services {
            all;
        }
    }
    interfaces {
        ge-5/0/0.0 {
```

```
            host-inbound-traffic {
                protocols {
                    all;
                }
            }
        }
    }
}
```

```
user@host# show security policies
...
from-zone UT-ZONE to-zone T-ZONE {
    policy P1 {
       match {
            source-address any;
            destination-address any;
            application any;
       }
       then {
            permit {
                firewall-authentication {
                    web-authentication {
                        client-match FWClient1;
                    }
                }
            }
       }
    }
}
```

```
user@host# show system services
...
ftp;
ssh;
telnet;
web-management {
    http {
        interface g-0/0/1.0;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table on page 5228

Verifying Firewall User Authentication and Monitoring Users and IP Addresses in the Authentication Table

Purpose Display firewall authentication user history and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter these **show** commands:

```
user@host> show security firewall-authentication history
user@host> show security firewall-authentication history identifier 1
user@host> show security firewall-authentication users
user@host> show security firewall-authentication users identifier 3
```

```
user@host> show security firewall-authentication history
History of firewall authentication data:
Authentications: 1
Id Source Ip Date Time Duration Status User
5 198.51.100.75      2010-04-24 01:08:57 0:10:30    Success  FWClient1
```

```
user@host> show security firewall-authentication history identifier 1
Username: FWClient1
Source IP: 198.51.100.752
Authentication state: Success
Authentication method: Web-authentication
Access start date: 2010-10-12
Access start time: 21:24:02
Duration of user access: 0:00:24
Source zone: N/A
Destination zone: N/A
Access profile: WEBAUTH
Bytes sent by this user: 0
Bytes received by this user: 2660
```

```
user@host> show security firewall-authentication users
Firewall authentication data:
Total users in table: 1
Id Source Ip Src zone Dst zone Profile Age Status User
4 198.51.100.75     N/A  N/A   WEBAUTH     1 Success   FWClient1
```

```
user@host> show security firewall-authentication users identifier 3
Username: FWClient1
Source IP: 198.51.100.75
Authentication state: Success
Authentication method: Web-authentication
Age: 3
Access time remaining: 9
Source zone: N/A
Destination zone: N/A
Access profile: WEBAUTH
Interface Name: ge-0/0/1.0
Bytes sent by this user: 0
Bytes received by this user: 1521
```

Related Documentation

- Understanding Web Authentication on page 5221
- Understanding Firewall Authentication Banner Customization on page 5251
- Security Zones and Interfaces Overview on page 1049

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring HTTPS Traffic to Trigger Web Authentication

This example shows how to configure HTTPS traffic to trigger Web authentication. HTTPS traffic is used to trigger Web authentication because HTTPS is more secure than HTTP, so it has become more popular and is more

- Requirements on page 5230
- Overview on page 5231
- Configuration on page 5232
- Verification on page 5234

Requirements

Before you begin:

This example uses the following hardware and software components:

- SRX Series device
- Two PCs with Linux and Open SSL installed. One PC acts as a client and another as an HTTPS server. The two PCs are used to create key files and to send traffic.
- Junos OS Release 12.1X44-D10 or later for high-end SRX Series Services Gateways and Junos OS Release 15.1X49-D40 or later for vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways.

An SRX Series device has to decode the HTTPS traffic to trigger Web authentication. The following list describes the steps to create and install a private key file and a certification key file.

NOTE: If you have an official .crt file and .key file, then you can directly upload and install the files on the SRX Series device. If you do not have a .crt file and .key file, then follow the procedure to create and install the files. Instructions specified in Step 1 and Step 2 must be run on a PC which has Linux and OpenSSL installed. Instructions specified in Step 3 and Step 4 must be run in operational mode.

1. From the PC, create the **.key** file.

   openssl genrsa -out /tmp/device.key 1024

2. From the PC, create the **.crt** file.

   openssl req -new -x509 -days 365 -key /tmp/server.key -out /tmp/device.crt -subj "/C=CN/ST=BJ/L=BJ/O=JNPR/OU=CNRD/CN=2.2.2.10/emailAddress=device@juniper.net"

3. From the SRX Series device, upload the **.key** and **.crt** files and install the files on the device using the following command:

   user@host> **request security pki local-certificate load filename /var/tmp/device.crt key /var/tmp/device.key certificate-id device**

Copyright © 2016, Juniper Networks, Inc.

Overview

Firewall authentication initiates a secure connection to be established across two devices. A network user must provide a username and password for authentication when initiating a connection across the firewall. Firewall authentication supports HTTPS traffic for pass-through authentication. HTTPS can secure HTTP firewall authentication traffic between users and the SRX Series device.

HTTPS is the secure version of HTTP, the protocol over which data is sent between the user and the device that the user is connected to. All communications between the user and the connected devices are encrypted. HTTPS is often used to protect highly confidential online transactions like online banking and online shopping order forms.

In this example, HTTPS traffic is used to trigger Web authentication because HTTPS is more secure than HTTP.

The user uses HTTPS to access an IP address on the device that is enabled for Web authentication. In this scenario, the user does not use HTTPS to access the IP address of the protected resource. The user is prompted for a username and password, which are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the results of this Web authentication.

Figure 233 shows an example of Web authentication using HTTPS traffic.

Figure 233: Web Authentication Using HTTPS Traffic

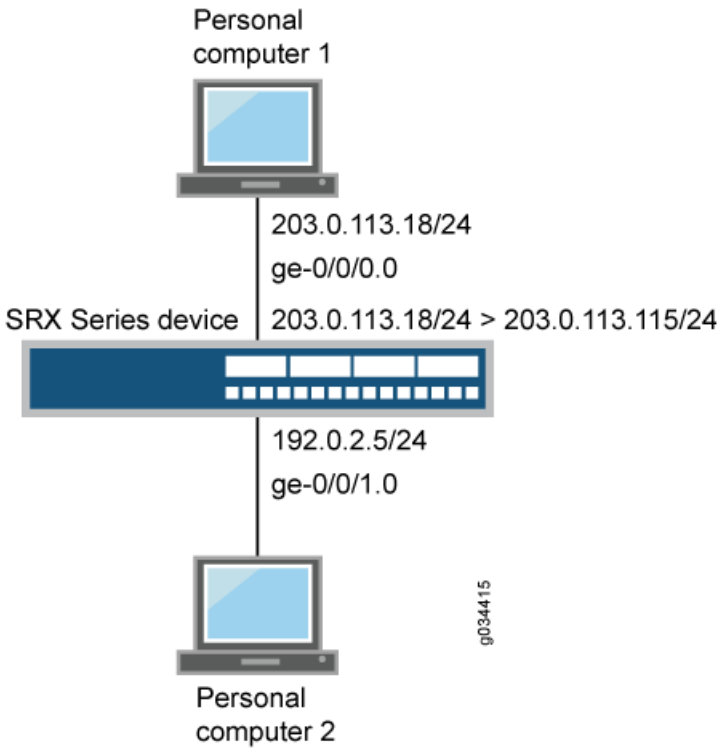

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set system services web-management https pki-local-certificate device
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/24
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.11/24 web-authentication https
set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
set security policies from-zone trust to-zone untrust policy p1 match source-address any
set security policies from-zone trust to-zone untrust policy p1 match destination-address
  any
set security policies from-zone trust to-zone untrust policy p1 match application any
set security policies from-zone trust to-zone untrust policy p1 then permit
set access profile local_pf client user1 firewall-user password user1
set access firewall-authentication web-authentication default-profile local_pf
set security policies from-zone trust to-zone untrust policy p1 then permit
  firewall-authentication web-authentication
```

Step-by-Step Procedure

To configure HTTPS traffic to trigger Web authentication:

1. Enable Web-management support to HTTPS traffic.

   ```
   [edit system services]
   user@host# set web-management https pki-local-certificate device
   ```

2. Configure interfaces and assign IP addresses. Enable Web authentication at ge-0/0/0 interface.

   ```
   [edit interfaces]
   user@host# set ge-0/0/0 unit 0 family inet address 1.1.1.1/24
   set ge-0/0/0 unit 0 family inet address 1.1.1.11/24 web-authentication https
   user@host# set ge-0/0/1 unit 0 family inet address 2.2.2.1/24
   ```

3. Configure security policies to permit firewall authenticated traffic from zone trust to zone untrust.

   ```
   [edit security policies]
   user@host# set from-zone trust to-zone untrust policy p1 match source-address
     any destination-address any application any
   user@host# set security policies from-zone trust to-zone untrust policy p1 then
     permit
   ```

4. Create an access profile, configure the client as a firewall user, and set the password.

   ```
   [edit access]
   user@host# set profile local_pf client user1 firewall-user password user1
   ```

5. Configure the type of firewall authentication settings.

   ```
   [edit access]
   user@host# set firewall-authentication web-authentication default-profile local_pf
   ```

6. Specify a policy action to take when a packet matches the criteria.

   ```
   [edit security policies]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set from-zone trust to-zone untrust policy p1 then permit
  firewall-authentication web-authentication
```

Results From configuration mode, confirm your configuration by entering the **show system services**, **show interfaces**, **show security policies**, and **show access** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show system services
  web-management {
  https {
  pki-local-certificate device;
  }
}
user@host# show interfaces
  ge-0/0/0 {
  unit 0 {
    family inet {
      address 1.1.1.11/24 {
        web-authentication https;
      }
    }
}
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 2.2.2.1/24;
      }
    }
  }
user@host# show security policies
from-zone trust to-zone untrust {
  policy p1 {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit {
        firewall-authentication {
          web-authentication;
        }
      }
    }
  }
}
user@host# show access
  profile local_pf {
    client user1 {
      firewall-user {
        password "user1";
      }
```

```
    }
  }
    firewall-authentication {
      web-authentication {
        default-profile local_pf;
      }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration

Purpose Verify that the configuration is correct.

Action From operational mode, enter the **show security firewall-authentication users identifier *identifier*** command.

Sample Output

```
user@host> show security firewall-authentication users identifier 1
  Username: user1
  Source IP: 1.1.1.10
  Authentication state: Success
  Authentication method: Web-authentication
  Age: 0
  Access time remaining: 10
  Lsys: root-logical-system
  Source zone: N/A
  Destination zone: N/A
  Access profile: local_pf
  Bytes sent by this user: 0
  Bytes received by this user: 0
```

Meaning The **show security firewall-authentication users identifier *identifier*** command displays the firewall authentication user information using the identifier ID of the user. If the authentication method parameter displays Web authentication and the authentication state parameter displays success in your output then your configuration is correct.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Understanding Web Authentication on page 5221
- Example: Configuring Web Authentication on page 5223

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 216

Configuring External Authentication Servers

- Understanding External Authentication Servers on page 5235
- Example: Configuring RADIUS and LDAP User Authentication on page 5236
- Example: Configuring SecurID User Authentication on page 5241
- Example: Deleting the SecurID Node Secret File on page 5244

Understanding External Authentication Servers

Authentication, authorization, and accounting (AAA) servers provide an extra level of protection and control for user access in the following ways:

- Authentication determines the firewall user.
- Authorization determines what the firewall user can do.
- Accounting determines what the firewall user did on the network.

You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

Once the user's credentials are collected, they are processed using firewall user authentication, which supports the following types of servers:

- Local authentication and authorization
- RADIUS authentication and authorization (compatible with Juniper Steel-Belted Radius server)
- LDAP authentication only (supports LDAP version 3 and is compatible with Windows AD)
- SecurID authentication only (using an RSA SecurID external authentication server)

NOTE: Junos OS also supports administrative authentication using local, RADIUS, and TACACS+ servers.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Understanding SecurID User Authentication on page 5236

Understanding SecurID User Authentication

SecurID is an authentication method that allows users to enter either static or dynamic passwords as their credentials. A dynamic password is a combination of a user's PIN and a randomly generated token that is valid for a short period of time, approximately one minute. A static password is set for the user on the SecurID server. For example, the SecurID server administrator might set a temporary static password for a user who lost his or her SecurID token.

When a user attempts to access a resource protected by a policy and SecurID is configured in the profile **authentication-order** parameter as either the only authentication mode or the first one to be used, the device forwards the user's credentials to the SecurID server for authentication. If the user enters valid values, the user is allowed access to the requested resource.

NOTE: **The SecurID server includes a feature that presents a user with a challenge if the user provides wrong credentials repeatedly. However, Junos OS does not support the challenge feature. Instead, the SecurID server administrator must resynchronize the RSA token for the user.**

For SecurID, you configure information about the Juniper Networks device on the SecurID server, and this information is exported to a file called sdconf.rec.

To install the sdconf.rec file on the device, you must use an out-of-band method such as FTP. Install the file in a directory whose files are not deleted regularly. Do not put it in a temporary directory. For example, you might install it in /var/db/secureid/server1/sdconf.rec.

The sdconf.rec file contains information that provides the Juniper Networks device with the address of the SecurID server. You do not need to configure this information explicitly when you configure the SecurID server to be used as the external authentication server.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Example: Configuring RADIUS and LDAP User Authentication on page 5236
- Example: Configuring SecurID User Authentication on page 5241
- Example: Deleting the SecurID Node Secret File on page 5244

Example: Configuring RADIUS and LDAP User Authentication

This example shows how to configure a device for external authentication.

- Requirements on page 5237
- Overview on page 5237

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 5237
- Verification on page 5240

Requirements

Before you begin, create an authentication user group.

Overview

You can put several user accounts together to form a user group, which you can store on the local database or on a RADIUS, an LDAP, or a SecurID server. When you reference an authentication user group and an external authentication server in a policy, the traffic matching the policy provokes an authentication check.

This example shows how access profile Profile-1 is configured for external authentication. Two RADIUS servers and one LDAP server are configured in the access profile. However, the order of authentication specifies RADIUS server only, so if the RADIUS server authentication fails, then the firewall user fails to authenticate. The local database is not accessed.

NOTE: If the firewall clients are authenticated by the RADIUS server, then the group-membership VSA returned by the RADIUS server should contain alpha, beta, or gamma client groups in the RADIUS server configuration or in the access profile, Profile-1. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set access profile Profile-1 authentication-order radius
set access profile Profile-1 client Client-1 client-group alpha
set access profile Profile-1 client Client-1 client-group beta
set access profile Profile-1 client Client-1 client-group gamma
set access profile Profile-1 client Client-1 firewall-user password pwd
set access profile Profile-1 client Client-2 client-group alpha
set access profile Profile-1 client Client-2 client-group beta
set access profile Profile-1 client Client-2 firewall-user password pwd
set access profile Profile-1 client Client-3 firewall-user password pwd
set access profile Profile-1 client Client-4 firewall-user password pwd
set access profile Profile-1 session-options client-group alpha
set access profile Profile-1 session-options client-group beta
set access profile Profile-1 session-options client-group gamma
set access profile Profile-1 session-options client-idle-timeout 255
set access profile Profile-1 session-options client-session-timeout 4
set access profile Profile-1 ldap-options base-distinguished-name
CN=users,DC=junos,DC=juniper,DC=net
set access profile Profile-1 ldap-options search search-filter sAMAccountName=
```

Copyright © 2016, Juniper Networks, Inc.

```
set access profile Profile-1 ldap-options search admin-search distinguished-name
cn=administrator,cn=users,dc=junos,dc=juniper,dc=net
set access profile Profile-1 ldap-options search admin-search password juniper
set access profile Profile-1 ldap-server 3.3.3.3
set access profile Profile-1 radius-server 4.4.4.4 secret juniper
set access profile Profile-1 radius-server 4.4.4.4 retry 10
set access profile Profile-1 radius-server 5.5.5.5 secret juniper
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a device for external authentication:

1. Specify the RADIUS server for external authentication order.

   ```
   [edit]

   user@host# set access profile Profile-1 authentication-order radius
   ```

2. Configure Client1-4 firewall users and assign the Client-1 firewall user and Client-2 firewall user to client groups.

   ```
   [edit access profile Profile-1]

   user@host# set client Client-1 client-group alpha

   user@host# set client Client-1 client-group beta

   user@host# set client Client-1 client-group gamma

   user@host# set client Client-1 firewall-user password pwd

   user@host# set client Client-2 client-group alpha

   user@host# set client Client-2 client-group beta

   user@host# set client Client-2 firewall-user password pwd

   user@host# set client Client-3 firewall-user password pwd

   user@host# set client Client-4 firewall-user password pwd
   ```

3. Configure client groups in the session options.

   ```
   [edit access profile Profile-1]

   user@host# set session-options client-group alpha

   user@host# set session-options client-group beta

   user@host# set session-options client-group gamma
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set session-options client-idle-timeout 255

user@host# set session-options client-session-timeout 4
```

4. Configure the IP address for the LDAP server and server options.

```
[edit access profile Profile-1]

user@host# set ldap-options base-distinguished-name
  CN=users,DC=junos,DC=juniper,DC=net

user@host# set ldap-options search search-filter sAMAccountName=

user@host# set ldap-options search admin-search password juniper

user@host# set ldap-options search admin-search distinguished-name
  cn=administrator,cn=users,dc=junos,dc=juniper,dc=net

user@host# set ldap-server 3.3.3.3
```

5. Configure the IP addresses for the two RADIUS servers.

```
[edit access profile Profile-1]

user@host# set radius-server 4.4.4.4 secret juniper

user@host# set radius-server 4.4.4.4 retry 10

user@host# set radius-server 5.5.5.5 secret juniper
```

Results From configuration mode, confirm your configuration by entering the **show access profile Profile-1** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show access profile Profile-1
authentication-order radius;
client Client-1 {
    client-group [ alpha beta gamma ];
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-2 {
    client-group [ alpha beta ];
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-3 {
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-4 {
    firewall-user {
```

Copyright © 2016, Juniper Networks, Inc.

```
        password "$ABC123"; ## SECRET-DATA
    }
}
session-options {
    client-group [ alpha beta gamma ];
    client-idle-timeout 255;
    client-session-timeout 4;
}
ldap-options {
    base-distinguished-name CN=users,DC=junos,DC=juniper,DC=net;
    search {
        search-filter sAMAccountName=;
        admin-search {
            distinguished-name
cn=administrator,cn=users,dc=junos,dc=juniper,dc=net;
            password "$ABC123"; ## SECRET-DATA
        }
    }
}
ldap-server {
    3.3.3.3;
}
radius-server {
    4.4.4.4 {
        secret "$ABC123"; ## SECRET-DATA
        retry 10;
    }
    5.5.5.5 {
        secret "$ABC123"; ## SECRET-DATA
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Troubleshooting with Logs on page 5240

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation

- Understanding External Authentication Servers on page 5235
- Example: Configuring SecurID User Authentication on page 5241
- Example: Deleting the SecurID Node Secret File on page 5244

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring SecurID User Authentication

This example shows how to configure SecurID as the external authentication server.

- Requirements on page 5241
- Overview on page 5241
- Configuration on page 5241
- Verification on page 5243
- Troubleshooting on page 5244

Requirements

Before you begin, create an authentication user group.

Overview

SecurID is an authentication method that allows users to enter either static or dynamic passwords as their credentials. A dynamic password is a combination of a user's PIN and a randomly generated token that is valid for a short period of time, approximately one minute. A static password is set for the user on the SecurID server. For example, the SecurID server administrator might set a temporary static password for a user who lost his or her SecurID token.

When a user attempts to access a resource protected by a policy and SecurID is configured in the profile **authentication-order** parameter as either the only authentication mode or the first one to be used, the device forwards the user's credentials to the SecurID server for authentication. If the user enters valid values, the user is allowed access to the requested resource.

Specify that Server-1 is to be used as the SecurID server and that its configuration file resides on the device in the **/var/db/securid/Server-1/sdconf.rec** file. From configuration mode, enter this command:

```
user@host# set access securid-server Server-1 configuration-file
  “/var/db/securid/Server-1/sdconf.rec”
```

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set access profile Profile-2 authentication-order securid
set access profile Profile-2 client Client-1 client-group alpha
set access profile Profile-2 client Client-1 client-group beta
set access profile Profile-2 client Client-1 client-group gamma
set access profile Profile-2 client Client-1 firewall-user password pwd
set access profile Profile-2 client Client-2 client-group alpha
set access profile Profile-2 client Client-2 client-group beta
set access profile Profile-2 client Client-2 firewall-user password pwd
set access profile Profile-2 client Client-3 firewall-user password pwd
```

Copyright © 2016, Juniper Networks, Inc.

```
set access profile Profile-2 client Client-4 firewall-user password pwd
set access profile Profile-2 session-options client-group alpha
set access profile Profile-2 session-options client-group beta
set access profile Profile-2 session-options client-group gamma
set access profile Profile-2 session-options client-idle-timeout 255
set access profile Profile-2 session-options client-session-timeout 4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure SecurID as the external authentication server:

1. For the Profile-2 profile, configure SecurID as the server to be used for external authentication.

   ```
   [edit]

   user@host# set access profile Profile-2 authentication-order securid
   ```

   To share a single SecurID server across multiple profiles, for each profile set the **authentication-order** parameter to include **securid** as the authentication mode.

2. Configure clients 1 through 4 as firewall users, and assign Client-1 and Client-2 to client groups.

   ```
   [edit access profile Profile-2]

   user@host# set client Client-1 client-group alpha

   user@host# set client Client-1 client-group beta

   user@host# set client Client-1 client-group gamma

   user@host# set client Client-1 firewall-user password pwd

   user@host# set client Client-2 client-group alpha

   user@host# set client Client-2 client-group beta

   user@host# set client Client-2 firewall-user password pwd

   user@host# set client Client-3 firewall-user password pwd

   user@host# set client Client-4 firewall-user password pwd
   ```

3. Configure client groups in the session options.

   ```
   [edit access profile Profile-2]

   user@host# set session-options client-group alpha

   user@host# set session-options client-group beta
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set session-options client-group gamma

user@host# set session-options client-idle-timeout 255

user@host# set session-options client-session-timeout 4
```

Results From configuration mode, confirm your configuration by entering the **show access profile Profile-2** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show access profile Profile-2
authentication-order securid;
client Client-1 {
    client-group [ alpha beta gamma ];
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-2 {
    client-group [ alpha beta ];
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-3 {
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
client Client-4 {
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
session-options {
    client-group [alpha beta gamma];
    client-idle-timeout 255;
    client-session-timeout 4;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Troubleshooting with Logs on page 5243

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Copyright © 2016, Juniper Networks, Inc.

Troubleshooting

- Troubleshooting Unsuccessful Authentication In a Dynamic VPN Configuration on page 5244

Troubleshooting Unsuccessful Authentication In a Dynamic VPN Configuration

Problem Device fails to locate client address in a dynamic VPN configuration.

Solution

1. Verify that the device host name, the domain-search, and the name server are configured properly.

   ```
   [edit system]

   user@host# set host-name srx101.uaclab.net

   user@host# set domain-search uaclab.net

   user@host# set name-server 10.204.91.25
   ```

2. Verify that the device host name is getting resolved on the RSA server.

Related Documentation

- Understanding External Authentication Servers on page 5235
- Example: Deleting the SecurID Node Secret File on page 5244

Example: Deleting the SecurID Node Secret File

This example shows how to delete the node secret file.

- Requirements on page 5244
- Overview on page 5244
- Configuration on page 5245
- Verification on page 5245

Requirements

Before you begin, confirm that it is necessary to delete the SecurID node secret file.

Overview

When the Juniper Networks device initially communicates successfully with the SecurID server, a node secret file is created for it automatically. The file is created as a result of the handshake between the Juniper Networks device and the SecurID server after the software authenticates the first user successfully. All subsequent communication between the Juniper Networks device and the SecurID server relies on this secret as a representation of trust between the two nodes instead of repeating the handshake with each authentication request.

Copyright © 2016, Juniper Networks, Inc.

Under normal circumstances you should not delete the node secret file. In the rare case that you must do so, for example, to debug a serious problem, you can use the **clear** command to remove the file.

WARNING: If you delete the file, you must deselect a box on the SecurID server to indicate that the node secret file for the Juniper Networks device and the SecurID server no longer exists. Otherwise, authentication attempts will fail.

Configuration

Step-by-Step Procedure

To delete the node secret file:

1. Use the **clear** command to remove the node secret file. During subsequent user authentication, the device reestablishes a shared secret with the SecurID server and re-creates the node secret file. From operational mode, enter the **clear network-access** command to clear the **securid-node-secret-file** for the Juniper Networks device.

   ```
   user@host> clear network-access securid-node-secret-file
   ```

2. From operational mode, confirm your deletion by entering the **show network-access securid-node-secret-file** command. If the output does not display, repeat the instructions in this example to correct it.

   ```
   user@host> show network-access securid-node-secret-file
   ```

Verification

Verify the deletion by entering the **show network-access securid-node-secret-file** command.

Related Documentation

- Understanding External Authentication Servers on page 5235
- Example: Configuring SecurID User Authentication on page 5241

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 217

Configuring Client Groups

- Understanding Client Groups for Firewall Authentication on page 5247
- Example: Configuring the Access Profile on page 5247
- Example: Configuring Local Users for Client Groups on page 5248

Understanding Client Groups for Firewall Authentication

To manage a number of firewall users, you can create user or client groups and store the information either on the local Juniper Networks device or on an external RADIUS or LDAP server.

A client group is a list of groups to which the client belongs. As with client-idle timeout, a client group is used only if the external authentication server does not return a value in its response. (For example, LDAP servers do not return such information.)

The RADIUS server sends the client's group information to the Juniper Networks device using Juniper VSA (46). The client-match portion of the policy accepts a string that can be either the username or the groupname to which the client belongs.

The reason to have a single database for different types of clients (except admins) is based on the assumption that a single client can be of multiple types. For example, a firewall user client can also be an L2TP client.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Example: Configuring Local Users for Client Groups on page 5248

Example: Configuring the Access Profile

The following example shows you how to configure the access profile:

```
[edit access]
profile westcoast_bldg_1 {
   client white {
      chap-secret "$ABC123";
      # SECRET-DATA
      ppp {
         idle-timeout 22;
         primary-dns 192.120.65.10;
```

Copyright © 2016, Juniper Networks, Inc.

```
            framed-ip-address 12.12.12.12/32;
          }
          group-profile westcoast_users;
        }
        client blue {
          chap-secret "$ABC123";
          # SECRET-DATA
          group-profile sunnyvale_users;
        }
        authentication-order password;
      }
      profile westcoast_bldg_1_tunnel {
        client test {
          l2tp {
            shared-secret "$ABC123";
            # SECRET-DATA
            maximum-sessions-per-tunnel 75;
            ppp-authentication chap;
          }
          group-profile westcoast_tunnel;
        }
        client production {
          l2tp {
            shared-secret "$ABC123"
            # SECRET-DATA
            ppp-authentication chap;
          }
          group-profile westcoast_tunnel;
        }
      }
```

Related Documentation

- Configuring Access Profiles for L2TP or PPP Parameters

Example: Configuring Local Users for Client Groups

This example shows how to configure a local user for client groups in a profile.

- Requirements on page 5248
- Overview on page 5248
- Configuration on page 5249
- Verification on page 5250

Requirements

Before you begin, create an access profile. See "Example: Configuring the Access Profile" on page 5247.

Overview

A client group is a list of groups to which the client belongs. As with client-idle timeout, a client group is used only if the external authentication server does not return a value in its response (for example, LDAP servers do not return such information).

Copyright © 2016, Juniper Networks, Inc.

This example shows how to configure a local user called Client-1 for client groups G1, G2, and G3 in a profile called Managers. Within this example, client groups are configured for a client. If a client group is not defined for the client, then the client group under the **access profile session-options** hierarchy is used.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set access profile Managers client Client-1 client-group G1
set access profile Managers client Client-1 client-group G2
set access profile Managers client Client-1 client-group G3
set access profile Managers client Client-1 firewall-user password pwd
set access profile Managers session-options client-group G1
set access profile Managers session-options client-group G2
set access profile Managers session-options client-group G3
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a local user for client groups in a profile:

1. Configure the firewall user profile Managers, and assign client groups to it.

```
user@host# edit access profile Managers
[edit access profile Managers]
user@host# set client Client-1 client-group G1
user@host# set client Client-1 client-group G2
user@host# set client Client-1 client-group G3
user@host# set client Client-1 firewall-user password pwd
```

2. Configure client groups in the session options.

```
[edit access profile Managers]
user@host# set session-options client-group G1
user@host# set session-options client-group G2
user@host# set session-options client-group G3
```

Results

Confirm your configuration by entering the **show access profile Managers** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show access profile Managers

client Client-1 {
    client-group [ G1 G2 G3 ];
    firewall-user {
        password "$ABC123"; ## SECRET-DATA
    }
}
session-options {
    client-group [ G1 G2 G3 ];
}
```

Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Troubleshooting with Logs on page 5250

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation

- Understanding Client Groups for Firewall Authentication on page 5247

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 218

Customizing the Firewall Authentication Banner

- Understanding Firewall Authentication Banner Customization on page 5251
- Example: Customizing a Firewall Authentication Banner on page 5251

Understanding Firewall Authentication Banner Customization

A banner is a message that appears on a monitor in different places depending on the type of login.

Figure 234: Banner Customization

- At the top of a browser screen after a user has successfully logged into a Web authentication address as shown Figure 234.
- Before or after a Telnet, an FTP, an HTTP, or and HTTPS login prompt, success message, and fail message for users

All banners, except for a console login banner, have default messages. You can customize the messages that appear on the banners to better suit the network environment in which you use the device.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Example: Customizing a Firewall Authentication Banner on page 5251

Example: Customizing a Firewall Authentication Banner

This example shows how to customize the banner text that appears in the browser.

- Requirements on page 5252
- Overview on page 5252
- Configuration on page 5252
- Verification on page 5253

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, create an access profile.

Overview

A banner is a message that appears on a monitor in different places depending on the type of login. This example shows how to change the banner that appears in the browser to indicate that a user has successfully authenticated after successfully logging in through Web authentication. The new message is “Web authentication is successful.” If the authentication fails, then the new message reads “Authentication failed.”

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set access firewall-authentication pass-through default-profile Profile-1
set access firewall-authentication pass-through ftp banner fail “ Authentication failed”
set access firewall-authentication web-authentication default-profile Profile-1
set access firewall-authentication web-authentication banner success “ Web authentication is successful”
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To customize the banner text that appears in the browser:

1. Specify the banner text for failed pass-through authentication through FTP.

   ```
   [edit]

   user@host# set access firewall-authentication pass-through default-profile Profile-1

   user@host# set access firewall-authentication pass-through ftp banner fail “
     Authentication failed”
   ```

2. Specify the banner text for successful Web authentication.

   ```
   [edit]

   user@host# set access web-authentication default-profile Profile-1

   user@host# set access web-authentication banner success “ Web authentication
     is successful”
   ```

Results

From configuration mode, confirm your configuration by entering the **show access firewall-authentication** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show access firewall-authentication
```


Copyright © 2016, Juniper Networks, Inc.

```
pass-through {
    default-profile Profile-1;
    ftp {
        banner {
            fail "Authentication failed";
        }
    }
}
web-authentication {
    default-profile Profile-1;
    banner {
        success "Web authentication is successful";
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Troubleshooting with Logs on page 5253

Troubleshooting with Logs

Purpose Use these logs to identify any issues.

Action From operational mode, enter the **show log messages** command and the **show log dcd** command.

Related Documentation

- Understanding Firewall Authentication Banner Customization on page 5251

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 68

Configuring Infranet Authentication

- Configuring UAC in a Junos OS Environment on page 5257
- Establishing Communications Between Devices on page 5261
- Configuring Policy Enforcement on page 5265
- Classifying Traffic with User Roles on page 5269
- Configuring Endpoint Security on page 5287
- Configuring IPsec on page 5289
- Configuring Captive Portal on page 5299

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 219

Configuring UAC in a Junos OS Environment

- Understanding UAC in a Junos OS Environment on page 5257
- Enabling UAC in a Junos OS Environment (CLI Procedure) on page 5259

Understanding UAC in a Junos OS Environment

A Unified Access Control (UAC) deployment uses the following components to secure a network and ensure that only qualified end users can access protected resources:

- IC Series UAC Appliances—An IC Series appliance is a policy decision point in the network. It uses authentication information and policy rules to determine whether or not to provide access to individual resources on the network. You can deploy one or more IC Series appliances in your network.
- Infranet Enforcers—An Infranet Enforcer is a policy enforcement point in the network. It receives policies from the IC Series appliance and uses the rules defined in those policies to determine whether or not to allow an endpoint access to a resource. You deploy the Infranet Enforcers in front of the servers and resources that you want to protect.
- Infranet agents—An Infranet agent is a client-side component that runs directly on network endpoints (such as users' computers). The agent checks that the endpoint complies to the security criteria specified in Host Checker policies and relays that compliance information to the Infranet Enforcer. The Infranet Enforcer then allows or denies the endpoint access based on the compliance results.

An SRX Series device can act as an Infranet Enforcer in a UAC network. Specifically, it acts as a Layer 3 enforcement point, controlling access by using IP-based policies pushed down from the IC Series appliance. When deployed in a UAC network, an SRX Series device is called a Junos OS Enforcer. See Figure 235.

Copyright © 2016, Juniper Networks, Inc.

Figure 235: Integrating a Junos OS Security Device into a Unified Access Control Network

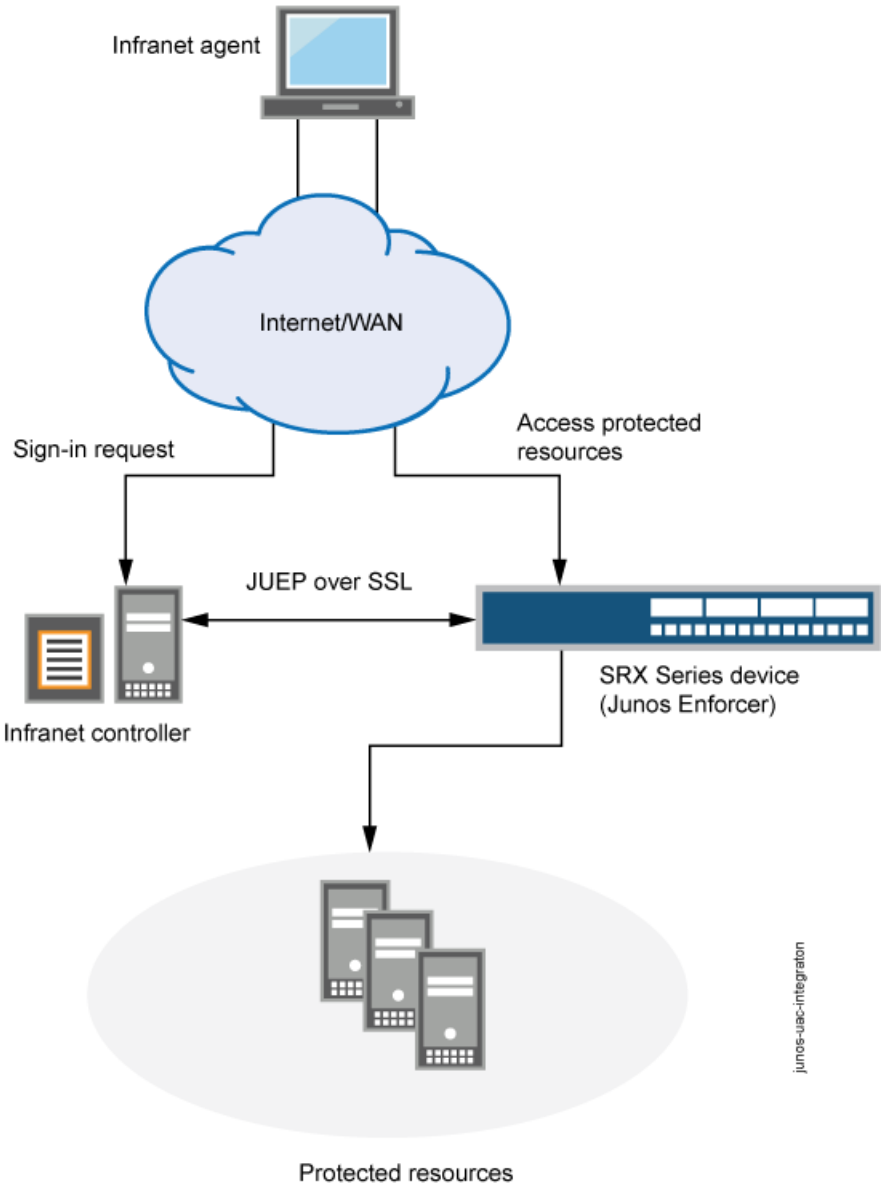

NOTE: You can use the Junos OS Enforcer with the IC Series appliance and Secure Access devices in an IF-MAP Federation network. In a federated network, multiple IC Series appliances and Secure Access devices that are not directly connected to the Junos OS Enforcer can access resources protected by the security device. There are no configuration tasks for IF-MAP Federation on the Junos OS Enforcer. You configure policies on IC Series appliances that can dynamically create authentication table entries on the Junos OS Enforcer.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Enabling UAC in a Junos OS Environment (CLI Procedure) on page 5259

Enabling UAC in a Junos OS Environment (CLI Procedure)

Junos OS security policies enforce rules for transit traffic, defining what traffic can pass through the Juniper Networks device. The policies control traffic that enters from one zone (from-zone) and exits another (to-zone). To enable an SRX Series device as a Junos OS Enforcer in a UAC deployment, you must:

- Identify the source and destination zones through which UAC traffic will travel. It also needs the list of interfaces, including which zones they are in. The IC Series UAC Appliance uses the destination zone to match its own IPsec routing policies configured on IC Series appliance.
- Identify Junos OS security policies that encompass those zones, and enable UAC for those policies.

Before you begin:

1. Set up the interfaces through which UAC traffic should enter the SRX Series device.
2. Group interfaces with identical security requirements into zones. See "Example: Creating Security Zones" on page 1051.
3. Create security policies to control the traffic that passes through the security zones. See "Example: Configuring a Security Policy to Permit or Deny All Traffic" on page 1094.

To configure UAC through a Junos OS security policy, enter the following configuration statement:

```
user@host# set security policies from-zone zone-name to-zone zone-name policy match
    then permit application-services uac-policy
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 220

Establishing Communications Between Devices

- Understanding Communications Between the Junos OS Enforcer and the IC Series UAC Appliance on page 5261
- Understanding Communications Between Junos OS Enforcer and a Cluster of IC Series UAC Appliances on page 5262
- Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure) on page 5262

Understanding Communications Between the Junos OS Enforcer and the IC Series UAC Appliance

When you configure an SRX Series device to connect to an IC Series UAC Appliance, the SRX Series device and the IC Series appliance establish secure communications as follows:

1. If more than one IC Series device are configured as Infranet Controllers on the SRX Series device, a round-robin algorithm determines which of the configured IC Series devices is the active Infranet Controller. The others are failover devices. If the active Infranet Controller becomes inoperative, the algorithm is reapplied to the remaining IC Series devices that are configured to establish the new active Infranet Controller.
2. The active IC Series appliance presents its server certificate to the SRX Series device. If configured to do so, the SRX Series device verifies the certificate. (Server certificate verification is not required; however, as an extra security measure you can verify the certificate to implement an additional layer of trust.)
3. The SRX Series device and the IC Series appliance perform mutual authentication using the proprietary challenge-response authentication. For security reasons, the password is not included in the message sent to the IC Series appliance.
4. After successfully authenticating with the SRX Series device, the IC Series appliance sends its user authentication and resource access policy information. The SRX Series device uses this information to act as the Junos OS Enforcer in the UAC network.
5. Thereafter, the IC Series appliance and the Junos OS Enforcer can communicate freely with one another over the SSL connection. The communications are controlled by a proprietary protocol called *Junos UAC Enforcer Protocol (JUEP)*.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding UAC in a Junos OS Environment on page 5257
- Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure) on page 5262

Understanding Communications Between Junos OS Enforcer and a Cluster of IC Series UAC Appliances

You can configure a Junos OS Enforcer to work with more than one IC Series UAC Appliance in a high availability configuration known as an IC Series appliance cluster. The Junos OS Enforcer communicates with only one IC Series appliance at a time; the other IC Series appliances are used for failover. If the Junos OS Enforcer cannot connect to the first IC Series appliance you added to a cluster, it tries to connect to the failed IC Series appliance again. Then it fails over to the other IC Series appliances in the cluster. It continues trying to connect to IC Series appliances in the cluster until a connection occurs.

When the Junos OS Enforcer cannot establish a connection to an Infranet Enforcer, it preserves all its existing authentication table entries and Unified Access Control (UAC) policies and takes the timeout action that you specify. Timeout actions include:

- **close**—Close existing sessions and block any further traffic. This is the default option.
- **no-change**—Preserve existing sessions and require authentication for new sessions.
- **open**—Preserve existing sessions and allow new sessions access.

Once the Junos OS Enforcer can reestablish a connection to an IC Series appliance, the IC Series appliance compares the authentication table entries and UAC policies stored on the Junos OS Enforcer with the authentication table entries and policies stored on the IC Series appliance and reconciles the two as required.

NOTE: The IC Series appliances configured on a Junos OS Enforcer should all be members of the same IC Series appliance cluster.

Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure)

To configure an SRX Series device to act as a Junos OS Enforcer in a UAC deployment, and therefore to enforce IC Series UAC Appliance policies, you must specify an IC Series appliance to which the SRX Series device should connect.

Copyright © 2016, Juniper Networks, Inc.

Before you begin:

1. Enable UAC through the relevant Junos OS security policies. See "Enabling UAC in a Junos OS Environment (CLI Procedure)" on page 5259.
2. (Optional) Create a profile for the certificate authority (CA) that signed the IC Series appliance's server certificate, and import the CA certificate onto the SRX Series device. See "Example: Loading CA and Local Certificates Manually" on page 6378.
3. Configure user authentication and authorization by setting up user roles, authentication and authorization servers, and authentication realms on the IC Series appliance.
4. Configure resource access policies on the IC Series appliance to specify which endpoints are allowed or denied access to protected resources.

To configure an SRX Series device to act as a Junos OS Enforcer:

1. Specify the IC Series appliance(s) to which the SRX Series device should connect.
   - To specify the IC Series appliance hostname:

     user@host# **set services unified-access-control infranet-controller *hostname***
   - To specify the IC Series appliance IP address:

     user@host# **set services unified-access-control infranet-controller *hostname* address *ip-address***

   NOTE: When configuring access to multiple IC Series appliances, you must define each separately. For example:

   ```
   user@host# set services unified-access-control infranet-controller IC1
   user@host# set services unified-access-control infranet-controller IC2
   user@host# set services unified-access-control infranet-controller IC3

   user@host# set services unified-access-control infranet-controller IC1
     address 10.10.10.1
   user@host# set services unified-access-control infranet-controller IC2
     address 10.10.10.2
   user@host# set services unified-access-control infranet-controller IC3
     address 10.10.10.3
   ```

   Make sure that all of the IC Series appliances are members of the same cluster.

   NOTE: By default, the IC Series appliance should select port 11123.

2. Specify the Junos OS interface to which the IC Series appliance should connect:

   user@host# **set services unified-access-control infranet-controller *hostname* interface *interface-name***

Copyright © 2016, Juniper Networks, Inc.

3. Specify the password that the SRX Series device should use to initiate secure communications with the IC Series appliance:

NOTE: Any change in the Unified Access Control's (UAC) contact interval and timeout values in the SRX Series device will be effective only after the next reconnection of the SRX Series device with the IC Series appliance.

user@host# **set services unified-access-control infranet-controller** *hostname* **password** *password*

4. (Optional) Specify information about the IC Series appliance's server certificate that the SRX Series device needs to verify the certificate.

   - To specify the server certificate subject that the SRX Series device checks:

     user@host# **set services unified-access-control infranet-controller** *hostname* **server-certificate-subject** *certificate-name*

   - To specify the CA profile associated with the certificate:

     user@host# **set services unified-access-control infranet-controller** *hostname* **ca-profile** *ca-profile*

NOTE: An IC Series appliance server certificate can be issued by an intermediate CA. There are two types of CAs—root CAs and intermediate CAs. An intermediate CA is secondary to a root CA and issues certificates to other CAs in the public key infrastructure (PKI) hierarchy. Therefore, if a certificate is issued by an intermediate CA, you need to specify the complete list of CA profiles in the certification chain.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 221

Configuring Policy Enforcement

- Understanding Junos OS Enforcer Policy Enforcement on page 5265
- Configuring Junos OS Enforcer Failover Options (CLI Procedure) on page 5266
- Testing Junos OS Enforcer Policy Access Decisions Using Test-Only Mode (CLI Procedure) on page 5267
- Verifying Junos OS Enforcer Policy Enforcement on page 5268

Understanding Junos OS Enforcer Policy Enforcement

Once the SRX Series device has successfully established itself as the Junos OS Enforcer, it secures traffic as follows:

1. First, the Junos OS Enforcer uses the appropriate Junos OS security policy to process the traffic. A *security policy* uses criteria such as the traffic's source IP address or the time of day that the traffic was received to determine whether or not the traffic should be allowed to pass.
2. Once it determines that the traffic may pass based on the Junos OS security policy, the Junos OS Enforcer maps the traffic flow to an authentication table entry. The Junos OS Enforcer uses the source IP address of the first packet in the flow to create the mapping.

   An *authentication table entry* contains the source IP address and user role(s) of a user who has already successfully established a UAC session. A *user role* identifies a group of users based on criteria such as type (for instance, "Engineering" or "Marketing") or status (for instance, "Antivirus Running"). The Junos OS Enforcer determines whether to allow or deny the traffic to pass based on the authentication results stored in the appropriate authentication table entry.

   The IC Series UAC Appliance pushes authentication table entries to the Junos OS Enforcer when the devices first connect to one another and, as necessary, throughout the session. For example, the IC Series appliance might push updated authentication table entries to the Junos OS Enforcer when the user's computer becomes noncompliant with endpoint security policies, when you change the configuration of a user's role, or when you disable all user accounts on the IC Series appliance in response to a security problem such as a virus on the network.

   If the Junos OS Enforcer drops a packet because of a missing authentication table entry, the device sends a message to the IC Series appliance, which in turn may

Copyright © 2016, Juniper Networks, Inc.

provision a new authentication table entry and send it to the Junos OS Enforcer. This process is called dynamic authentication table provisioning.

3. Once it determines that the traffic may pass based on the authentication table entries, the Junos OS Enforcer maps the flow to a resource. The Junos OS Enforcer uses the destination IP address specified in the flow to create the mapping. Then the device uses that resource as well as the user role specified in the authentication table entry to map the flow to a resource access policy.

   A *resource access policy* specifies a particular resource to which you want to control access based on user role. For instance, you might create a resource access policy that allows only users who are members of the Engineering and Antivirus Running user roles access to the Engineering-Only server. Or you might create a resource access policy that allows members of the No Antivirus Running user role access to the Remediation server on which antivirus software is available for download.

   The IC Series appliance pushes resource access policies to the Junos OS Enforcer when the devices first connect to one another and when you modify your resource access policy configurations on the IC Series appliance.

   If the Junos OS Enforcer drops the packet because of a “deny” policy, the Junos OS Enforcer sends a message to the IC Series appliance, which in turn sends a message to the endpoint’s Odyssey Access Client (if available). (The IC Series appliance does not send “deny” messages to the agentless client.)

4. Once it determines that the traffic may pass based on the resource access policies, the Junos OS Enforcer processes the traffic using the remaining application services defined in the Junos OS policy. The Junos OS Enforcer runs the remaining services in the following order: Intrusion Detection and Prevention (IDP), URL filtering, and Application Layer Gateways (ALGs).

Configuring Junos OS Enforcer Failover Options (CLI Procedure)

To configure IC Series UAC Appliance failover processing, you must configure the Junos OS Enforcer to connect to a cluster of IC Series appliances. The Junos OS Enforcer communicates with one of these IC Series appliances at a time and uses the others for failover processing.

Before you begin:

1. Enable UAC through the relevant Junos OS security policies.
2. Configure the SRX Series device as a Junos OS Enforcer. During the configuration, define a cluster of IC Series appliances to which the Junos OS Enforcer should connect. See “Enabling UAC in a Junos OS Environment (CLI Procedure)” on page 5259.

To configure failover processing:

1. Specify how often (in seconds) the Junos OS Enforcer should expect a heartbeat signal from the IC Series appliance indicating an active connection:

   user@host# **set services unified-access-control interval *seconds***

 Copyright © 2016, Juniper Networks, Inc.

2. Specify the interval (in seconds) at which the Junos OS Enforcer should consider the current connection timed out:

NOTE: Any change in the Unified Access Control's (UAC) contact interval and timeout values in the SRX Series device will be effective only after the next reconnection of the SRX Series device with the IC Series appliance.

```
user@host# set services unified-access-control timeout seconds
```

3. Specify how the Junos OS Enforcer should handle all current and subsequent traffic sessions when its connection to an IC Series appliance cluster times out:

```
user@host# set services unified-access-control timeout-action (close | no-change |
  open)
```

Related Documentation

- Understanding Junos OS Enforcer Policy Enforcement on page 5265
- Testing Junos OS Enforcer Policy Access Decisions Using Test-Only Mode (CLI Procedure) on page 5267
- Verifying Junos OS Enforcer Policy Enforcement on page 5268

Testing Junos OS Enforcer Policy Access Decisions Using Test-Only Mode (CLI Procedure)

When configured in test-only mode, the SRX Series device enables all UAC traffic to go through regardless of the UAC policy settings. The device logs the UAC policy's access decisions without enforcing them so you can test the implementation without impeding traffic.

Before you begin:

1. Enable UAC through the relevant Junos OS security policies. See "Enabling UAC in a Junos OS Environment (CLI Procedure)" on page 5259
2. Configure the SRX Series devices as a Junos OS Enforcer. See "Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure)" on page 5262.
3. If you are connecting to a cluster of IC Series UAC Appliances, enable failover options. See "Configuring Junos OS Enforcer Failover Options (CLI Procedure)" on page 5266.

To activate or deactivate test-only mode, enter the following configuration statement:

```
user@host#  set services unified-access-control test-only-mode (true | false)
```

Copyright © 2016, Juniper Networks, Inc.

Verifying Junos OS Enforcer Policy Enforcement

- Displaying IC Series UAC Appliance Authentication Table Entries from the Junos OS Enforcer on page 5268
- Displaying IC Series UAC Appliance Resource Access Policies from the Junos OS Enforcer on page 5268

Displaying IC Series UAC Appliance Authentication Table Entries from the Junos OS Enforcer

Purpose Display a summary of the authentication table entries configured from the IC Series UAC Appliance.

Action Enter the **show services unified-access-control authentication-table** CLI command.

Displaying IC Series UAC Appliance Resource Access Policies from the Junos OS Enforcer

Purpose Display a summary of UAC resource access policies configured from the IC Series UAC Appliance.

Action Enter the **show services unified-access-control policies** CLI command.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 222

Classifying Traffic with User Roles

- Understanding Unified Access Control on page 5269
- Acquiring User Role Information from an Active Directory Authentication Server on page 5269

Understanding Unified Access Control

In Junos OS Release 12.1 and later, user role firewall security policies let you classify traffic based on the roles to which a user is assigned. Based on match criteria, which includes the user's role, you create policies to apply services that allow or block access to resources. The user role firewall is similar to the identity-based network access control (NAC) solution available with UAC on the SRX Series device. A user role firewall, however, does not require the Junos Pulse/Odyssey installation, and it supports agentless transparent authentication.

User role information can be collected in several ways: locally on the SRX Series device, from a Junos Pulse Access Control Service device, or by relaying authentication data from a third-party authentication server through a Junos Pulse Access Control Service device to the SRX Series device.

Related Documentation

- Acquiring User Role Information from an Active Directory Authentication Server on page 5269

Acquiring User Role Information from an Active Directory Authentication Server

Networks have used the IP address as a way of identifying users and servers. The strategy is based on the assumption that users or groups of users connect to the network from fixed locations and use one device at a time.

Wireless networking and mobile devices require a different strategy. Individuals can connect to the network using multiple devices simultaneously. The way in which devices connect to the network changes rapidly. It is no longer possible to identify a user with a group of statically allocated IP addresses.

In Junos OS Release 12.1 and later, user role firewall security policies let you classify traffic based on the roles to which a user is assigned. Based on match criteria, which includes the user's role, you create policies to apply services that allow or block access to resources. The user role firewall is similar to the identity-based network access control (NAC)

Copyright © 2016, Juniper Networks, Inc.

solution available with UAC on the SRX Series device. A user role firewall, however, does not require the Junos Pulse/Odyssey installation, and it supports agentless transparent authentication.

User role information can be collected in several ways: locally on the SRX Series device, from a Junos Pulse Access Control Service device, or by relaying authentication data from a third-party authentication server through a Junos Pulse Access Control Service device to the SRX Series device.

Incorporating a third-party authentication server into a user role firewall configuration can also provide single sign-on (SSO) support. This allows a browser-based user to authenticate once and have that authentication communicated to other trusted servers in the domain as needed.

- Requirements on page 5270
- Overview on page 5271
- Configuration on page 5272

Requirements

This solution uses the following hardware and software components:

- One MAG Series Junos Pulse Gateway device with software release 4.2 or later
- The MAGx600-UAC-SRX license installed on the MAG Series device
- One SRX Series device with Junos OS Release 12.1 or later
- One Microsoft Active Directory server using version 2008

NOTE: **Microsoft Windows 2003 is also compatible with this functionality, but terminology, pathways, and settings might differ from what is presented in this document.**

Before you begin:

- Ensure that the MAG Series device is configured as an Access Control Service and is accessible to the network. See the *MAG Series Junos Pulse Gateway Hardware Guide* for configuration details.
- Ensure that the MAGx600-UAC-SRX license is installed on the MAG Series device.
- Ensure that the SRX Series device is configured and initialized with Junos OS version 12.1 or later.
- Ensure that the Active Directory authentication server is configured for standard Junos Pulse Access Control Service authentication. See your third-party documentation.
- Ensure that the administrator has the appropriate capabilities for configuring the roles, users, and device interactions.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this solution an SRX Series device obtains user role information dynamically from a Microsoft Active Directory authentication server. Authentication verification and user role information from the Active Directory server is relayed by the Access Control Service on the MAG Series device to the SRX Series device.

Users within the same domain are connected to a LAN segment. They are associated with user role groups, such as developer or manager, depending on their work in the organization. When a user authenticates to the AD authentication server, the user should be able to access protected resources without having to authenticate a second time.

The SRX Series device is configured as an enforcer for the MAG Series device. It receives user role information from the MAG Series device and applies user role firewall policies accordingly to incoming and outgoing traffic.

When the SRX Series device has no user role information for a user, the user's browser is redirected to the MAG Series device. Transparently to the user, the MAG Series device requests verification from the browser. The browser retrieves a token from the Active Directory server confirming authentication and passes it to the MAG Series device. With the information provided by the token, the MAG Series device retrieves user role information for the user from the Active Directory server and creates an authentication table entry consisting of the current IP address and the user role data. The MAG Series device pushes the updated table to the SRX Series device and redirects the browser back to the SRX to request access again. This time, the table does contain user role information which is then retrieved and used as part of the match criteria for applying user role firewall services.

The user is not aware of the process unless the Active Directory (AD) server has no current authentication for the user. When that is the case, the server prompts the user for name and password. Once authentication occurs, the server returns a token to the browser.

The procedure documented here initially configures the MAG Series device as the authenticator. The configuration is later modified to retrieve authentication information from the AD server. This solution uses SPNEGO negotiation and Kerberos authentication to secure communications among the SRX Series device, the MAG Series device, the browser, and the authentication server.

Topology

Figure 236 shows the topology for this deployment in which the MAG Series device is used initially as the authenticationn source. Later, the AD server is used transparently unless the user is not authenticated, in which case he is prompted for a user name and password.

Copyright © 2016, Juniper Networks, Inc.

Figure 236: Single Sign-On Support Topology

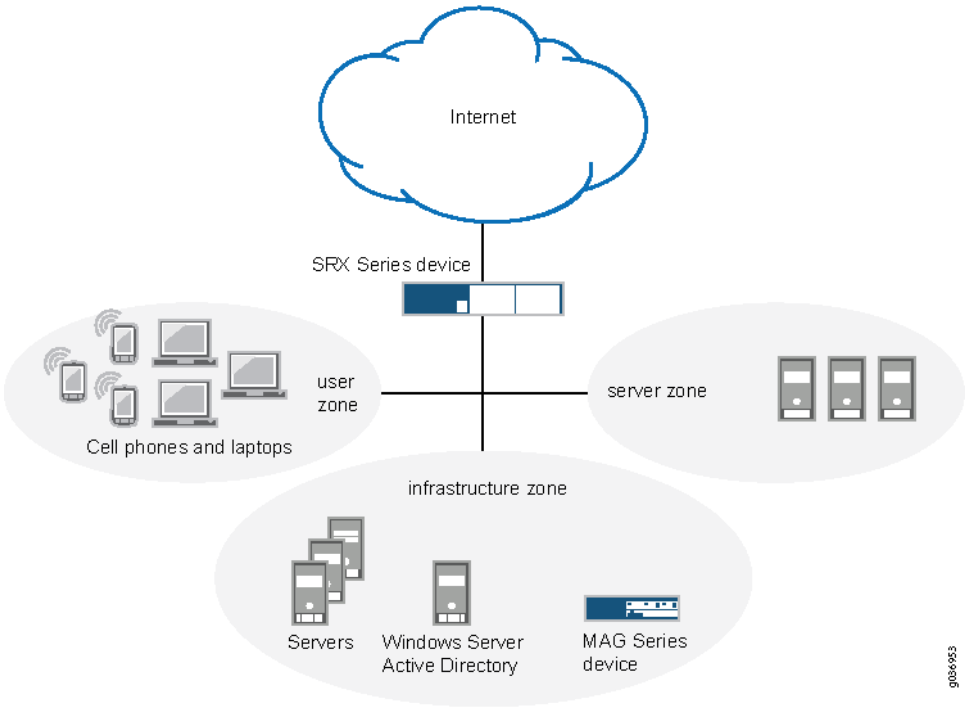

A user's request to access another resource is controlled by roles and groups associated with the user. For example, a user belonging to a group of developers named Dev might have access to a particular test server. The same user might also be the manager and belong to the Mgr group that can access certain HR resources. A contractor working for this manager might require access to the test server as well but not to the HR resources. In this case, the user would be added to the Dev group and perhaps a Contractor group, but not the Mgr group.

User role firewall policies defined on the SRX Series device control the groups and user roles that can access various resources. In this configuration, if user role data does not exist for a user requesting access, a policy redirects the user's browser to the MAG Series device to authenticate the user and retrieve any associated user role data.

A token exchange among the Access Control Service, the browser, and the Active Directory server remains transparent to the user while it verifies the user's authentication. The exchange uses SPNEGO negotiation and Kerberos authentication for encrypting and decrypting messages among the devices.

With information obtained from the response token, the MAG Series device retrieves the user's roles and groups directly from the Active Directory server. It then creates an authentication table entry and passes it to the SRX Series device.

Configuration

Configure the devices for this solution by performing the following tasks.

Copyright © 2016, Juniper Networks, Inc.

- Connect the SRX Series device and the MAG Series device in an enforcer configuration.
- Configure the Access Control Service on the MAG Series device for local user authentication and verify that authentication information is transferred between the devices.
- Configure a captive portal policy on the SRX Series device to redirect any unauthenticated user to the Access Control Service and verify that redirection is functioning properly.
- Configure the Microsoft Active Directory authentication server to interact with the Access Control Service and the endpoints.
- Reconfigure the Access Control Service for remote authentication by the Active Directory server and redefine Active Directory groups for the SRX Series device.
- Configure endpoint browsers for the SPNEGO protocol

NOTE: **Configuring the Access Control Service using local authentication is not necessary for this solution. However, by configuring local authentication first you can verify the captive portal interaction between the MAG Series device and the SRX Series device.**

The following solution requires you to navigate various levels in the configuration hierarchy on the SRX Series device. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

Connecting the SRX Series Device to the Access Control Service

Step-by-Step Procedure

In an enforcer configuration, the Access Control Service on the MAG Series device and the SRX Series device communicate over a secure channel. When the SRX Series device first connects with the Access Control Service, the devices exchange information to ensure secure communication. Optionally, you can use digital security certificates as an enhanced mechanism for establishing trust.

See the *Unified Access Control Administration Guide* for details about configuring certificate trust between the SRX Series device and the Access Control Service.

To connect the SRX Series device and the Access Control Service on the MAG Series device:

1. Configure the SRX Series device.

   a. Configure the zones and interfaces of the devices.

      ```
      user@host# set security zones security-zone user interfaces ge-0/0/0
      user@host# set security zones security-zone infrastructure interfaces ge-0/0/1
      user@host# set security zones security-zone untrust interfaces ge-0/0/2
      ```

   b. Configure the IP addresses of the interfaces.

      ```
      user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.12.12.1/24
      user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.20/24
      user@host# set interfaces ge-0/0/2 unit 0 family inet address 1.1.1.1/24
      ```

Copyright © 2016, Juniper Networks, Inc.

c. Identify the Access Control Service as a new Infranet Controller, and configure the interface for the connection to it.

```
[edit]
user@host# set services unified-access-control infranet-controller mag123
  address 10.0.0.22
user@host# set services unified-access-control infranet-controller mag123
  interface fxp0.0
```

d. Specify the password for securing interactions between the Access Control Service and the SRX Series device.

```
[edit]
user@host# set services unified-access-control infranet-controller mag123
  password "InSub321"
```

NOTE: The same password must be configured on both devices.

e. (Optional) Specify the full name of the Access Control Service certificate that the SRX Series device must match during connection.

```
user@host# set services unified-access-control infranet-controller mag123
  ca-profile ca-mag123-enforcer
```

f. If you are done configuring the SRX Series device, enter commit from configuration mode.

2. Configure the Access Control Service from the administrator console on the MAG Series device.

a. Navigate to the Infranet Enforcer page, and click **New Enforcer**.

b. Select **Junos**, enter the password set previously on the SRX Series device (InSub321), and enter the serial number of the SRX Series device.

c. Click **Save Changes**.

Results When both devices are configured, the SRX Series device connects automatically to the Access Control Service.

- From the Access Control Service, select **System>Status>Overview** to view the status of the connection to the SRX Series device. The diode in the display is green if the connection is functioning. To display additional information, click the device name.
- From operational mode on the SRX Series device, confirm your connection by entering the **show services unified-access-control status** command. If the output does not display the intended configuration, repeat the instructions in this section to correct the configuration.

```
user@host> show services unified-access-control status

Host           Address         Port   Interface     State
       mag123        10.0.0.22        11123  fxp0.0        connected
```


Copyright © 2016, Juniper Networks, Inc.

Configuring the Access Control Service for Local User Authentication

Step-by-Step Procedure

When a user is authenticated, the Access Control Service on the MAG Series device updates its authentication table with the IP address and associated roles of the user, and pushes the updated table to the SRX Series device. If this user data is deleted or modified, the Access Control Service updates the authentication table with the new information and again pushes it to the SRX Series device.

To test the successful transfer and content of the authentication table, this task configures the Access Control Service on the MAG Series device for local authentication. Within this configuration you can test the user role firewall from the SRX Series device without affecting other network operations. A later task modifies this configuration to provide user role retrieval from the remote Active Directory server.

NOTE: It is not a requirement to configure the Access Control Service for local user authentication. It is provided so that you can test each task in the configuration.

To configure the Access Control Service for local authentication:

1. Define roles on the Access Control Service.
   a. From the administrator console of the Access Control Service, select **Users>User Roles>New User Role**.
   b. Enter **dev** as the role name.

      In this solution, use the default values for other role settings.
   c. Click **Save Changes**.

NOTE: This solution assumes that the MAGx600-UAC-SRX license is installed on the Access Control Service. If the full-feature license is installed, you will need to disable OAC Install and enable Agentless Access.

2. Configure the default authentication server.
   a. Select **Authentication>Auth. Servers**.
   b. Select **System Local**. This establishes the MAG Series device as the default authentication server.
3. Create users.
   a. Select the **Users** tab, and click **New**.
   b. Create **user-a** by entering the following details.

Copyright © 2016, Juniper Networks, Inc.

- Username
- User's full name
- Password
- Password confirmation

c. Repeat the previous step to create **user-b**.

d. Click **Save Changes**.

4. Create a realm.

a. Select **Users>User Realms>New User Realm**.

b. Enter **REALM6** as the realm name.

c. Select **System Local** in the Authentication box.

d. Click **Save Changes**.

5. From the same page, create role mapping rules.

a. Select the **Role Mapping** tab, and click **New Rule**.

b. Define two rules with the following details.

- Enter username user-a, and assign it to role dev.
- Enter username user-b, and assign it to role dev.

c. Click **Save Changes**.

6. Set up the default sign-in page.

a. Select **Authentication>Signing In>Sign-in Policies**.

b. Click the default **Sign-in policy (*/)**.

c. In the **Sign-in URL** box, enter the IP address of this device.

d. In **Authentication realm, Available realms**, select REALM6.

e. Click **Save Changes**.

Results Verify the results of the configuration. If the output does not display the intended configuration, repeat the instructions in this section to correct the configuration.

1. Verify that local authentication on the Access Control Service is functioning properly.

- Open a browser window from an endpoint in the network.
- Enter the fully qualified domain name for the Access Control Service.

  The default sign-in page should display.
- Sign in as user-a, and provide the defined password.

Copyright © 2016, Juniper Networks, Inc.

2. From operational mode on the SRX Series device:

   a. Confirm that the authentication table on the SRX Series device was updated with **user-a**.

      user@host> **show services unified-access-control authentication-table**

      ```
      Id    Source IP       Username     Age      Role identifier
      1     203.0.113.102    user-a       0        0000000001.000005.0
      Total: 1
      ```

   b. Confirm that the correct role has been associated with the role identifier.

      user@host> **show services unified-access-control roles**

      ```
      Name                Identifier

      dev                 0000000001.000005.0
      ```

   c. List all roles associated with user-a.

      user@host> **show services unified-access-control authentication-table detail**

      ```
      Identifier: 1
        Source IP: 203.0.113.102
        Username: user-a
        Age: 0
        Role identifier       Role name
          0000000001.000005.0 dev
      ```

Configuring Redirection from the SRX Series Device to the Access Control Service

Step-by-Step Procedure

Local authentication, as configured in the previous task, requires users to log on to the Access Control Service directly to gain access to network resources. The SRX Series device can be configured to automatically redirect the browser of an unauthenticated user to the Access Control Service if a user requests access to a protected resource directly. You can define a user role firewall policy to redirect an unauthenticated user to a captive portal on the Access Control Service for sign-in.

NOTE: Other services, such as IDP, UTM, AppFW, and AppQoS, can be configured as well as the UAC captive portal implementation. The solution focuses on captive portal for authentication for user role implementation only.

To configure redirection from the SRX Series device to the Access Control Service:

1. From configuration mode on the SRX Series device, configure the profile for the captive portal acs-device.

   ```
   [edit]
   user@host# set services unified-access-control captive-portal acs-device
     redirect-traffic unauthenticated
   ```

Copyright © 2016, Juniper Networks, Inc.

2. Add either the redirection URL for the Access Control Service or a default URL.

```
[edit]
user@host# set services unified-access-control captive-portal acs-device
redirect-url “https://%ic-url%/?target=%dest-url%&enforcer=%enforcer-id%”
```

   This command specifies the default target and enforcer variables so that the browser is returned to the SRX Series device after authentication.

3. Allow traffic to the Acitve Directory (AD) server, the Access Control Service, and the other infrastructure servers.

```
[edit]
user@host# set security policies from-zone user to-zone infrastructure policy
  Allow-AD-UAC match source-address any
user@host# set security policies from-zone user to-zone infrastructure policy
  Allow-AD-UAC match destination-address any
user@host# set security policies from-zone user to-zone infrastructure policy
  Allow-AD-UAC application any
user@host# set security policies from-zone user to-zone infrastructure policy
  Allow-AD-UAC then permit
```

4. Configure a security policy that redirects HTTP traffic from zone user to zone untrust if the source-identity is unauthenticated-user.

```
[edit]
user@host# set security policies from-zone user to-zone untrust policy user-role-fw1
  match source-address any
user@host# set security policies from-zone user to-zone untrust policy user-role-fw1
  match destination-address any
user@host# set security policies from-zone user to-zone untrust policy user-role-fw1
  match application http
user@host# set security policies from-zone user to-zone untrust policy user-role-fw1
  match source-identity unauthenticated-user
```

5. Configure the action to be taken when traffic matches the criteria for user-role-fw1.

   In this case, traffic meeting the specified criteria is allowed access to the UAC captive portal defined by the acs-device profile.

```
user@host# set security policies from-zone user to-zone untrust policy user-role-fw1
  then permit application-services uac-policy captive-portal acs-device
```

6. Configure a security policy allowing access to any HTTP traffic from zone user to zone untrust.

```
[edit]
user@host# set security policies from-zone user to-zone untrust policy user-role-fw2
  match source-address any
user@host# set security policies from-zone user to-zone untrust policy user-role-fw2
  match destination-address any
user@host# set security policies from-zone user to-zone untrust policy user-role-fw2
  match application http
user@host# set security policies from-zone user to-zone untrust policy user-role-fw2
  match source-identity any
user@host# set security policies from-zone user to-zone untrust policy user-role-fw2
  then permit
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: It is important to position the redirection policy for unauthenticated users before a policy for “any” user so that UAC authentication is not shadowed by a policy intended for authenticated users.

7. If you are done configuring the policies, commit the changes.

```
[edit]
user@host# commit
```

Results Confirm your configuration with the following procedures. If the output does not display the intended configuration, repeat the instructions in this section to correct the configuration.

1. From configuration mode, confirm your captive portal profile configuration by entering the **show services** command.

```
[edit]
user@host# show services

...
unified-access-control {
 captive-portal acs-device {
  redirect-traffic unauthenticated;
  redirect-url “https://%ic-url%/?target=%dest-url%&enforcer=%enforcer-id%”
...
```

2. From configuration mode, confirm your policy configuration by entering the **show security policies** command.

```
user@host# show security policies

...
from-zone user to-zone infrastructure {
 policy Allow-AD-UAC  {
  match {
   source-address any;
   destination-address any;
   application any;
  }
  then {
   permit
  }
 }
}
from-zone user to-zone untrust {
 policy user-role-fw1 {
  match {
   source-address any;
   destination-address any;
   application http;
   source-identity unauthenticated-user
  }
  then {
   permit {
    application-services {
```

Copyright © 2016, Juniper Networks, Inc.

```
      uac-policy {
       captive-portal acs-device;
      }
     }
    }
   }
  }
 }
}
from-zone user to-zone untrust {
 policy user-role-fw2 {
  match {
   source-address any;
   destination-address any;
   application http;
   source-identity any
  }
  then {
   permit
  }
 }
}
...
```

3. Verify that the redirection policy is functioning correctly.

   a. Open a browser window from a second endpoint in the network.

   b. Enter a third-party URL, such as www.google.com.

      The default sign-in page from the Access Control Service prompts for a user and password.

   c. Enter the username **user-b** and its password.

      The browser should display the requested URL.

      NOTE: If a pop-up blocker is set on the endpoint, it could interfere with this functionality.

   d. From operational mode on the SRX Series device, verify that the authentication data and roles from the Access Control Service were pushed to the SRX Series device successfully.

      user@host> **show services unified-access-control authentication-table**

```
user@host>show services unified-access-control authentication-table
Id    Source IP       Username     Age      Role identifier
1     203.0.113.112    user-a       0        0000000001.000005.0
2     203.0.113.15    user-b       0        0000000001.000005.0
Total: 2
```


Copyright © 2016, Juniper Networks, Inc.

Configuring Active Directory Settings

Step-by-Step Procedure

SPNEGO negotiation and Kerberos authentication are transparent to the user and network administrator, but certain configuration options enable the use of these protocols. This section identifies configuration requirements when using Active Directory as the authentication server. To interact in SPNEGO negotiation, the Access Control Service requires a keytab file created by Active Directory. Refer to your third-party documentation for more information about enabling SPNEGO and Kerberos usage.

This section is not intended to be a tutorial for Active Directory. However, there are specific configuration details required for this solution. See your third-party documentation to set up Active Directory as a domain controller.

To configure the Active Directory authentication server:

1. Add a DNS entry as the UAC service account in the **Forward Lookup Zones**. In this way clients can refer to the MAG Series device by name or by IP address.

   This UAC service account name will be used in the next section when reconfiguring the UAC service on the MAG Series device.

2. Single sign-on authentication requires that the UAC service account password never expires. To modify user settings:

   a. From the Active Directory Users and Computers application in DNS, select **Users>New>User** and select the UAC service account created in step 1.

   b. Select the **Account** tab.

   c. In user settings, click **Password Never Expires**.

3. On the Domain Controller, open a command line, and enter the **ktpass** command to create the SPNEGO keytab file.

   The keytab file created on the Active Directory server contains the full service principal name (SPN) and other encryption information from the server. The keytab file is then uploaded to the Access Control Service on the MAG Series device. This shared information identifies one device to the other whenever encrypted messages and responses are sent.

   Use the following syntax.

   ktpass -out *output-file-name* -mapuser *uac-service-acccount-name* -prin *service://fqdn@REALM*

   ktpass—Third-party Kerberos utility that maps an SPN to a user, in this case, to the UAC service account. The executable is available for download. Refer to your third-party documentation for the source for this utility.

   -out ***output-file-name***—The name for the SPNEGO keytab file you are creating.

   -mapuser ***uac-service-acccount-name***—The name of the UAC service account created in step 1.

Copyright © 2016, Juniper Networks, Inc.

-prin *service://fqdn*@*REALM*—The service principal name. The Kerberos authentication uses the SPN in its communication. It does not use an IP address.

service—The HTTP service.

fqdn—The hostname of the Junos Pulse Access Control Service. The *service://FQDN* portion of the name is provided by the Access Control Service when registering with the Active Directory server.

REALM—The realm of the Active Directory authentication server. It is the same as the domain name. The Kerberos realm name is always in uppercase letters following the recommendation in RFC 1510. This affects interoperability with other Kerberos-based environments.

The following command creates an SPNEGO keytab file named ic.ktpass.

```
ktpass -out ic.ktpass -mapuser icuser@UCDC.COM -princ
  HTTP/mag123.ucdc.com@UCDC.COM -pass Doj73096
```

This file is copied to the Access Control Service on the MAG Series device in the next section when SPNEGO is configured for remote authentication.

Reconfiguring Remote Authentication on the Access Control Service

Step-by-Step Procedure

This section reconfigures the Access Control Service on the MAG Series device to query the remote Active Directory server instead of the local authentication table when authenticating a user. The following steps add services and authentication options to the Access Control Service on the MAG Series device. The configuration of the SRX Series device remains unchanged.

When you reconfigure the realm's authentication server, the Access Control Service displays all roles or groups from the configured domain controller and its trusted domains. Establishing role mapping rules equates the authentication server's roles or groups to those defined on the Access Control Service.

To reconfigure remote authentication on the Access Control Service:

1. From the administrator console of the Access Control Service on the MAG Series device, select **Authentication>Auth. Servers**.
2. Choose the **Active Directory/Windows NT** server type, and click **Add New Server**.
3. Enter the profile of the new authentication server.
   a. Name the Active Directory server.
   b. Enter its NetBIOS domain name in the domain box.

NOTE: You might receive the following message: "Either the server is not a domain controller of the domain, or the NetBIOS name of the domain is different from the Active Directory (LDAP) name." This message is informational and does not affect the processing of the authentication.

Copyright © 2016, Juniper Networks, Inc.

c. Enter the Kerberos Realm name.

   The Kerberos realm name is the FQDN of the Active Directory domain. For example, if “my-company” is the domain or NetBIOS name, my-company.net is the Kerberos realm name.

d. In the Domain Join Configuration section, enter the username and password of the UAC services account which has permission to join computers to the Active Directory domain.

   Select the Save credentials box.

e. Enter the Container name.

   This is the name of the container in Active Directory where you created the UAC services account for the Access Control Service.

f. Enter the Computer Name.

   Specify the machine ID that the Access Control Service uses to join the specified Active Directory domain as a computer. This name is derived from the licence hardware ID of the Access Control Service in the following format: 0161MT2L00K2C0.

g. Verify that the join operation has succeeded.

   The Join Status indicator provides a color-coded status for the domain join operation as follows:

   - Gray: Not started
   - Yellow: In progress
   - Red: Failed to join
   - Green: Joined the domain

h. Select **Kerberos** and **NTLM v2** as the authentication protocols.

i. In the Trusts section, select the Allow trusted domains box.

j. Select **Enable SPNEGO**.

k. Use the Browse button to upload the keytab file that you created in the previous section.

l. Click **Save Changes** and **Test Configuration**.

4. Ensure that SSO is enabled.

   a. Select **Users>User Realms** and the realm name.

   b. Select the Active Directory server name from the **Auth Server** list.

   c. Select the **Authentication Policy** tab.

Copyright © 2016, Juniper Networks, Inc.

d. Verify that the **SSO** option is selected.

e. Click **Save Changes**.

5. Create role-mapping policies for groups acquired from the authentication server.

   Groups from the Active Directory authentication server need to be mapped to roles on the Access Control Service. You first need to create roles, and then map one or more groups to the appropriate role.

   a. Select the Role Mapping tab.

   b. Click **New Rule**, enter a role name, and click **Save Changes**.

      You do not need to add users to the role. Create as many roles as needed to map the groups from the Active Directory authentication server.

   c. Click **Groups**, and select **Search** to list the groups defined in the domain controller.

   d. Select the group names that you want to map to the new role.

   e. Repeat steps b through d to create and map other groups.

   f. Click **Save Changes**.

Copyright © 2016, Juniper Networks, Inc.

Configuring Endpoint Browsers for the SPNEGO

Step-by-Step Procedure

Ensure that endpoint browsers have SPNEGO enabled. For further information, see your third-party documentation.

- Internet Explorer

  From **Security>Local Intranet>Sites>Advanced** add the trusted URL.

  IE performs SPNEGO without any further endpoint configuration but the user is prompted for a username and password. The username and password can be cached.

  To provide single sign-on support, an Internet Explorer configuration can be pushed by configuring a group policy on the Active Directory server. See your third-party documentation for further information.

  Integrated Windows Authentication must be enabled. Use the **Tools>Internet Options>Advanced>Security>Enable Integrated Windows Authentication** path to verify that IWA is enabled.

- Firefox (Windows and MacOS)

  The configuration is in a hidden location. For the URL, type **about:config** and search for the word **trusted**. The required key is the comma separated parameter named **network.negotiate-auth.trusted-uris**.

  NOTE: You need to specify the URL of the resource (in this solution, the FQDN or domain controller value UCDC.com).

- Chrome

  Use the Internet Explorer setting. From **Security>Local Intranet>Sites>Advanced** add the trusted URL.

  An internet Explorer configuration can also be pushed by configuring a group policy on the Active Directory server. This configuration is honored by Chrome.

After successful authentication, the standard agentless page is shown along with a second window with the protected resource (unless a pop-up blocker prevents this).

Related Documentation

- *Unified Access Control Design and Implementation Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 223

Configuring Endpoint Security

- Understanding Endpoint Security Using the Infranet Agent with the Junos OS Enforcer on page 5287
- Configuring Endpoint Security Using the Infranet Agent with the Junos OS Enforcer on page 5288

Understanding Endpoint Security Using the Infranet Agent with the Junos OS Enforcer

An Infranet agent helps you secure traffic on your network starting with the endpoints that initiate communications as follows:

1. The Infranet agent, which runs directly on the endpoint, checks that the endpoint is compliant with your Unified Access Control (UAC) Host Checker policies. You can use a wide variety of criteria within a UAC Host Checker policy to determine compliance. For example, you can configure the Host Checker policy to confirm that the endpoint is running antivirus software or a firewall or that the endpoint is not running specific types of malware or processes.
2. The Infranet agent transmits the compliance information to the Junos OS Enforcer.
3. The Junos OS Enforcer allows or denies the endpoint access to the resources on your network based on the Host Checker compliance results.

Because the Infranet agent runs directly on the endpoint, you can use the Infranet agent to check the endpoint for security compliance at any time. For instance, when a user tries to sign into the IC Series UAC Appliance, you can require the Infranet agent to send compliance results immediately—the user will not even see the sign-in page until the Infranet agent returns positive compliance results to the IC Series appliance. You can also configure the Infranet agent to check for compliance after the user signs in or periodically during the user session.

If the endpoints running the Infranet agent have appropriate access, they will automatically send their compliance results to the IC Series appliance, and the IC Series appliance will update the authentication table entries accordingly and push them to the Junos OS Enforcer. The Junos OS Enforcer supports connections with the Odyssey Access Client and “agentless” Infranet agents.

Copyright © 2016, Juniper Networks, Inc.

Configuring Endpoint Security Using the Infranet Agent with the Junos OS Enforcer

To integrate the Infranet agent into a Junos OS-UAC deployment, no special configuration is required on the Junos OS Enforcer. You simply need to create security policies enabling access to the appropriate endpoints as you would for any other Junos OS-UAC deployment.

Related Documentation

- Understanding Endpoint Security Using the Infranet Agent with the Junos OS Enforcer on page 5287
- *Building Blocks Feature Guide for Security Devices*
- Understanding User Authentication for Security Devices on page 5173

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 224

Configuring IPsec

- Understanding Junos OS Enforcer Implementations Using IPsec on page 5289
- Example: Configuring the Device as a Junos OS Enforcer Using IPsec (CLI) on page 5290

Understanding Junos OS Enforcer Implementations Using IPsec

To configure an SRX Series device to act as a Junos OS Enforcer using IPsec, you must:

- Include the identity configured under the security IKE gateway. The identity is a string such as "gateway1.example.com", where gateway1.example.com distinguishes between IKE gateways. (The identities specify which tunnel traffic is intended.)
- Include the preshared seed. This generates the preshared key from the full identity of the remote user for Phase 1 credentials.
- Include the RADIUS shared secret. This allows the IC Series UAC Appliance to accept RADIUS packets for extended authentication (XAuth) from the Junos OS Infranet Enforcer.

When configuring IPsec between the IC Series appliance, the Odyssey Access Client, and the SRX Series device, you should note that the following are IKE (or Phase 1) proposal methods or protocol configurations that are supported from the IC Series appliance to the Odyssey Access Client:

- IKE proposal: **authentication-method pre-shared-keys** (you must specify **pre-shared-keys**)
- IKE policy:
  - **mode aggressive** (you must use aggressive mode)
  - **pre-shared-key ascii-text *key*** (only ASCII text preshared-keys are supported)
- IKE gateway: dynamic
  - **hostname *identity*** (you must specify a unique identity among gateways)
  - **ike-user-type group-ike-id** (you must specify **group-ike-id**)
  - **xauth access-profile *profile*** (you must specify **xauth**)

Copyright © 2016, Juniper Networks, Inc.

The following are IPsec (or Phase 2) proposal methods or protocol configurations that are supported from the IC Series appliance to the Odyssey Access Client.

- IPsec proposal: **protocol esp** (you must specify **esp**)
- IPsec VPN: **establish-tunnels immediately** (you must specify **establish-tunnels immediately**)

NOTE:

- Only one IPsec VPN tunnel is supported per from-zone to to-zone security policy. This is a limitation on the IC Series appliance.
- Junos OS security policies enable you to define multiple policies differentiated by different source addresses, destination addresses, or both. The IC Series appliance, however, cannot differentiate such configurations. If you enable multiple policies in this manner, the IC Series appliance could potentially identify the incorrect IKE gateway.

Example: Configuring the Device as a Junos OS Enforcer Using IPsec (CLI)

To configure an SRX Series device to act as a Junos OS Enforcer using IPsec:

1. Set system and syslog information using the following configuration statements:

```
system {
host-name test_host;
   domain-name test.juniper.net;
   host-name test_host;
   root-authentication {
   encrypted-password "$ABC123";
   }
   services {
   ftp;
      ssh;
      telnet;
      web-management {
      http {
         interface ge-0/0/0.0;
      }
   }
}
syslog {
user * {
   any emergency;
      }
      file messages {
      any critical;
      authorization info;
      }
      file interactive-commands {
      interactive-commands error;
   }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    max-configurations-on-flash 5;
    max-configuration-rollbacks 5;
    license {
    autoupdate {
      url https://ae1.juniper.net/junos/key_retrieval;
    }
    }
    ntp {
    boot-server 1.2.3.4;
    server 1.2.3.4;
}
}
```

NOTE: On SRX Series devices, the factory default for the maximum number of backup configurations allowed is five. Therefore, you can have one active configuration and a maximum of five rollback configurations. Increasing this backup configuration number will result in increased memory usage on disk and increased commit time.

To modify the factory defaults, use the following commands:

```
root@host# set system max-configurations-on-flash number
root@host# set system max-configuration-rollbacks number
```

where **max-configurations-on-flash** indicates backup configurations to be stored in the configuration partition and **max-configuration-rollbacks** indicates the maximum number of backup configurations.

2. Configure the interfaces using the following configuration statements:

```
interfaces {
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.64.75.135/16;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.100.54.1/16;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      10.101.54.1/16;
    }
  }
}
```

3. Configure routing options using the following configuration statements:

Copyright © 2016, Juniper Networks, Inc.

```
routing-options {
static {
  route 0.0.0.0/0 next-hop 10.64.0.1;
    route 10.11.0.0/16 next-hop 10.64.0.1;
    route 172.0.0.0/8 next-hop 10.64.0.1;
  route 10.64.0.0/16 next-hop 10.64.0.1;
}
}
```

4. Configure security options using the following configuration statements:

```
security {
ike {
  traceoptions {
    file ike;
    flag all;
  }
  proposal prop1 {
  authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
  encryption-algorithm 3des-cbc;
  }
  policy pol1 {
  mode aggressive;
    proposals prop1;
  pre-shared-key ascii-text "$ABC123";
  }
  gateway gateway1 {
  ike-policy pol1;
    dynamic {
    hostname gateway1.juniper.net;
      connections-limit 1000;
    ike-user-type group-ike-id;
    }
    external-interface ge-0/0/0;
  xauth access-profile infranet;
  }
  gateway gateway2 {
  ike-policy pol1;
    dynamic {
    hostname gateway2.juniper.net;
      connections-limit 1000;
    ike-user-type group-ike-id;
    }
    external-interface ge-0/0/0;
  xauth access-profile infranet;
}
}
```

5. Configure IPsec parameters using the following configuration statements:

```
ipsec {
proposal prop1 {
  protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
```


Copyright © 2016, Juniper Networks, Inc.

```
    lifetime-seconds 86400;
    }
    policy pol1 {
    proposals prop1;
    }
    vpn vpn1 {
    ike {
      gateway gateway1;
      ipsec-policy pol1;
    }
    }
    vpn vpn2 {
    ike {
      gateway gateway2;
      ipsec-policy pol1;
    }
  }
  }
```

6. Configure screen options using the following configuration statements:

```
  screen {
  ids-option untrust-screen {
    icmp {
      ping-death;
      }
      ip {
      source-route-option;
      tear-drop;
      }
      tcp {
      syn-flood {
        alarm-threshold 1024;
          attack-threshold 200;
          source-threshold 1024;
          destination-threshold 2048;
          queue-size 2000;
        timeout 20;
        }
      land;
    }
  }
  }
```

7. Configure zones using the following configuration statements:

```
  zones {
  security-zone trust {
    tcp-rst;
      host-inbound-traffic {
      system-services {
        all;
        }
        protocols {
        all;
      }
      }
```

```
    interfaces {
    ge-0/0/0.0;
  }
  }
  security-zone untrust {
  host-inbound-traffic {
    system-services {
      all;
      }
      protocols {
      all;
    }
    }
    interfaces {
    ge-0/0/1.0;
  }
  }
  security-zone zone101 {
  host-inbound-traffic {
    system-services {
      all;
      }
      protocols {
      all;
    }
    }
    interfaces {
    ge-0/0/2.0;
  }
}
}
```

8. Configure policies for UAC using the following configuration statements:

```
policies {
from-zone trust to-zone trust {
  policy default-permit {
    match {
      source-address any;
        destination-address any;
      application any;
      }
      then {
      permit;
    }
  }
  }
  from-zone trust to-zone untrust {
  policy default-permit {
    match {
      source-address any;
        destination-address any;
      application any;
      }
      then {
      permit;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    }
    policy default-deny {
    match {
      source-address any;
        destination-address any;
      application any;
      }
      then {
      permit;
    }
    }
    policy pol1 {
    match {
      source-address any;
        destination-address any;
      application any;
      }
      then {
      permit {
        tunnel {
          ipsec-vpn vpn1;
          }
          application-services {
          uac-policy;
        }
        }
        log {
        session-init;
        session-close;
      }
    }
}
}
from-zone untrust to-zone trust {
policy pol1 {
    match {
      source-address any;
        destination-address any;
      application any;
      }
      then {
      permit;
        log {
        session-init;
        session-close;
      }
    }
}
}
from-zone trust to-zone zone101 {
policy pol1 {
    match {
      source-address any;
        destination-address any;
      application any;
```

```
        }
        then {
        permit {
          tunnel {
            ipsec-vpn vpn2;
            }
            application-services {
            uac-policy;
          }
          }
          log {
          session-init;
          session-close;
        }
      }
      }
      policy test {
      match {
        source-address any;
          destination-address any;
        application any;
        }
        then {
        permit;
      }
    }
    }
    default-policy {
    deny-all;
  }
  }
  }
```

9. Configure RADIUS server authentication access using the following configuration statements:

```
access {
profile infranet {
  authentication-order radius;
    radius-server {
    10.64.160.120 secret "$ABC123";
  }
}
}
```

10. Configure services for UAC using the following configuration statements:

```
services {
unified-access-control {
  infranet-controller IC27 {
    address 3.23.1.2;
      interface ge-0/0/0.0;
    password "$ABC123";
    }
    infranet-controller prabaIC {
    address 10.64.160.120;
      interface ge-0/0/0.0;
```


Copyright © 2016, Juniper Networks, Inc.

```
        password "$ABC123";
        }
        certificate-verification optional;
        traceoptions {
        flag all;
    }
}
}
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 225

Configuring Captive Portal

- Understanding the Captive Portal on the Junos OS Enforcer on page 5299
- Understanding Captive Portal Configuration on the Junos OS Enforcer on page 5301
- Understanding the Captive Portal Redirect URL Options on page 5301
- Example: Creating a Captive Portal Policy on the Junos OS Enforcer on page 5302
- Example: Configuring a Redirect URL for Captive Portal on page 5305

Understanding the Captive Portal on the Junos OS Enforcer

In a Unified Access Control (UAC) deployment, users might not be aware that they must first sign in to the IC Series UAC Appliance for authentication and endpoint security checking before they are allowed to access a protected resource behind the Junos OS Enforcers. To help users sign in to the IC Series appliance, you can configure the captive portal feature. The captive portal feature allows you to configure a policy in the Junos OS Enforcer that automatically redirects HTTP traffic destined for protected resources to the IC Series appliance or to a URL configured in the Junos OS Enforcer.

You can configure a captive portal for deployments that use either source IP enforcement or IPsec enforcement, or a combination of both enforcement methods.

Figure 237 shows the captive portal feature enabled on a Junos OS Enforcer. Users accessing protected resources are automatically redirected to the IC Series appliance:

1. Users point to a protected resource using the browser.
2. The Junos OS Enforcer determines that the user is not authenticated and redirects the request to the IC Series appliance or another server.
3. Users enter their Infranet username and password to log in.
4. The IC Series appliance passes the user credentials to an authentication server.
5. After authentication, the IC Series appliance redirects the users to the protected resource they wanted to access.

Copyright © 2016, Juniper Networks, Inc.

Figure 237: Enabling the Captive Portal Feature on a Junos OS Enforcer

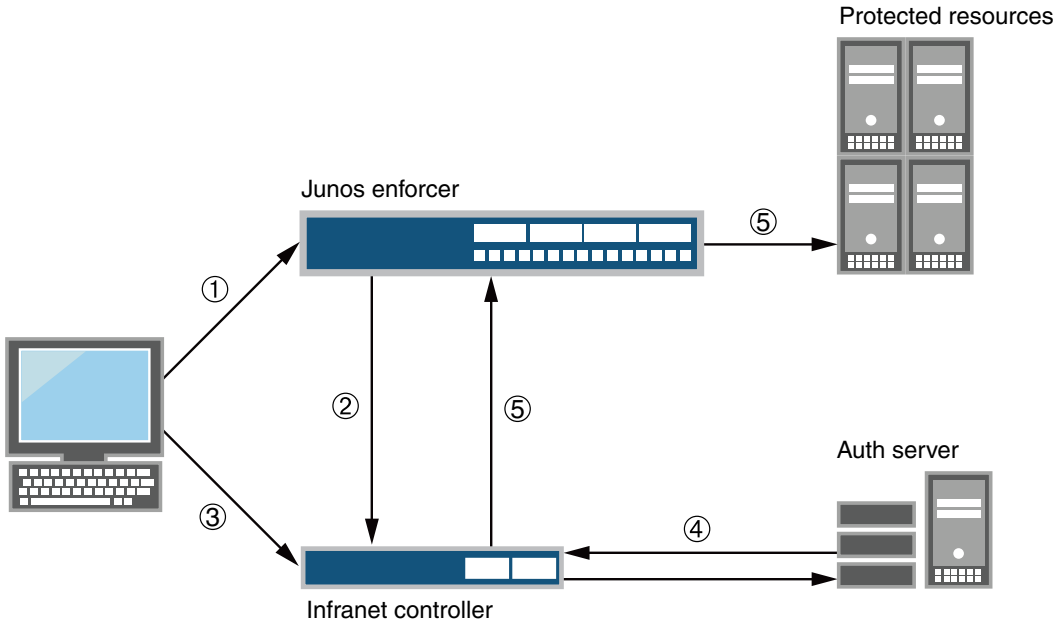

By default, the Junos OS Enforcer encodes and forwards to the IC Series appliance the protected resource URL that the user entered. The IC Series appliance uses the protected resource URL to help users navigate to the protected resource. The manner in which the IC Series appliance uses the protected resource URL depends on whether or not the user's endpoint is running the Odyssey Access Client or Junos Pulse. If the user's endpoint is not running the Odyssey Access Client or Junos Pulse (that is, it is in an agentless or Java agent configuration), the IC Series appliance automatically opens a new browser window and uses HTTP to access the protected resource after the user signs in. If the endpoint is using the Odyssey Access Client, the IC Series appliance inserts a hypertext link in the webpage that automatically opens after the user signs in. The user must then click that hypertext link to access the protected resource by means of HTTP in the same browser window.

The Junos OS Enforcer supports the captive portal feature only for HTTP traffic. If you attempt to access a protected resource by using HTTPS or a non-browser application (such as an e-mail application), the Junos OS Enforcer does not redirect the traffic. When using HTTPS or a non-browser application, you must manually sign in to the IC Series appliance first before attempting to access protected resources.

Copyright © 2016, Juniper Networks, Inc.

Understanding Captive Portal Configuration on the Junos OS Enforcer

To configure the captive portal feature, you create a security policy on the Junos OS Enforcer and then specify a redirection option for the captive portal security policy. You can choose to redirect traffic to an external server or to the IC Series UAC Appliance. You can also choose to redirect all traffic or unauthenticated traffic only.

- Redirecting traffic to an external webserver—You can configure the Junos OS Enforcer to redirect HTTP traffic to an external webserver instead of the IC Series appliance. For example, you can redirect HTTP traffic to a webpage that explains to users the requirement to sign in to the IC Series appliance before they can access the protected resource. You could also include a link to the IC Series appliance on that webpage to help users sign in.
- Redirecting unauthenticated traffic—Select this option if your deployment uses source IP only or a combination of source IP and IPsec. The Junos OS Enforcer redirects clear-text traffic from unauthenticated users to the currently connected IC Series appliance or to an IP address or domain name that you specify in a redirect URL. After a user signs in to the IC Series appliance and the user's endpoint system meets the requirements of the IC Series appliance security policies, the Junos OS Enforcer allows the user's clear-text traffic to pass through in source IP deployments. For IPsec deployments, the Odyssey Access Client creates a VPN tunnel between the user and the Junos OS Enforcer. The Junos OS Enforcer then applies the VPN policy, allowing the encrypted traffic to pass through.
- Redirecting all traffic—Specify this option if you want to redirect all traffic to the URL that you specify in a redirect URL.
- Redirecting traffic with multiple IC Series appliances—You can configure multiple IC Series appliances on your Junos OS Enforcer, but it is connected to only one IC Series appliance at any given time. If the connection to the IC Series appliance fails, the Junos OS Enforcer tries to connect to next configured IC Series appliance. As a result, you cannot be sure which IC Series appliance is connected to the Junos OS Enforcer at any given time. To ensure that the Junos OS Enforcer redirects traffic to the connected IC Series appliance, configure the default redirect URL or the **%ic-ip%** option in the URL.

Related Documentation

- Example: Creating a Captive Portal Policy on the Junos OS Enforcer on page 5302

Understanding the Captive Portal Redirect URL Options

By default, after you configure a captive portal policy, the Junos OS Enforcer redirects HTTP traffic to the currently connected IC Series UAC Appliance by using HTTPS. To perform the redirection, the Junos OS Enforcer uses the IP address or domain name that you specified when you configured the IC Series appliance instance on the Junos OS Enforcer. The format of the URL that the Junos OS Enforcer uses for default redirection is:

```
https://%ic-ip%/?target = %dest-url% &enforcer = %enforcer-id% &policy =
%policy-id% &dest-ip = %dest-ip%
```

Copyright © 2016, Juniper Networks, Inc.

If you configured your Junos OS Enforcer to work with multiple IC Series appliances in a cluster, and the current IC Series appliance becomes disconnected, the Junos OS Enforcer automatically redirects HTTP traffic to the next active IC Series appliance in its configuration list. The Junos OS Enforcer redirects traffic to only one IC Series appliance at a time.

Otherwise, the browser displays a certificate warning to users when they sign in. You do not need to override the default redirection destination except in these situations:

- You are using a VIP for a cluster of IC Series appliances, and the Junos OS Enforcer is configured to connect to the IC Series appliance physical IP addresses.
- You want to redirect traffic to a webserver instead of the IC Series appliance.
- If, because of split DNS or IP routing restrictions at your site, the Junos OS Enforcer uses a different address for the IC Series appliance than endpoints, you must specify the domain name or IP address that endpoints must use to access the IC Series appliance.

NOTE: If a captive portal policy is configured with the IC Series UAC Appliance URL as the target, then use only HTTPS to redirect traffic.

Table 470 lists different options that you can configure in the redirect URL string.

Table 470: Redirect URL String Options

URL String	Description
%dest-url%	Specifies the protected resource which the user is trying to access.
%enforcer-id%	Specifies the ID assigned to the Junos OS Enforcer by the IC Series appliance.
%policy-id%	Specifies the encrypted policy ID for the captive portal security policy that redirected the traffic.
%dest-ip%	Specifies the IP address or hostname of the protected resource which the user is trying to access.
%ic-ip%	Specifies the IP address or hostname of the IC Series appliance to which the Junos OS Enforcer is currently connected.

Example: Creating a Captive Portal Policy on the Junos OS Enforcer

This example shows how to create a captive portal policy on the Junos OS Enforcer. In this example, you deploy a Junos OS Enforcer in front of the data center resources you want to protect and configure the captive portal feature on the Junos OS Enforcer. The

Copyright © 2016, Juniper Networks, Inc.

Junos OS Enforcer then automatically redirects HTTP traffic destined for the protected resource to the IC Series UAC Appliance for authentication.

- Requirements on page 5303
- Overview on page 5303
- Configuration on page 5303
- Verification on page 5305

Requirements

Before you begin:

- Deploy the IC Series appliance in the network so that users can access the device. Use the internal port on the IC Series appliance to connect users, the Junos OS Enforcer, and authentication servers. See "Configuring Communications Between the Junos OS Enforcer and the IC Series UAC Appliance (CLI Procedure)" on page 5262.
- Set up security zones and interfaces on the Junos OS Enforcer. Make sure that end users are in a different security zone than protected resources. For example, protected resources in the data center are configured in the trusted zone and users in an untrusted zone. See "Example: Creating Security Zones" on page 1051.
- Add individual users to either an external authentication server or the local authentication server. Set up roles and realms for individual users. You can provision access to protected resources based on your network security needs.

Overview

In this example, you want to protect the trusted zone from users on the LAN by making sure that only compliant and authenticated users are granted access. New users join your network every month. You want to configure the captive portal feature on your system so that unauthenticated users are redirected to the IC Series appliance automatically without requiring new users to remember to log in to the IC Series appliance.

The configuration instructions in this topic describe how to create a security policy called **my-policy**, specify a match condition for this policy, specify the captive portal policy as a part of the UAC policy, and set criteria for redirecting traffic to the IC Series appliance. In this example, the policy **my-policy**:

- Specifies the match condition to include any traffic from a previously configured zone called **trust** to another previously configured zone called **untrust**.
- Specifies the captive portal policy called **my-captive-portal-policy** as part of the UAC policy.
- Specifies the redirect-traffic criteria as **unauthenticated**.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy

Copyright © 2016, Juniper Networks, Inc.

and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone untrust to-zone trust policy my-policy match
  destination-address any source-address any application any
set security policies from-zone untrust to-zone trust policy my-policy then permit
  application-services uac-policy captive-portal my-captive-portal-policy
set services unified-access-control captive-portal my-captive-portal-policy redirect-traffic
  unauthenticated
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To create a captive portal policy on the Junos OS Enforcer:

1. Specify the match condition for the policy.

   ```
   [edit security policies from-zone untrust to-zone trust policy my-policy]
   user@host# set match destination-address any source-address any application
     any
   ```

2. Specify the captive portal policy as part of the UAC policy to be applied on the traffic that matches the conditions specified in the security policy.

   ```
   [edit security policies from-zone untrust to-zone trust policy my-policy]
   user@host# set then permit application-services uac-policy captive-portal
     my-captive-portal-policy
   ```

3. Redirect all unauthenticated traffic to the IC Series appliance.

   ```
   [edit services unified-access-control]
   user@host# set captive-portal my-captive-portal-policy redirect-traffic
     unauthenticated
   ```

Results

Confirm your configuration by entering the **show services** and **show security policies** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  user@host# show services
  unified-access-control {
  captive-portal my-captive-portal-policy {
    redirect-traffic unauthenticated;
  }
}

[edit]
  user@host# show security policies
  ...
  from-zone untrust to-zone trust {
  policy my-policy {
```


Copyright © 2016, Juniper Networks, Inc.

```
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit {
            application-services {
              uac-policy {
                captive-portal my-captive-portal-policy;
              }
            }
          }
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying the Captive Portal Policy on page 5305

Verifying the Captive Portal Policy

Purpose Verify that the captive portal policy was created.

Action From operational mode, enter the **show security policies detail** command.

Related Documentation

- Understanding Captive Portal Configuration on the Junos OS Enforcer on page 5301

Example: Configuring a Redirect URL for Captive Portal

This example shows how to redirect traffic to the currently connected IC Series UAC Appliance or to an IP address or domain name that you specify in a redirect URL. We recommend the default configuration that redirects traffic to the IC Series appliance for authentication.

- Requirements on page 5305
- Overview on page 5306
- Configuration on page 5306
- Verification on page 5306

Requirements

Before you specify the redirect URL, make sure you configure the captive portal policy. For information about creating the captive portal policy, see "Example: Creating a Captive Portal Policy on the Junos OS Enforcer" on page 5302.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you configure the URL to redirect traffic to the IC Series appliance and after authentication to forward the traffic automatically to the protected resource. The configuration instructions in this topic describe how to set the URL to https://my-website.com.

You can redirect traffic to the currently connected IC Series appliance or to an IP address or domain name that you specify in a redirect URL. We recommend the default configuration that redirects traffic to the IC Series appliance for authentication.

If you need to override the default redirection destination, you can specify any combination of redirect options:

- **https://*IP or domain name/URL path*/target=%dest-url%**—Forwards users to the protected resource automatically after authentication with the IC Series appliance or webserver. The Junos OS Enforcer replaces the **%dest-url%** parameter with the protected resource URL and then forwards the protected resource URL in encrypted form to the IC Series appliance.
- **https://*IP or domain name/*target=*URL path***—Forwards users to the specified URL automatically after authentication with the IC Series appliance or webserver.
- **https://*IP or domain name/URL path***—Redirects users to the IC Series appliance authentication page but not be forwarded to the protected resource after authentication. Users must manually open a new browser window and enter the protected resource URL again after signing in.
- **redirect-all**—Redirects all traffic to the URL that you specify in a redirect URL.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the redirect URL for the captive portal feature on the Junos OS Enforcer:

1. Specify the redirect URL for the preconfigured captive portal policy.

```
[edit services unified-access-control]
user@host# set captive-portal my-captive-portal-policy redirect-url
  https://192.168.0.100/target=my-website.com
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show services unified-access-control captive-portal my-captive-portal-policy** command.

Copyright © 2016, Juniper Networks, Inc.

PART 69

Configuring Integrated User Firewall

- Understanding Integrated User Firewalls on page 5309
- Managing Event Logs on page 5329

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 226

Understanding Integrated User Firewalls

- Overview of Integrated User Firewall on page 5309
- Understanding Active Directory Authentication Tables on page 5312
- LDAP Functionality in Integrated User Firewall on page 5317
- Example: Configuring Integrated User Firewall on page 5320

Overview of Integrated User Firewall

This topic includes the following sections:

- Integrated User Firewall and Authentication Sources on page 5309
- Benefits of Integrated User Firewall on page 5310
- How the Integrated User Firewall Works on page 5310
- Deployment Scenario for User Firewall Integration with Windows Active Directory on page 5311
- Limitations on page 5311

Integrated User Firewall and Authentication Sources

The SRX Series device already supports Unified Access Control (UAC) integration with Network Access Control (NAC) and a user firewall that can derive its authentication source from Windows Active Directory via the UAC MAG Series Junos Pulse Gateway. However, many customers want simple user firewall functionality without full NAC, and do not want the additional cost or complexity of user role firewall (which has Active Directory dependencies such as Kerberos, SPNEGO on Browsers, Active Directory DNS/Certs, and UAC configuration).

The integrated user firewall feature fulfills the requirement for simplicity. It retrieves user-to-IP address mappings from the Windows Active Directory to use in firewall policies as match criteria. This feature consists of the SRX Series polling the event log of the Active Directory controller to determine, by username and source IP address, who has logged in to the SRX Series device. Then the username and group information are queried from the LDAP service in the Active Directory controller. Once the SRX Series has the IP address, username, and group relationship information, it generates authentication entries. With the authentication entries, the SRX Series UserFW module enforces user-based and group-based policy control over traffic.

Copyright © 2016, Juniper Networks, Inc.

For a comparison of integrated user firewall, user role firewall, and UAC NAC, see "Understanding the Three-Tiered User Firewall Features" on page 5173.

Benefits of Integrated User Firewall

The integrated user firewall feature introduces an authentication source via integration with Microsoft Active Directory technology.

- It provides visibility into who is accessing the SRX Series and best-effort security for access to the SRX Series.
- It is a single-box solution, requiring only an SRX Series.
- It requires fewer configuration steps than the UAC integration with NAC, which uses the UAC MAG Series.
- It does not require the configuration of a captive portal, although that option is available to enforce on users who do not authenticate.
- It is ideal for small-to-medium businesses and low-scale deployments.
- It supports High Availability (HA).

How the Integrated User Firewall Works

At a high level, this feature involves the UserID process in the SRX Series Routing Engine, which reads the Windows event log from the Active Directory controller and abstracts IP address-to-user mapping information. The process correlates users to the groups to which they belong, via the LDAP protocol with LDAP service in the Active Directory controller. Thus, the process has gathered enough information to generate authentication entries. The network administrator then references the authentication entries in user firewall security policies to control traffic.

A more detailed explanation of how this feature works is as follows:

1. The SRX Series reads the Active Directory event log to get source IP address-to-username mapping information. To do so, a process in the SRX Series Routing Engine implements a Windows Management Instrumentation (WMI) client with Microsoft Distributed COM/Microsoft RPC stacks and an authentication mechanism to communicate with a Windows Active Directory controller in an Active Directory domain. Using event log information retrieved from the Active Directory controller, the process knows the IP addresses of active Active Directory users and abstracts IP-to-Active Directory username mapping information. The process monitors Active Directory event log changes via the same WMI DCOM interface to adjust local mapping information to reflect any change in the Active Directory server.
2. The process uses LDAP to query the LDAP service interface of the Active Directory to identify the groups to which users belong. Having the IP address, the Active Directory user, and the groups, the process can generate authentication entries accordingly.
3. The process pushes the authentication entries to the Packet Forwarding Engine authentication table. The Packet Forwarding Engine uses the entries and user policy to apply user firewall access control to traffic.

 Copyright © 2016, Juniper Networks, Inc.

This feature supports two domains and up to 10 Active Directory controllers in a domain.

Deployment Scenario for User Firewall Integration with Windows Active Directory

Figure 238 illustrates a typical scenario where the integrated user firewall feature is deployed. Users in the Active Directory domain and users outside the Active Directory domain want access to the Internet through an SRX Series device. The domain controller might also act as the LDAP server.

Figure 238: Scenario for Integrated User Firewall

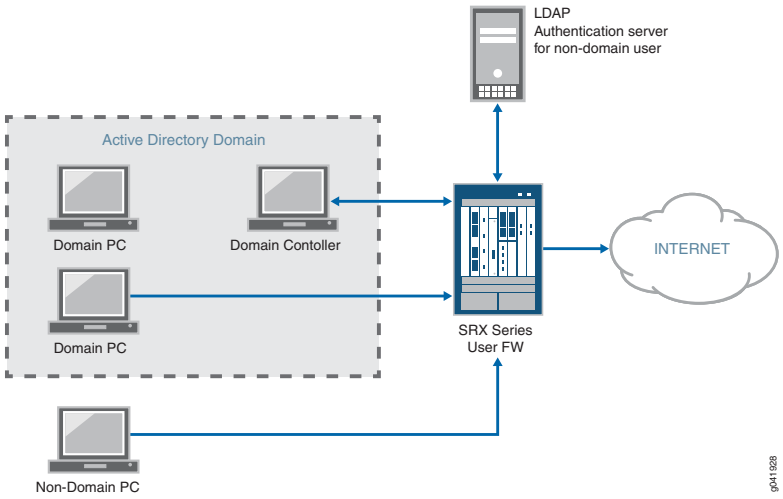

The SRX Series device reads and analyzes the event log of the domain controller and generates an authentication table as an Active Directory authentication source for this feature. The user firewall is aware of any domain user on an Active Directory domain device via the Active Directory authentication source. The SRX Series device administrator configures a user firewall policy that enforces the desired user-based or group-based access control.

For any non-domain user or domain user on a non-domain machine, the administrator specifies a captive portal to force the user to do firewall authentication (if the SRX Series supports captive portal for the traffic type). After the user enters a name and password and passes firewall authentication, the SRX Series gets firewall authentication user/group information and can enforce user firewall policy to control the user accordingly.

In addition to captive portal, if the IP address or user information is not available from the event log, the user can again log in to the Windows PC to generate an event log entry. Then the system generates the user's authentication entry accordingly.

Limitations

- Windows Active Directory controllers older than Windows 2003 are not supported.
- Tracking the status of non-Windows Active Directory users is not supported.
- IPv6 addresses are not supported.

Copyright © 2016, Juniper Networks, Inc.

- Logical systems are not supported.
- The WMIC does not support multiple users logged onto the same PC.
- Domain controllers and domain PCs must be running Windows OS. The minimum support for a Windows client is Windows XP. The minimum support for a server is Windows Server 2003.

Related Documentation

- Understanding the Three-Tiered User Firewall Features on page 5173
- Understanding How the WMIC Reads the Event Log on the Domain Controller on page 5329
- Understanding Active Directory Authentication Tables on page 5312
- Understanding Integrated User Firewall Domain PC Probing on page 5332
- Example: Configuring Integrated User Firewall on page 5320
- **user-identification (Services) on page 5499**

Understanding Active Directory Authentication Tables

This topic includes the following sections:

- Active Directory Authentication as an Authentication Source on page 5312
- Active Directory Authentication Tables on page 5313
- State Information for Active Directory Authentication Table Entries on page 5314
- Active Directory Authentication Table Management on page 5315
- Timeout Interval for Table Entries on page 5316

Active Directory Authentication as an Authentication Source

On an SRX Series device, user information tables serve as the authentication source for information required by firewall security policies. The SRX Series device supports various user information tables including local, user firewall, and Unified Access Control (UAC) types. The integrated user firewall feature introduces another type of authentication source—Active Directory authentication.

The integrated user firewall feature gathers user and group information for Active Directory authentication by reading domain controller event logs, probing domain PCs, and querying Lightweight Directory Access Protocol (LDAP) services within the configured Windows domain. Up to two Windows domains are supported.

From the user and group information, the integrated user firewall feature generates an Active Directory authentication table on the Routing Engine of the SRX Series device, which then pushes the authentication table to the Packet Forwarding Engine. Security policies use the information in the table to authenticate users and to provide access control for traffic through the firewall.

Copyright © 2016, Juniper Networks, Inc.

Active Directory Authentication Tables

The Active Directory authentication table contains the IP address, username, and group mapping information that serves as the authentication source for the SRX Series integrated user firewall feature. Information in the table is obtained by reading Window Active Directory domain controller event logs, probing domain PCs, and querying LDAP services within a specified Windows domain.

Reading domain controller event logs generates a list of IP address-to-user mapping information that is used to create entries in the Active Directory authentication table. Once entries have been added in the table, a query is sent to the LDAP server for user-to-group mapping information.

The LDAP server returns all group information; this includes not only information about the groups you directly belong to, but also all the parent (and parent of the parent and so on) groups that you belong to. Group information returned from the LDAP server is compared with the source identity in security policies. If there is a match, Active Directory authentication table entries are updated to include only the group information provided in the security policy. In this way, only relevant group information is listed in the authentication table. Whenever source identity is updated, the authentication table is also updated to reflect the up-to-date relevant group information for all listed users.

When user traffic arrives at the firewall, the Active Directory authentication table is searched for an entry corresponding to the source IP address of the traffic. If an entry exists, policies matching that entry are applied to the traffic and access is allowed or denied.

Table 1 lists Active Directory authentication table support by SRX Series devices:

Table 471: Active Directory Authentication Table Support by SRX Series Devices

SRX Series Devices	Active Directory Authentication Table Entries	Domains	Active Directory Controllers
SRX300	500	1	5
SRX320	500	1	5
SRX340, 345	1000	1	5
SRX550HM	5000	2	10
SRX1500	20,000	2	10
SRX5000 line	100,000	2	10
vSRX	5000	2	10

Once the maximum number of authentication table entries is reached, no additional entries are created.

Copyright © 2016, Juniper Networks, Inc.

To be compliant with the Active Directory authentication table, entries must adhere to the following parameters:

- Usernames are limited to 64 characters.
- Group names are limited to 64 characters.
- Each entry can be associated with up to 200 relevant groups (configured in the source identity field). For example, if you belong to 1000 groups in LDAP and out of these, no more than 200 groups are configured in the source identity field, you are compliant with the Active Directory authentication table.

The Active Directory Authentication table must be enabled as the authentication source for integrated user firewall information retrieval.

```
user@host# set security user-identification authentication source
  active-directory-authentication-table priority priority
```

NOTE: The priority option specifies the sequence in which user information tables are checked. Using the lowest setting for the Active Directory authentication source specifies the highest priority, meaning that the Active Directory authentication source is searched first.

State Information for Active Directory Authentication Table Entries

Active Directory authentication table entries can be in one of four states:

Initial—Specifies that IP address-to-user mapping information was obtained by reading domain controller event logs and an entry was added to the authentication table. Entries in this state are changed to valid when the table is pushed from the Routing Engine to the Packet Forwarding Engine.

Valid—Specifies that a valid entry was obtained by reading domain controller event logs or that a valid response was received from a domain PC probe and the user is a valid domain user.

Invalid—Specifies that an invalid response was received from a domain PC probe and the user is an invalid domain user.

Pending—Specifies that a probe event generated an entry in the authentication table, but no probe response has been received from the domain PC. If a probe response is not received within 90 seconds, the entry is deleted from the table.

For a list of probe responses, see "Understanding Integrated User Firewall Domain PC Probing" on page 5332.

To display Active Directory authentication entries, along with their state information, use the following command:

```
user@host>show services user-identification active-directory-access
  active-directory-authentication-table all
```

```
Domain: www.example1.net
Total count: 2
```


Copyright © 2016, Juniper Networks, Inc.

```
Source IP       Username      Groups          State
 192.168.10.2     u2           r1, r3, r4      initial
 192.168.10.3     u3           r5, r6, r4      pending


Domain: www.example2.net
Total count: 2
Source IP       Username      Groups          State
 10.1.1.2         u4           r1, r3, r4      valid
 10.1.1.3         u5           r5, r6, r4      invalid
```

Command options allow you to display information by **user** or **group**, and to define additional output levels—**brief**, **domain**, **extensive**, **node**.

Active Directory Authentication Table Management

Windows domain environments are constantly changing as users log in and out of the network and as network administrators modify user group information. The integrated user firewall feature manages changes in the Windows domain by periodically reading domain controller event logs and querying the LDAP server for user-to-group mapping information. That information is used in updating the Active Directory authentication table as appropriate.

Additionally, a probe function is provided to address changes that occur between reading event logs, or to address the case where event log information is lost. An on-demand probe is triggered when client traffic arrives at the firewall but a source IP address for that client cannot be found in the table. And at any point, manual probing is available to probe a specific IP address.

Changes to the active Directory Authentication table also occur due to source identity changes in the security policy configuration.

Table 472 describes events that trigger an Active Directory authentication table update.

Table 472: Events Triggering Active Directory Authentication Table Updates

Event	Active Directory Authentication Table Update
A domain controller event log is read at configured intervals.	New IP address-to-user entries are added in the authentication table in initial state. Group information is retrieved from the LDAP server. When the authentication entry is pushed to Packet Forwarding Engine, the state is changed to valid.
An on-demand or manual probe is sent to a domain PC.	An entry is added in the authentication table in pending state. If a probe response is not returned within 90 seconds, the state of the entry is deleted.
An on-demand or manual probe response is received from a domain PC.	Based on the response, entries in pending state are changed to valid or invalid. For valid responses, the group information is retrieved from the LDAP server. For invalid responses, the entry is marked as invalid.

Copyright © 2016, Juniper Networks, Inc.

Table 472: Events Triggering Active Directory Authentication Table Updates *(continued)*

Event	Active Directory Authentication Table Update
An LDAP server query identifies new user-to-group mapping information.	Entries are updated with the group information.
An LDAP server query identifies deleted user information.	Entries associated with that user are deleted from the table.
An LDAP server query identifies deleted group information.	The affected group information is updated. For example, user2 belongs to group2, and group2 belongs to group1. And, group1 is listed as a source-identity for group2. For any authentication entry of user2, group1 is listed in its relevant groups. However, if group2 is removed from the LDAP server, user2 loses the connection with group1, and as a result, group1 is removed from the user2 authentication table.
An LDAP server query identifies added group information.	If the group is referenced in a security policy, entries associated with this group are updated to add the group information.
The source identity information is removed from a security policy configuration.	Entries associated with the source identity are deleted from Active Directory authentication table.

NOTE: If an entry is deleted from the table, any sessions attached to that entry are also deleted. If an entry in the table is updated to add or remove group information, there is no impact to existing sessions for that entry.

To manually delete an entry from the table, use the **request services user-identification active-directory-access active-directory-authentication-table** command. Options exist for deleting a specific IP address, domain, group, or user.

To clear the contents of the Active Directory authentication table, use the **clear services user-identification active-directory access active-directory-authentication-table** command.

Timeout Interval for Table Entries

When a user is no longer active, a timer is started for that user's entry in the Active Directory authentication table. When time is up, the user's entry is removed from the table. Entries in the table remain active as long as there are sessions associated with the entry.

To set the timeout value, use the following statement:

```
user@host# set services user-identification active-directory-access
  authentication-entry-timeout minutes
```

The default **authentication-entry-timeout** interval is 30 minutes. To disable timeouts, set the interval to 0.

Copyright © 2016, Juniper Networks, Inc.

NOTE: We recommend that you disable timeouts when disabling on-demand probing in order to prevent someone from accessing the Internet without logging in again.

To view timeout information for Active Directory authentication table entries, use the following command:

```
user@host>show services user-identification active-directory-access
    active-directory-authentication-table all extensive

Domain: www.example1.net
Total entries: 2
Source IP: 192.168.1.2
Username: u2
Groups: r1, r3, r4
State: initial
Access start date: 2014-03-22
Access start time: 10:56:58
Age time: 20 min

Source IP: 192.168.1.3
Username: u3
Groups: r5, r6, r4
State: pending
Access start date: 2014-03-22
Access start time: 10:46:58
Age time: 10 min
```

This example shows that the timer has started for two entries—the entry for user u2 will time out in 20 minutes, while the entry for user u3 will time out in 10 minutes. When session traffic is associated with an entry, the age time value changes to “infinite.”

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Understanding Integrated User Firewall Domain PC Probing on page 5332
- LDAP Functionality in Integrated User Firewall on page 5317
- **active-directory-authentication-table on page 5402**
- **user-identification (Services) on page 5499**

LDAP Functionality in Integrated User Firewall

This topic includes the following sections:

- Role of LDAP in Integrated User Firewall on page 5318
- LDAP Server Configuration and Base Distinguished Name on page 5318
- LDAP’s Authentication Method on page 5318
- LDAP Server’s Username, Password, and Server Address on page 5318
- Caching and Calculation of User-to-Group Mappings on page 5319
- Updating Group Information in the Authentication Entry Table on page 5319

Copyright © 2016, Juniper Networks, Inc.

- LDAP Server Status and Statistics on page 5319
- Active Directory Autodiscovery on page 5319

Role of LDAP in Integrated User Firewall

In order to get the user and group information necessary to implement the Integrated User Firewall feature, the SRX Series device uses the Lightweight Directory Access Protocol (LDAP). The SRX Series acts as an LDAP client communicating with an LDAP server. In a common implementation scenario of the integrated user firewall feature, the domain controller acts as the LDAP server. The LDAP module in the SRX Series, by default, queries the Active Directory in the domain controller.

The SRX Series downloads user and group lists from the LDAP server. The device also queries the LDAP server for user and group updates. The SRX Series downloads a first-level, user-to-group mapping relationship and then calculates a full user-to-group mapping.

The use of "LDAP" in this section applies specifically to LDAP functionality within the integrated user firewall feature.

LDAP Server Configuration and Base Distinguished Name

Most of the LDAP server configuration is optional, leveraging the common implementation scenario where the domain controller acts as the LDAP server. The SRX Series periodically (every two minutes) queries the LDAP server to get the user and group information changed since the last query.

The only required LDAP server configuration is the LDAP base distinguished name (DN), which is the top level of the LDAP directory tree. Microsoft Active Directory follows the convention of deriving the base DN from a company's DNS domain components. An example of a base DN is dc=juniper, dc=net.

LDAP's Authentication Method

By default, the LDAP authentication method uses simple authentication. The client's username and password are sent to the LDAP server in plaintext. Keep in mind that the password is clear and can be read from the network.

To avoid exposing the password, you can use simple authentication within an encrypted channel [namely Secure Sockets layer (SSL)], as long as the LDAP server supports LDAP over SSL (LDAPS). After enabling SSL, the data sent from the LDAP server to the SRX Series is encrypted. To enable SSL, see the **user-group-mapping** statement.

LDAP Server's Username, Password, and Server Address

The LDAP server's username, password, IP address, and port are all optional, but they can be configured.

- If the username and password are not configured, the system uses the configured domain controller's username and password.
- If the LDAP server's IP address is not configured, the system uses the address of one of the configured Active Directory domain controllers.

Copyright © 2016, Juniper Networks, Inc.

- If the port is not configured, the system uses port 389 for plaintext or port 636 for encrypted text.

Caching and Calculation of User-to-Group Mappings

The SRX Series device caches user-to-group mappings in its local database when the **show services user-identification active-directory-access user-group-mapping** operation is performed. This command displays the users who belong to a group or the groups to which a user belongs.

Three events cause a user-to-group mapping to be removed from the cache:

- A source-identity is removed from a referenced firewall policy (because only source-identities referenced in a policy are stored in the authentication table).
- The LDAP configuration is deleted from the customer's configuration, so all cached Active Directory user-to-group mappings for the domain are removed.
- The user-to-group mapping is deleted from the LDAP server.

The SRX periodically queries to get user and group information from the LDAP server in real time. The user list and the group list show only cached users or groups, not all users or groups in the LDAP server. From this information, the SRX Series calculates one-level mapping relationships. The user list, group list, and mapping are cached in the local database.

Updating Group Information in the Authentication Entry Table

The SRX Series device queries to get the changed users and groups based on the prior query results from the LDAP server. The SRX Series updates the local database and triggers an authentication entry update. Only user/group mappings that are already cached are updated. Other users and groups that are not in the database do not have their mapping relationships cached.

LDAP Server Status and Statistics

You can verify the LDAP connection status by issuing the **show services user-identification active-directory-access user-group-mapping status** command.

You can see counts of queries made to the LDAP server by issuing the **show services user-identification active-directory-access statistics user-group-mapping** command.

Active Directory Autodiscovery

The integrated user firewall feature provides the IP address and Active Directory name of the domain. The auto-discovery feature can use the Active Directory's global catalog feature and then query DNS for a list of global catalogs. The global catalogs in the list are typically provided in a weighted order based on criteria such as network location, system-set weights based on global catalog server size, and so on. Once the customer has the list of Active Directories, the customer can configure it for both event log reading and LDAP search.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Understanding How the WMIC Reads the Event Log on the Domain Controller on page 5329
- **show services user-identification active-directory-access statistics on page 5572**
- **show services user-identification active-directory-access user-group-mapping on page 5575**
- **user-group-mapping on page 5497**

Example: Configuring Integrated User Firewall

This example shows how to implement the integrated user firewall feature by configuring a Windows Active Directory domain, an LDAP base, unauthenticated users to be directed to captive portal, and a security policy based on a source identity.

- Requirements on page 5320
- Overview on page 5320
- Configuration on page 5321
- Verification on page 5326

Requirements

This example uses the following hardware and software components:

- One SRX Series device
- Junos OS Release 12.1X47-D10 or later for SRX Series devices

No special configuration beyond device initialization is required before configuring this feature.

Overview

In a typical scenario for the integrated user firewall feature, domain and non-domain users want to access the Internet through an SRX Series device. The SRX Series device reads and analyzes the event log of the domain controllers configured in the domain. Thus, the SRX Series device detects domain users on an Active Directory domain controller. Active Directory domain generates an authentication table as the Active Directory authentication source for the integrated user firewall. The SRX Series device uses this information to enforce the policy to achieve user-based or group-based access control.

For any non-domain user or domain user on a non-domain device, the network administrator can specify a captive portal to force the user to submit to firewall authentication (if the SRX Series device supports captive portal for the traffic type. For example, HTTP). After the user enters a name and password and passes firewall authentication, the SRX Series device gets firewall authentication user-to-group mapping information from the LDAP server and can enforce user firewall policy control over the user accordingly.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands to a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter **commit** from configuration mode.

```
set services user-identification active-directory-access domain example.net
  user-group-mapping ldap base DC=example,DC=net
set services user-identification active-directory-access domain example.net user
  administrator password xxxxx
set services user-identification active-directory-access domain example.net
  domain-controller ad1 address 192.0.2.15
set access profile profile1 authentication-order ldap
set access profile profile1 authentication-order password
set access profile profile1 ldap-options base-distinguished-name
  CN=Users,DC=example,DC=net
set access profile profile1 ldap-options search search-filter sAMAccountName=
set access profile profile1 ldap-options search admin-search distinguished-name
  CN=Administrator,CN=Users,DC=example,DC=net
set access profile profile1 ldap-options search admin-search password "$ABC123"
set access profile profile1 ldap-server 192.0.2.3
set security policies from-zone trust to-zone untrust policy p1 match source-address any
set security policies from-zone trust to-zone untrust policy p1 match destination-address
  any
set security policies from-zone trust to-zone untrust policy p1 match application any
set security policies from-zone trust to-zone untrust policy p1 match source-identity
  unauthenticated-user
set security policies from-zone trust to-zone untrust policy p1 match source-identity
  unknown-user
set security policies from-zone trust to-zone untrust policy p1 then permit
  firewall-authentication user-firewall access-profile profile1
set security policies from-zone trust to-zone untrust policy p2 match source-address any
set security policies from-zone trust to-zone untrust policy p2 match destination-address
  any
set security policies from-zone trust to-zone untrust policy p2 match application any
set security policies from-zone trust to-zone untrust policy p2 match source-identity
  “example.net\galenrikka”
set security policies from-zone trust to-zone untrust policy p2 match source-identity
  “example.net\domainuser_1_0000001”
set security policies from-zone trust to-zone untrust policy p2 then permit
set security user-identification authentication source active-directory-authentication-table
  priority 125
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To establish a Windows Active Directory domain, to configure captive portal, and to configure another security policy, perform the steps in this section.

Once configured, when traffic arrives, the SRX Series device consults the user firewall process, which in turn consults the Active Directory authentication source to determine whether the source is in its authentication table. If the user firewall hits an authentication entry, the SRX Series device checks the policy configured in Step 4 for further action. If the user firewall does not hit any authentication entry, the SRX Series device checks the policy configured in Step 3 to enforce the user to do captive portal.

1. Configure the LDAP base distinguished name.

```
[edit services user-identification]
user@host# set active-directory-access domain example.net user-group-mapping
  ldap base DC=example,DC=net
```

2. Configure a domain name, the username and password of the domain, and the name and IP address of the domain controller in the domain.

```
[edit services user-identification]
user@host# set active-directory-access domain example.net user administrator
  password xxxxx
user@host# set active-directory-access domain example.net domain-controller
  ad1 address 192.0.2.15
```

3. Configure an access profile and set the authentication order and LDAP options.

```
[edit access profile profile1]
user@host# set authentication-order ldap
user@host# set authentication-order password
user@host# set ldap-options base-distinguished-name
  CN=Users,DC=example,DC=net
user@host# set ldap-options search search-filter sAMAccountName=
user@host# set ldap-options search admin-search distinguished-name
  CN=Administrator,CN=Users,DC=example,DC=net
user@host# set ldap-options search admin-search password pw1593
user@host# set ldap-server 192.0.2.3
```

4. Configure a policy for the source-identity “unauthenticated-user” and “unknown-user” and enable the firewall authentication captive portal. Configuring the source identity is required in case there is no authentication sources configured, it is disconnected.

```
[edit security policies from-zone trust to-zone untrust policy p1]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set match source-identity unauthenticated-user
user@host# set match source-identity unknown-user
user@host# set then permit firewall-authentication user-firewall access-profile
  profile1
user@host#set then permit firewall-authentication user-firewall domain
  ad-userfw.net
```


Copyright © 2016, Juniper Networks, Inc.

5. Configure a second policy to enable a specific user.

```
[edit security policies from-zone trust to-zone untrust policy p2]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application any
user@host# set match source-identity “example.net\domainuser_1_0000001”
user@host# set match source-identity “example.net\galenrikka”
user@host# set then permit
```

NOTE: When you specify a source identity in a policies statement, prepend the domain name and a backslash to the group name or username. Enclose the combination in quotation marks.

6. Set the Active Directory authentication table as the authentication source for integrated user firewall information retrieval and specify the sequence in which user information tables are checked.

```
[edit security ]
user@host# set user-identification authentication source
  active-directory-authentication-table priority 125
```

NOTE: You must set the Active Directory authentication table as the authentication source for integrated user firewall information retrieval and specify the sequence in which user information tables are checked using the command **set security user-identification authentication source active-directory-authentication-table priority** *value*.

The default value of this option is 125. The default priority for all the authentication sources is as follows:

- Local authentication: 100
- Integrated user firewall: 125
- User role firewall: 150
- Unified Access Control (UAC): 200

The field **priority** specifies the sources for the Active Directory authentication table. The value set determines the sequence for searching among various supported authentication tables to retrieve a user role. Note that these are the only currently supported values. You can enter any value from 0 through 65,535. The default priority of the Active Directory authentication table is 125. This means that even if you do not specify a priority value, the Active Directory authentication table will be searched starting at sequence of value 125 (integrated user firewall).

For more details, see “Understanding Active Directory Authentication Tables” on page 5312 and **active-directory-authentication-table**.

Copyright © 2016, Juniper Networks, Inc.

(Optional) Configuration of PKI and SSL Forward Proxy to Authenticate Users

Step-by-Step Procedure

Optionally, for non-domain users, you can configure public key infrastructure (PKI) to validate integrity, confidentiality, and authenticity of traffic. PKI includes digital certificates issued by the Certificate Authority (CA), certificate validity and expiration dates, details about the certificate owner and issuer, and security policies.

NOTE: For any non-domain user or domain user on a non-domain machine, the administrator specifies a captive portal to force the user to do firewall authentication (if the SRX Series device supports captive portal for the traffic type). After the user enters a name and password and passes firewall authentication, the SRX Series device gets firewall authentication user/group information and can enforce the user firewall policy to control the user accordingly. In addition to captive portal, if the IP address or user information is not available from the event log, the user can again log in to the Windows PC to generate an event log entry. Then the system generates the user's authentication entry accordingly.

To enable the SRX Series device to authenticate the users through HTTPs, the SSL forward proxy must be configured and enabled. You need to generate a local certificate, add an SSL termination profile, add an SSL proxy profile, and reference the SSL proxy profile in the security policy. If the SSL forward proxy is not enabled, the SRX Series device cannot authenticate users who are using HTTPS, but for users who are using HTTP, FPT, and Telnet, the authentication can be performed as expected.

To generate PKI and enable SSL forward proxy, perform the following steps:

1. Generate a PKI public/private key pair for a local digital certificate.

   ```
   user@host# request security pki generate-key-pair certificate-id ssl-inspect-ca size 2048 type rsa
   ```

2. Manually generate a self-signed certificate for the given distinguished name.

   ```
   user@host# request security pki local-certificate generate-self-signed certificate-id ssl-inspect-ca domain-name www.example.com subject "CN=www.example.com,OU=IT,O=MyCompany, Sunnyvale,ST=CA,C=US" email security-admin@example.com
   ```

3. Define the access profile to be used for SSL termination services. This option is available only on high-end SRX Series devices.

   ```
   user@host# set services ssl termination profile for_userfw server-certificate ssl-inspect-ca
   ```

4. Configure the loaded certificate as root-ca in the SSL proxy profile. This option is available only on high-end SRX Series devices.

   ```
   user@host# set services ssl proxy profile ssl-inspect-profile root-ca ssl-inspect-ca
   ```


Copyright © 2016, Juniper Networks, Inc.

5. Specify the **ignore-server-auth-failure** option if you do not want to import the entire CA list and you do not want dropped sessions. This option is available only on high-end SRX Series devices.

   ```
   user@host# set services ssl proxy profile ssl-inspect-profile actions
     ignore-server-auth-failure
   ```

6. Add an SSL termination profile into security policies. This option is available only on high-end SRX Series devices.

   ```
   user@host# set security policies from-zone untrust to-zone trust policy p1 then
     permit firewall-authentication user-firewall ssl-termination-profile for_userfw
   ```

Results

From configuration mode, confirm your integrated user firewall configuration by entering the **show services user-identification active-directory-access** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show services user-identification active-directory-access
  domain example.net {
    user {
      administrator;
      password "$ABC123"; ## SECRET-DATA
    }
    domain-controller ad1 {
      address 192.0.2.15;
    }
    user-group-mapping {
      ldap {
        base DC=example,DC=net;
      }
    }
  }
```

From configuration mode, confirm your policy configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show security policies
  from-zone trust to-zone untrust {
    policy p1 {
      match {
        source-address any;
        destination-address any;
        application any;
        source-identity unauthenticated-user;
        source-identity unknown-user;
      }
      then {
        permit {
          firewall-authentication {
            user-firewall {
              access-profile profile1;
            }
```

Copyright © 2016, Juniper Networks, Inc.

```
            }
          }
        }
      }
      policy p2 {
        match {
          source-address any;
          destination-address any;
          application any;
          source-identity “example.net\galenrikka”;
          source-identity “example.net\domainuser_1_0000001”
        }
        then {
          permit;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Connectivity to a Domain Controller on page 5326
- Verifying the LDAP Server on page 5326
- Verifying Authentication Table Entries on page 5327
- Verifying IP-to-User Mapping on page 5327
- Verifying IP Probe Counts on page 5327
- Verifying User-to-Group Mapping Queries on page 5327

Verifying Connectivity to a Domain Controller

Purpose Verify that at least one domain controller is configured and connected.

Action From operational mode, enter the **show services user-identification active-directory-access domain-controller status** command.

Meaning The domain controller is shown to be connected or disconnected.

Verifying the LDAP Server

Purpose Verify that the LDAP server is providing user-to-group mapping information.

Action From operational mode, enter the **show services user-identification active-directory-access user-group-mapping status** command.

Meaning The LDAP server address, port number, and status are displayed.

Copyright © 2016, Juniper Networks, Inc.

Verifying Authentication Table Entries

Purpose See which groups users belong to and the users, groups, and IP addresses in a domain.

Action From operational mode, enter the **show services user-identification active-directory-access active-directory-authentication-table all** command.

Meaning The IP addresses, usernames, and groups are displayed for each domain.

Verifying IP-to-User Mapping

Purpose Verify that the event log is being scanned.

Action From operational mode, enter the **show services user-identification active-directory-access statistics ip-user-mapping** command.

Meaning The counts of the queries and failed queries are displayed.

Verifying IP Probe Counts

Purpose Verify that IP probes are occurring.

Action From operational mode, enter the **show services user-identification active-directory-access statistics ip-user-probe** command.

Meaning The counts of the IP probes and failed IP probes are displayed.

Verifying User-to-Group Mapping Queries

Purpose Verify that user-to-group mappings are being queried.

Action From operational mode, enter the **show services user-identification active-directory-access statistics user-group-mapping** command.

Meaning The counts of the queries and failed queries are displayed.

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Understanding the Three-Tiered User Firewall Features on page 5173
- policies on page 702
- show services user-identification active-directory-access active-directory-authentication-table on page 5565
- show services user-identification active-directory-access domain-controller status on page 5569
- show services user-identification active-directory-access statistics on page 5572

Copyright © 2016, Juniper Networks, Inc.

- show services user-identification active-directory-access user-group-mapping on page 5575

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 227

Managing Event Logs

- Understanding How the WMIC Reads the Event Log on the Domain Controller on page 5329
- Using Firewall Authentication as an Alternative to WMIC on page 5331
- Understanding Integrated User Firewall Domain PC Probing on page 5332

Understanding How the WMIC Reads the Event Log on the Domain Controller

This topic includes the following sections:

- Windows Management Instrumentation Client on page 5329
- WMIC Reads the Event Log on the Domain Controller on page 5330
- Specifying IP Filters to Limit IP-to-User Mapping on page 5330
- Event Log Verification and Statistics on page 5330

Windows Management Instrumentation Client

When you configure the integrated user firewall feature on an SRX Series device, the SRX Series establishes a Windows Management Instrumentation (WMI)/Distributed Component Object Module (DCOM) connection to the domain controller. The SRX Series acts as a WMI client (WMIC). It reads and monitors the security event log on the domain controller. The SRX Series analyzes the event messages to generate IP address-to-user mapping information.

All configuration regarding the WMIC is optional; it will function with default values. After the domain is configured (by the **set services user-identification active-directory-access domain** statement), the WMIC starts to work. The WMIC connection to the domain controller uses the same user credentials as those configured for the domain.

CAUTION: Integrated user firewall uses NTLMv2 as the default WMIC authentication protocol for security reasons. NTLMv1 exposes the system to attacks in which authentication hashes could be extracted from NTLMv1 authentication responses.

For compatibility with integrated user firewall, you must apply the latest version of the Microsoft SP2 patch if you are running an older version of Windows OS, including Windows 2000, Windows XP, and Windows 2003.

Copyright © 2016, Juniper Networks, Inc.

WMIC Reads the Event Log on the Domain Controller

The following SRX Series behaviors apply to reading the event log:

- The SRX Series monitors the event log at a configurable interval, which defaults to 10 seconds.
- The SRX Series reads the event log for a certain timespan, which you can configure. The default timespan is one hour. Each time at WMIC startup, the SRX Series checks the last timestamp and the timespan. If the last timestamp is older than the current timespan, then the timespan takes effect. After the WMIC and the UserID process start working, the timespan does not apply; the SRX Series simply reads the latest event log.
- During WMIC startup, the SRX Series has a maximum count of events it will read from the event log, and that maximum is not configurable.
  - On SRX Series branch devices, the maximum count is 100,000.
  - On SRX high-end devices, the maximum count is 200,000.

  During WMIC startup, this maximum is used with the timespan setting, so that if either limit is reached, the WMIC stops reading the event log.
- After a failover, the SRX Series reads the event log from the latest event log timestamp.
- In a chassis cluster environment, the WMIC works on the primary node only.

Specifying IP Filters to Limit IP-to-User Mapping

You can specify IP filters to limit the IP address-to-user mapping information that the SRX Series generates from the event log.

To understand when a filter is useful for such mapping, consider the following scenario. A customer deploys 10 SRX Series devices in one domain, and each SRX Series controls a branch. All 10 SRX Series devices read all 10 branch user login event logs in the domain controller. However, the SRX Series is configured to detect only whether the user is authenticated on the branch it controls. By configuring an IP filter on the SRX Series, the SRX Series reads only the IP event log under its control.

You can configure a filter to include or exclude IP addresses or prefixes. You can specify a maximum of 20 addresses for each filter.

Event Log Verification and Statistics

You can verify that the authentication table is getting IP address and user information by issuing the **show services user-identification active-directory-access active-directory-authentication-table all** command. A list of IP address-to-user mappings is displayed for each domain. The table contains no group information until LDAP is running.

You can see statistics about reading the event log by issuing the **show services user-identification active-directory-access ip-user-mapping statistics domain** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- show services user-identification active-directory-access active-directory-authentication-table on page 5565
- Overview of Integrated User Firewall on page 5309
- LDAP Functionality in Integrated User Firewall on page 5317
- Using Firewall Authentication as an Alternative to WMIC on page 5331

Using Firewall Authentication as an Alternative to WMIC

This topic includes the following sections:

- WMIC Limitations on page 5331
- Firewall Authentication as a Backup Method for IP Address-to-User Mappings on page 5331

WMIC Limitations

The primary method for the integrated user firewall feature to get IP address-to-user mapping information is for the SRX Series device to act as a WMI client (WMIC). However, the WMIC has limitations, such as the following:

- On Windows XP or Server2003, the Windows firewall does not allow the WMIC request to pass through because of the dynamic port allocation of the Distributed Component Object Model (DCOM). Therefore, for these operating systems when Windows firewall is enabled, the PC does not respond to the WMIC probe.
- Because the event-log-reading and PC probe functions both use WMI, using a global policy to disable the WMI-to-PC probe also affects event log reading.

Because these cases might result in the failure of the PC probe, a backup method for getting IP address-to-user mappings is needed. That method is to use firewall authentication to identify users.

Firewall Authentication as a Backup Method for IP Address-to-User Mappings

If you want to use firewall authentication to identify users for the integrated user firewall feature, specify a domain name in the **set security policies from-zone trust to-zone untrust policy <policy-name> then permit firewall-authentication user-firewall domain <domain-name>** statement.

If a domain is configured in that statement, fwauth recognizes that the domain is for a domain authentication entry, and will send the domain name to the fwauth process along with the authentication request. After it receives the authentication response, fwauth deletes that domain authentication entry. The fwauth process sends the source IP address, username, domain, and other information to the USERID process, which verifies that it is a valid domain user entry. The subsequent traffic will hit this user firewall entry.

NOTE: The Active Directory authentication entry that comes from the fwauth process is not subject to the IP filters.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- user-firewall on page 1330
- Overview of Integrated User Firewall on page 5309
- Understanding How the WMIC Reads the Event Log on the Domain Controller on page 5329

Understanding Integrated User Firewall Domain PC Probing

This topic includes the following sections:

- Overview of Domain PC Probing on page 5332
- Probing Domain PCs for User Information on page 5332
- Probe Response on page 5333
- Probe Configuration on page 5334
- Probe Rate and Statistics on page 5334

Overview of Domain PC Probing

At a high level, the integrated user firewall feature gathers IP address, user, and group information from Windows Active Directory domain controller event logs and LDAP services. This information is used to generate Active Directory authentication table entries on an SRX Series device. Authentication entries serve as the authentication source for security policies that enforce user-based or group-based access control.

PC probing acts as a supplement of event log reading. When a user logs in to the domain, the event log contains that information. The PC probe is triggered only when there is no IP-to-address mapping from the event log.

Domain information constantly changes as users log in and out of domain PCs. The integrated user firewall probe functionality provides a mechanism for tracking and verifying information in the authentication tables by directly probing domain PCs for IP address-to-user mapping information. New and changed information identified by the probe serves to update Active Directory authentication table entries, which is critical to maintaining firewall integrity.

The IP address filter also impacts the PC probe. Once you configure the IP address filter, only the IP address specified in the filter is probed.

Probing Domain PCs for User Information

The integrated user firewall feature tracks the online status of users by probing domain PCs. If a user is not online or is not an expected user, the Active Directory authentication table is updated as appropriate. The following probe behaviors apply:

On-demand probing—On-demand probing occurs when a packet is dropped due to a missing entry in the Active Directory authentication table. In this case, an entry is added in pending state to the authentication table, and the domain PC identified by the source IP field of the dropped packet is probed for IP address and user information. The entry remains in pending state until a response is received from the probe.

Copyright © 2016, Juniper Networks, Inc.

Manual probing—Manual probing is used to verify and troubleshoot the online status of a user or a range of users, and is at the discretion of the system administrator. To initiate a manual probe, use the **request services user-identification active-directory-access ip-user-probe address ip-address** *address* **domain** *domain-name* command. If a domain name is not specified, the probe looks at the first configured domain for the IP address. To specify a range, use the appropriate network address.

NOTE: Manual probing can cause entries to be removed from the Active Directory authentication table. For example, if there is no response from your PC due to a network issue, such as when the PC is too busy, the IP address entry of the PC is marked as *invalid* and your access is blocked.

If the SRX Series device cannot access a domain PC for some reason, such as a network configuration or Windows firewall issue, the probe fails.

Probe Response

Based on the domain PC probe response, updates are made to the Active Directory authentication table, and associated firewall policies take effect. If no response is received from the probe after 90 seconds, the authentication entry times out. The timed-out authentication entry is the pending state authentication entry, which is generated when you start the PC probe.

If the probe is successful, the state of the authentication entry is updated from pending to valid. If the probe is unsuccessful, the state of the authentication entry is marked as invalid. The invalid entry has the same lifetime as a valid entry and is overwritten by upcoming fwauth (firewall authentication process) authentication results or by the event log. Table 473 lists probe responses and corresponding authentication table actions.

Table 473: Probe Responses and Associated Active Directory Authentication Table Actions

Probe Response from Domain PC	Active Directory Authentication Table Action
Valid IP address and username	Add IP-related entry.
Logged on user changed	Update IP-related entry.
Connection timeout	Update IP-related entry as invalid.
Access denied	Update IP-related entry as invalid.
Connection refused	Update IP-related entry as invalid.
Authentication failed (The configured username and password have no privilege to probe the domain PC.)	Update IP-related entry as invalid.

Copyright © 2016, Juniper Networks, Inc.

Probe Configuration

On-demand probing is enabled by default. To disable on-demand probing, use the **set services user-identification active-directory-access no-on-demand-probe** statement. Delete this statement to reenable probing. When on-demand probing is disabled, manual probing is available.

The probe timeout value is configurable. The default timeout is 10 seconds. To configure the timeout value, use the following statement:

user@host# **set services user-identification active-directory-access wmi-timeout** ***seconds***

If no response is received from the domain PC within the **wmi-timeout** interval, the probe fails and the system either creates an invalid authentication entry or updates the existing authentication entry as invalid. If an authentication table entry already exists for the probed IP address, and no response is received from the domain PC within the **wmi-timeout** interval, the probe fails and that entry is deleted from the table.

NOTE: To probe domain PCs, you must configure the integrated user firewall feature with the username and password credentials. You do not necessarily need a username and password account for each PC; instead you could set up one administrator account with privileges to access information on multiple PCs.

Probe Rate and Statistics

The maximum probe rate for the integrated user firewall feature is set by default and cannot be changed. For high-end SRX Series devices, the probe rate is 600 times per minute. For branch SRX Series devices, the probe rate is 100 times per minute. Probe functionality supports 5000 users, or up to 10 percent of the total supported authentication entries, whichever is smaller. Supporting 10 percent means that at any time, the number of IP addresses waiting to be probed cannot exceed 10 percent. For more information about the number of supported Active Directory authentication table entries, see "Understanding Active Directory Authentication Tables" on page 5312.

High-level statistics covering probe activity are available for the total number of probes and the number of failed probes. Table 473 describes the reasons for probe failures. To display probe statistics, use the **show services user-identification active-directory-access statistics ip-user-probe** command.

```
user@host> show services user-identification active-directory-access statistics ip-user-probe
Domain: www.example1.net
    Total user probe number                     : 176116
    Failed user probe number                    : 916

Domain: www.example2.net
    Total user probe number                     : 17632
    Failed user probe number                    : 342
```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Overview of Integrated User Firewall on page 5309
- Understanding How the WMIC Reads the Event Log on the Domain Controller on page 5329
- Understanding Active Directory Authentication Tables on page 5312

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 70

Configuration Statements and Operational Commands

- Configuration Statements on page 5339
- Operational Commands on page 5511

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 228

Configuration Statements

- Access Configuration Statement Hierarchy on page 5342
- Security Configuration Statement Hierarchy on page 5351
- Services Configuration Statement Hierarchy on page 5352
- System Configuration Statement Hierarchy on page 5360
- [edit security firewall-authentication] Hierarchy Level on page 5390
- [edit security policies] Hierarchy Level on page 5391
- [edit security user-identification] Hierarchy Level on page 5395
- [edit security zones] Hierarchy Level on page 5395
- actions (Services SSL Proxy) on page 5398
- active-directory-access on page 5400
- active-directory-authentication-table on page 5402
- address (Services) on page 5403
- admin-search on page 5403
- application (Security Policies) on page 5404
- application-services (Security Policies) on page 5405
- assemble on page 5406
- authentication-source on page 5406
- banner (Access FTP HTTP Telnet Authentication) on page 5407
- banner (Access Web Authentication) on page 5407
- base-distinguished-name on page 5408
- ca-profile (Services) on page 5408
- captive-portal (Services UAC) on page 5409
- captive-portal (Services UAC Policy) on page 5409
- certificate-verification on page 5410
- client-group on page 5411
- client-idle-timeout (Access Profile) on page 5411
- client-name-filter on page 5412
- client-session-timeout (Access Profile) on page 5412

Copyright © 2016, Juniper Networks, Inc.

- configuration-file on page 5413
- count on page 5413
- custom-ciphers on page 5414
- default-profile on page 5414
- distinguished-name (Access) on page 5415
- domain-name (Access Profile) on page 5415
- enable-flow-tracing (Services) on page 5416
- enable-session-cache on page 5416
- fail on page 5417
- file (Services) on page 5417
- files (Services) on page 5418
- file (System Logging) on page 5419
- firewall-authentication on page 5422
- firewall-authentication (Security) on page 5423
- firewall-authentication (Security Policies) on page 5424
- firewall-authentication (User Identification) on page 5425
- firewall-authentication-service on page 5425
- firewall-user on page 5426
- flag (Services) on page 5426
- from-zone (Security Policies) on page 5427
- ftp (Access) on page 5429
- group-profile (Access) on page 5430
- http (Access) on page 5431
- http (Services) on page 5432
- https (Services) on page 5433
- infranet-controller on page 5434
- interface (Services) on page 5435
- interval (Services) on page 5435
- ip-address (Access Profile) on page 5436
- ip-user-mapping on page 5437
- ldap-options on page 5438
- ldap-server on page 5439
- level (Services) on page 5439
- lifetime-seconds (Security IKE) on page 5440
- link (Access) on page 5440
- local-authentication-table on page 5441
- log (Services) on page 5442

Copyright © 2016, Juniper Networks, Inc.

- login (Access) on page 5443
- match (Services) on page 5443
- network (Access) on page 5444
- no-remote-trace (Services) on page 5444
- pass-through on page 5445
- password (Access) on page 5446
- password (Services) on page 5446
- permit (Security Policies) on page 5447
- policies on page 5449
- pool (Access) on page 5454
- port (Access LDAP) on page 5455
- port (Services) on page 5456
- preferred-ciphers on page 5456
- prefix (Access IPv6) on page 5457
- protocol-version on page 5457
- radius-options (Access) on page 5458
- radius-server (Access) on page 5459
- range (Access) on page 5460
- redirect-traffic on page 5461
- redirect-url on page 5462
- retry (Access LDAP) on page 5463
- retry (Access RADIUS) on page 5463
- revert-interval (Access LDAP) on page 5464
- revert-interval (Access RADIUS) on page 5464
- root-ca (Services) on page 5465
- routing-instance (Access LDAP) on page 5465
- routing-instance (Access RADIUS) on page 5466
- search on page 5466
- search-filter on page 5467
- secret (Access Profile) on page 5467
- securid-server on page 5468
- separator on page 5469
- server-certificate (Services) on page 5469
- server-certificate-subject on page 5470
- session-options (Access Profile) on page 5470
- size (Services) on page 5471
- source-address (Access LDAP) on page 5471

Copyright © 2016, Juniper Networks, Inc.

- source-address (Access RADIUS) on page 5472
- ssl (Services) on page 5473
- ssl-termination-profile on page 5475
- success on page 5475
- telnet (Access) on page 5476
- termination (Services) on page 5477
- test-only-mode on page 5477
- then (Security Policies) on page 5478
- timeout (Access LDAP) on page 5480
- timeout (Access RADIUS) on page 5480
- timeout (Services) on page 5481
- timeout-action on page 5482
- to-zone (Security Policies) on page 5483
- traceoptions (Access) on page 5486
- traceoptions (Active Directory Access) on page 5488
- traceoptions (Security Firewall Authentication) on page 5490
- traceoptions (Services SSL) on page 5491
- traceoptions (Services UAC) on page 5492
- trusted-ca (Services) on page 5493
- uac-policy (Application Services) on page 5493
- uac-service on page 5494
- unified-access-control (Security) on page 5495
- unified-access-control (Services) on page 5496
- user-group-mapping on page 5497
- user-identification (Services) on page 5499
- web-authentication on page 5501
- web-authentication (Access) on page 5502
- web-authentication (Interfaces) on page 5503
- web-management (System Services) on page 5504
- web-redirect on page 5507
- web-redirect-to-https on page 5508
- whitelist (Services) on page 5508
- wins-server (Access) on page 5509

Access Configuration Statement Hierarchy

Use the statements in the **access** configuration hierarchy to configure access to the device and authentication methods, including address assignment and address pool, user and firewall authentication, a group profile, LDAP options and LDAP server configuration, an

Copyright © 2016, Juniper Networks, Inc.

access profile, RADIUS options and RADIUS server configuration, and SecurID server configuration.

```
access {
  address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-name;
    pool pool-name {
      family {
        inet {
          dhcp-attributes {
            boot-file boot-file-name;
            boot-server boot-server-name;
            domain-name domain-name;
            grace-period seconds;
            maximum-lease-time (seconds | infinite);
            name-server ipv4-address;
            netbios-node-type (b-node | h-node | m-node | p-node);
            next-server next-server-name;
            option dhcp-option-identifier-code {
              array {
                byte [8-bit-value];
                flag [ false| off |on |true];
                integer [32-bit-numeric-values];
                ip-address [ip-address];
                short [signed-16-bit-numeric-value];
                string [character string value];
                unsigned-integer [unsigned-32-bit-numeric-value];
                unsigned-short [16-bit-numeric-value];
              }
              byte 8-bit-value;
              flag (false | off | on | true);
              integer 32-bit-numeric-values;
              ip-address ip-address;
              short signed-16-bit-numeric-value;
              string character string value;
              unsigned-integer unsigned-32-bit-numeric-value;
              unsigned-short 16-bit-numeric-value;
            }
            option-match {
              option-82 {
                circuit-id match-value {
                  range range-name;
                }
                remote-id match-value;
                  range range-name;
                }
              }
            }
            propagate-ppp-settings [interface-name];
            propagate-settings interface-name;
            router ipv4-address;
            server-identifier ip-address;
```

```
        sip-server {
            ip-address ipv4-address;
            name sip-server-name;
        }
        tftp-server server-name;
        wins-server ipv4-address;
    }
    host hostname {
        hardware-address mac-address;
        ip-address reserved-address;
    }
    network network address;
    range range-name {
        high upper-limit;
        low lower-limit;
    }
    xauth-attributes {
        primary-dns ip-address;
        primary-wins ip-address;
        secondary-dns ip-address;
        secondary-wins ip-address;
    }
}
inet6 {
    dhcp-attributes {
        dns-server ipv6-address;
        grace-period seconds;
        maximum-lease-time (seconds | infinite);
        option dhcp-option-identifier-code {
            array {
                byte [8-bit-value];
                flag [ false| off |on |true];
                integer [32-bit-numeric-values];
                ip-address [ip-address];
                short [signed-16-bit-numeric-value];
                string [character string value];
                unsigned-integer [unsigned-32-bit-numeric-value];
                unsigned-short [16-bit-numeric-value];
            }
            byte 8-bit-value;
            flag (false | off | on | true);
            integer 32-bit-numeric-values;
            ip-address ip-address;
            short signed-16-bit-numeric-value;
            string character string value;
            unsigned-integer unsigned-32-bit-numeric-value;
            unsigned-short 16-bit-numeric-value;
        }
        propagate-ppp-settings [interface-name];
        sip-server-address ipv6-address;
        sip-server-domain-name domain-name;
    }
    prefix ipv6-network-prefix;
    range range-name {
        high upper-limit;
        low lower-limit;
```


Copyright © 2016, Juniper Networks, Inc.

```
            prefix-length delegated-prefix-length;
          }
        }
      link pool-name;
    }
  }
  address-pool pool-name {
    (address address-or-address-prefix ) {
    address-range {
      high upper-limit;
      low lower-limit;
      mask network-mask;
    }
    primary-dns name;
    primary-wins name;
    secondary-dns name;
    secondary-wins name;
  }
  address-protection;
  domain {
    delimiter delimiter;
    map domain-map-name {
      aaa-logical-system logical-system-name;
      aaa-routing-instance routing-instance-name;
      access-profile access-profile-name;
      address-pool address-pool-name;
      dynamic-profile dynamic-profile-name;
      padn destination-address; {
        mask destination-mask;
        metric metric-value
      }
      strip-domain;
      target-logical-system logical-system-name;
      target-routing-instance target-routing-instance;
    }
    parse-direction (left-to-right | right-to-left);
  }
  firewall-authentication {
    pass-through {
      default-profile profile-name;
      ftp {
        banner {
          fail string;
          login string;
          success string;
        }
      }
      http {
        banner {
          fail string;
          login string;
          success string;
        }
      telnet {
        banner {
          fail string;
```

```
            login string;
            success string;
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            flag flag;
            match regular-expression;
            no-remote-trace;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
    }
    web-authentication {
        banner {
            success string;
        }
        default-profile profile-name;
    }
}
group-profile profile-name {
    ppp {
        cell-overhead;
        encapsulated-overhead encapsulated-overhead-value;
        framed-pool address-pool-name;
        idle-timeout seconds;
        interface-id interface-identifier;
        keepalive seconds;
        ppp-options {
            chap;
            pap;
        }
        primary-dns name;
        primary-wins name;
        secondary-dns name;
        secondary-wins name;
    }
}
gx-plus {
    global {
        max-outstanding-requests max-outstanding-requests;
    }
    partition partition-name {
        destination-host gx-plus-destination-host;
        destination-realm gx-plus-destination-realm;
        diameter-instance gx-plus-diameter-instance;
    }
}
ldap-options {
    assemble {
        common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
```


Copyright © 2016, Juniper Networks, Inc.

```
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter filter-name;
    }
  }
  ldap-server hostname-or-address; {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  ppp-options {
    compliance {
      rfc(2486 | [rfc-number]);
    }
  }
  profile profile-name  {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      coa-immediate-update;
      duplication;
      immediate-update;
      order [accounting-method];
      statistics (time | volume-time);
      update-interval minutes;
    }
    accounting-order [accounting-method];
    address-assignment pool pool-name;
    authentication-order [ldap | none | password | radius | securid];
    authorization-order [jsrc];
    client client-name  {
      chap-secret  chap-secret;
      client-group [ group-names ];
      firewall-user {
        password password;
      }
      no-rfc2486;
      pap-password pap-password;
      x-auth ip-address;
    }
    client-name-filter {
      count number;
      domain-name domain-name;
      separator special-character;
    }
    ldap-options {
      assemble {
        common-name common-name;
      }
      base-distinguished-name base-distinguished-name;
      revert-interval seconds;
```

```
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
  radius {
    accounting-server [server];
    attributes {
      exclude {
        acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
        acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
        acc-loop-cir-id [access-request | accounting-start | accounting-stop];
        accounting-authentic [accounting-off | accounting-on | accounting-start |
          accounting-stop];
        accounting-delay-time [accounting-off | accounting-on | accounting-start |
          accounting-stop];
        accounting-session-id [access-request];
        accounting-terminate-cause [accounting-off];
        act-data-rate-dn [access-request | accounting-start | accounting-stop];
        act-data-rate-up [access-request | accounting-start | accounting-stop];
        act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
        act-interlv-delay-up [access-request | accounting-start | accounting-stop];
        att-data-rate-dn [access-request | accounting-start | accounting-stop];
        att-data-rate-up [access-request | accounting-start | accounting-stop];
        called-station-id [access-request | accounting-start | accounting-stop];
        calling-station-id [access-request | accounting-start | accounting-stop];
        class [access-request | accounting-start | accounting-stop];
        delegated-ipv6-prefix [accounting-start | accounting-stop];
        dhcp-gi-address [access-request | accounting-start | accounting-stop];
        dhcp-mac-address [access-request | accounting-start | accounting-stop];
        dhcp-options [access-request | accounting-start | accounting-stop];
        downstream-calculated-qos-rate [access-request | accounting-start |
          accounting-stop];
        dsl-forum-attributes [access-request | accounting-start | accounting-stop];
        dsl-line-state [access-request | accounting-start | accounting-stop];
        dsl-type [access-request | accounting-start | accounting-stop];
        dynamic-iflset-name [accounting-start | accounting-stop];
        event-time-stamp [accounting-off | accounting-on | accounting-start |
          accounting-stop];
        framed-interface-id [access-request | accounting-start | accounting-stop];
        framed-ip-address [access-request | accounting-start | accounting-stop];
        framed-ip-netmask [access-request | accounting-start | accounting-stop];
        framed-ip-route [access-request | accounting-start | accounting-stop];
        framed-ipv6-pool [accounting-start | accounting-stop];
        framed-ipv6-prefix [accounting-start | accounting-stop];
```


Copyright © 2016, Juniper Networks, Inc.

```
            framed-ipv6-route [accounting-start | accounting-stop];
            framed-pool [accounting-start | accounting-stop];
            input-filter [accounting-start | accounting-stop];
            input-gigapackets [accounting-stop];
            input-gigawords [accounting-stop];
            input-ipv6-gigawords [accounting-stop];
            input-ipv6-octets [accounting-stop];
            input-ipv6-packets [accounting-stop];
            interface-description [access-request | accounting-start | accounting-stop];
            l2c-downstream-data [access-request | accounting-start | accounting-stop];
            l2c-upstream-data [access-request | accounting-start | accounting-stop];
            max-data-rate-dn [access-request | accounting-start | accounting-stop];
            max-data-rate-up [access-request | accounting-start | accounting-stop];
            max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
            max-interlv-delay-up [access-request | accounting-start | accounting-stop];
            min-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-data-rate-up [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
            nas-identifier [access-request | accounting-start | accounting-stop];
            nas-port [access-request | accounting-off | accounting-on | accounting-start |
              accounting-stop];
            nas-port-id [access-request | accounting-start | accounting-stop];
            nas-port-type [access-request | accounting-start | accounting-stop];
            output-filter [accounting-start | accounting-stop];
            output-gigapackets [accounting-stop];
            output-gigawords [accounting-stop];
            output-ipv6-gigawords [accounting-stop];
            output-ipv6-octets [accounting-stop];
            output-ipv6-packets [accounting-stop];
            upstream-calculated-qos-rate [access-request | accounting-start |
              accounting-stop];
        }
        ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
        }
    }
authentication-server [server];
radius-options {
    request-rate number;
    revert-interval seconds;
}
radius-server server-address {
    accounting-port port-number
    max-outstanding-requests number-of--outstanding-requests;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
    service {
      accounting-order {
        activation-protocol;
        radius;
      }
    }
    session-options {
      client-group [group-name];
      client-idle-timeout minutes;
      client-session-timeout minutes;
    }
  }
  radius-options {
    request-rate number;
    revert-interval seconds;
  }
  radius-server server-address {
    accounting-port port-number;
    max-outstanding-requests number-of-max-outstanding-requests;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
  }
  securid-server server-name {
    configuration-file filepath;
  }
  terminate-code {
    aaa {
      deny {
        authentication-denied {
          radius acct-terminate-cause-value;
        }
        no-resources {
          radius acct-terminate-cause-value;
        }
        server-request-timeout {
          radius acct-terminate-cause-value;
        }
      }
      shutdown {
        administrative-reset {
          radius acct-terminate-cause-value;
        }
        remote-reset {
          radius acct-terminate-cause-value;
        }
      }
    }
    dhcp {
      client-request {
        radius acct-terminate-cause-value;
      }
      lost-carrier {
```


Copyright © 2016, Juniper Networks, Inc.

```
            radius acct-terminate-cause-value;
        }
        nak {
            radius acct-terminate-cause-value;
        }
        nas-logout {
            radius acct-terminate-cause-value;
        }
        no-offers {
            radius acct-terminate-cause-value;
        }
    }
  }
}
```

Related Documentation

- *Understanding User Authentication Methods*
- Layer 2 Transparent and Switching Overview on page 2893
- Understanding User Authentication for Security Devices on page 5173

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837

Copyright © 2016, Juniper Networks, Inc.

- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

Services Configuration Statement Hierarchy

Use the statements in the **services** configuration hierarchy to configure application identification, probes, and Unified Access Control.

```
services {
  application-identification {
    application application-name {
      address-mapping address-name {
        filter {
          ip ip-address-and-prefix-length;
          port-range {
            tcp [port];
            udp [port];
          }
        }
      }
      cacheable;
```


Copyright © 2016, Juniper Networks, Inc.

```
        description;
        icmp-mapping {
            code number;
            type number;
        }
        ip-protocol-mapping {
            protocol number;
        }
        order number;
        over protocol-type ;
        priority [high | low];
    }
    application-group group-name {
        application-groups application-group-name;
        applications application-name;
    }
    application-system-cache-timeout value;
    download {
        automatic {
            interval hours;
            start-time  MM-DD.hh:mm;
        }
        url url;
    }
    enable-performance-mode max-packet-threshold number;
    no-application-identification;
    no-application-system-cache;
    statistics {
        interval minutes;
    }
    traceoptions {
        file {
            filename ;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        level [all | error | info | notice | verbose | warning]
        no-remote-trace;
    }
}
captive-portal {
    authentication-profile-name authentication-profile-name;
    custom-options {
        banner-message string;
        footer-bgcolor hex-color-value;
        footer-message string;
        footer-text-color hex-color-value;
        form-header-bgcolor hex-color-value;
        form-header-message string;
        form-header-text-color hex-color-value;
        form-reset-label label name;
        form-submit-label label name;
        header-bgcolor hex-color-value;
```

Copyright © 2016, Juniper Networks, Inc.

```
        header-logo filename;
        header-message string;
        header-text-color hex-color-value;
        post-authentication-url url-string;
      }
      interface (all | interface-name) {
        quiet-period seconds;
        retries number-of-retries;
        server-timeout seconds;
        session-expiry seconds;
        supplicant (multiple | single | single-secure);
      }
      secure-authentication (http | https);
      traceoptions {
        file {
          filename ;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
      }
    }
    flow-monitoring {
      version9 {
        template template-name {
          flow-active-timeout seconds;
          flow-inactive-timeout seconds;
          ipv4-template;
          ipv6-template;
          option-refresh-rate {
            packets packets;
            seconds seconds;
          }
          template-refresh-rate {
            packets packets;
            seconds seconds;
          }
        }
      }
    }
    ip-monitoring {
      policy policy-name {
        match {
          rpm-probe [probe-name];
        }
        no-preempt ;
        then {
          interface interface-name (disable | enable);
          preferred-route {
            route destination-address {
              next hop next-hop;
              preferred-metric metric;
              }
            routing-instances name;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  rpm {
    bgp {
      data-fill data;
      data-size size;
      destination-port port;
      history-size size;
      logical-system logical-system-name <routing-instances routing-instance-name>;
      moving-average-size number-of-samples;
      probe-count count;
      probe-interval seconds;
      probe-type type;
      routing-instances {
        routing-instance-name;
      }
      test-interval seconds;
    }
    probe owner {
      test test-name {
        data-fill data;
        data-size size;
        destination-interface interface-name;
        destination-port port;
        dscp-code-point dscp-bits;
        hardware-timestamp;
        history-size size;
        inet6-options {
          source-address address;
        }
        moving-average-size number;
        next-hop next-hop;
        one-way-hardware-timestamp;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instance instance-name;
        source-address address;
        target {
          address ipv4-address;
          url url;
          inet6-address ipv6-address;
          inet6-url url;
```

```
        }
        test-interval interval;
        thresholds {
          egress-time microseconds;
          ingress-time microseconds;
          jitter-egress microseconds;
          jitter-ingress microseconds;
          jitter-rtt microseconds;
          rtt microseconds;
          std-dev-egress microseconds;
          std-dev-ingress microseconds;
          std-dev-rtt microseconds;
          successive-loss count;
          total-loss count;
        }
        traps [ trap-names];
      }
    }
  probe-limit number;
    probe-server {
      icmp {
        destination-interface interface-name;
      }
      tcp {
        destination-interface interface-name;
        port port-number;
      }
      udp {
        destination-interface interface-name;
        port port-number;
      }
    }
  service-device-pools {
    pool pool-name {
      interface service-device-name;
    }
  }
  service-interface-pools {
    pool pool-name {
      interface service-interface-name;
    }
  }
  ssl {
    initiation {
      profile profile-name {
        actions {
          ignore-server-auth-failure;
        }
        client-certificate;
        custom-ciphers [cipher];
        enable-flow-tracing;
        enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | tls1 | tls11 | tls12);
        trusted-ca (all | [ca-profile] );
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
proxy {
  global-config {
    session-cache-timeout seconds;
  }
  profile profile-name {
    actions {
      crl {
        disable;
        if-not-present (allow | drop);
        ignore-hold-instruction-code;
      }
      disable-session-resumption;
      ignore-server-auth-failure;
      logs {
        all;
        errors;
        info;
        sessions-allowed;
        sessions-dropped;
        sessions-ignored;
        sessions-whitelisted;
        warning;
      }
      renegotiation {
        (allow | allow-secure | drop);
      }
    }
    custom-ciphers [cipher];
    enable-flow-tracing;
    preferred-ciphers (custom | medium | strong | weak);
    root-ca root-certificate;
    trusted-ca (all | [ca-profile] );
    whitelist [global-address-book-addresses];
  }
}
termination {
  profile profile-name {
    custom-ciphers [cipher];
    enable-flow-tracing;
    enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | tls1 | tls11 | tls12);
    server-certificate certificate-identifier;
  }
}
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  level [brief | detail | extensive | verbose];
```

Copyright © 2016, Juniper Networks, Inc.

```
    no-remote-trace;
  }
}
unified-access-control {
  captive-portal redirect-policy-name {
    redirect-traffic (all | unauthenticated);
    redirect-url redirect-url;
  }
  certificate-verification [ optional | required | warning ];
  infranet-contoller host-name {
    address ip-address;
    ca-profile [ca-profile];
    interface interface-name;
    password password;
    port port-number;
    server-certificate-subject subject;
  }
  interval seconds;
  test-only-mode;
  timeout seconds;
  timeout-action (close | no-change | open);
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
user-identification {
  active-directory-access {
    domain domain-name {
      user username;
        password password;
      domain-controller domain-controller-name {
        address domain-controller-address;
      }
      ip-user-mapping {
        discovery-method {
          wmi {
            event-log-scanning-interval seconds;
            initial-event-log-timespan hours;
          }
        }
      }
      user-group-mapping {
        ldap {
          address ip-address {
            port port;
          }
          authentication-algorithm {
            simple;
```


Copyright © 2016, Juniper Networks, Inc.

```
                }
                base base;
                ssl;
                user username {
                  password password;
                }
              }
            }
          }
          authentication-entry-timeout minutes;
          filter {
            include address;
            exclude address;
          }
          no-on-demand-probe;
          wmi-timeout seconds;
          traceoptions {
            file file;
            flag {
              active-directory-authentication;
              all;
              configuration;
              db;
              ip-user-mapping;
              ip-user-probe;
              ipc;
              user-group-mapping;
              wmic;
            }
            level {
              all;
              error;
              info;
              notice;
              verbose;
              warning;
            }
            no-remote-trace;
          }
        }
      wireless-wan {
        adapter adapter-name {
          adapter-type cx-bridge;
          ip-address ip-address;
          modem {
            usb1 description description;
            usb2 description description;
            usb3 description description;
          }
        }
      }
    }
```

Related Documentation

- Understanding AppSecure Services on page 483
- Understanding Unified Access Control on page 5269

Copyright © 2016, Juniper Networks, Inc.

System Configuration Statement Hierarchy

Use the statements in the **system** configuration hierarchy to configure system management functions including addresses of the Domain Name System (DNS) servers; device's hostname, address, and domain name; health monitoring; interface filtering; properties of the device's auxiliary and console ports; security profiles for logical systems; time zones and Network Time Protocol (NTP) properties; trace options; and user login accounts, including user authentication and the root-level user account. Statement descriptions that are exclusive to the SRX Series devices running Junos OS are described in this section.

```
system {
  accounting {
    destination {
      radius {
        server server-address {
          accounting-port port-number;
          max-outstanding-requests number;
          port number;
          retry number;
          secret password;
          source-address address;
          timeout seconds;
        }
      }
      tacplus {
        server server-address {
          port port-number;
          secret password;
          single-connection;
          source-address source-address;
          timeout seconds;
        }
      }
    }
    events [change-log interactive-commands login];
    traceoptions {
      file {
        filename;
        files number;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  allow-v4mapped-packets;
  archival {
    configuration {
      archive-sites url {
        password password;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
      transfer-interval interval;
      transfer-on-commit;
    }
  }
  arp {
    aging-timer minutes;
    gratuitous-arp-delay seconds;
    gratuitous-arp-on-ifup;
    interfaces {
      interface name {
        aging-timer minutes;
      }
    }
    passive-learning;
    purging;
  }
  authentication-order [password radius tacplus];
  auto-configuration {
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  auto-snapshot;
  autoinstallation {
    configuration-servers {
      url {
        password password;
      }
    }
    interfaces {
      interface-name {
        bootp;
        rarp;
      }
    }
    usb {
      disable;
    }
  }
  auto-snapshot;
  backup-router {
    address;
    destination [network];
  }
  commit {
    server {
      commit-interval seconds;
```

```
    days-to-keep-error-logs days;
    maximum-aggregate-pool number;
    maximum entries number;
    traceoptions {
      file {
        filename;
        files number;
        microsecond-stamp;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  synchronize;
}
compress-configuration-files;
default-address-selection;
diag-port-authentication {
  encrypted-password passsword;
  plain-text-password;
}
domain-name domain-name;
domain-search [domain-list];
donot-disable-ip6op-ondad;
dump-device (boot-device | compact-flash | usb);
dynamic-profile-options {
  versioning;
}
encrypt-configuration-files;
extensions {
  providers {
    provider-id {
      license-type license deployment-scope [deployments];
    }
  }
  resource-limits {
    package package-name {
      resources {
        cpu {
          priority number;
          time seconds;
        }
        file {
          core-size bytes;
          open number;
          size bytes;
        }
        memory {
          data-size mbytes;
          locked-in mbytes;
          resident-set-size mbytes;
          socket-buffers mbytes;
          stack-size mbytes;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      process process-ui-name {
        resources {
          cpu {
            priority number;
            time seconds;
          }
          file {
            core-size bytes;
            open number;
            size bytes;
          }
          memory {
            data-size mbytes;
            locked-in mbytes;
            resident-set-size mbytes;
            socket-buffers mbytes;
            stack-size mbytes;
          }
        }
      }
    }
  }
  fips {
    level (0 | 1 | 2 | 3 | 4);
  }
  host-name hostname;
  inet6-backup-router {
    address;
    destination destination;
  }
  internet-options {
    icmpv4-rate-limit {
      bucket-size seconds;
      packet-rate packets-per-second;
    }
    icmpv6-rate-limit {
      bucket-size seconds;
      packet-rate packets-per-second;
    }
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    ipv6-duplicate-addr-detection-transmits number;
    (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
    ipv6-path-mtu-discovery-timeout minutes;
    no-tcp-reset (drop-all-tcp | drop-tcp-with-syn-only);
    no-tcp-rfc1323;
    no-tcp-rfc1323-paws;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port upper-limit upper-limit;
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
    tcp-mss bytes;
  }
  kernel-replication;
  license {
```

```
    autoupdate {
    url url;
    password password;
    }
    renew {
      before-expiration number;
      interval interval-hours;
    }
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
  }
  login {
    announcement text;
    class class-name {
      access-end hh:mm;
      access-start hh:mm;
      allow-commands regular-expression;
      allow-configuration regular-expression;
      allow-configuration-regexps [regular-expression];
      allowed-days [day];
      deny-commands regular-expression;
      deny-configuration regular-expression;
      deny-configuration-regexps [regular-expression];
      idle-timeout minutes;
      logical-system logical-system;
      login-alarms;
      login-script script;
      login-tip;
      permissions [permissions ];
      security-role (audit-administrator | crypto-administrator | ids-administrator |
        security-administrator);
    }
    deny-sources {
```


Copyright © 2016, Juniper Networks, Inc.

```
    address [address-or-hostname];
  }
  message text;
  }
  password {
    change-type (character-set | set-transitions);
    format (des | md5 | sha1);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
  }
  retry-options {
    backoff-factor seconds;
    backoff-threshold number;
    lockout-period time;
    maximum-time seconds;
    minimum-time seconds;
    tries-before-disconnect number;
  }
  user username {
    authentication {
      encrypted-password password;
      load-key-file url;
      plain-text-password;
      ssh-dsa public-key;
      ssh-rsa public-key;
    }
    class class-name;
    full-name complete-name;
    uid uid-value;
  }
}
log-vital {
  interval minutes;
  files days;
  storage-limit percentage;
  file-size Mbytes;
  add oid{
    comment comment;
  }
  group {
    operating;
    idp;
    storage;
    cluster-counter;
    screen zone-name;
    spu spu-name;
    }
  }
max-configuration-rollbacks number;
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server ip-address;
nd-maxmcast-solicit value;
nd-retrasmit-timer value;
no-compress-configuration-files;
```

Copyright © 2016, Juniper Networks, Inc.

```
no-debugger-on-alt-break;
no-multicast-echo;
no-neighbor-learn;
no-ping-record-route;
no-ping-time-stamp;
no-redirects;
no-saved-core-context;
ntp {
  authentication-key key-number {
    type md5;
    value password;
  }
  boot-server address;
  broadcast broadcast-address  {
    key key;
    ttl value;
    version version;
  }
  broadcast-client;
  multicast-client {
    address;
  }
  peer peer-address {
    key key;
    prefer;
    version version;
  }
  server server-address {
    key key;
    prefer;
    version version;
  }
  source-address source-address;
  trusted-key [key-number];
}
pic-console-authentication {
  encrypted-password password;
  plain-text-password;
}
ports {
  auxiliary {
    disable;
    insecure;
    type (ansi | small-xterm | vt100 | xterm);
  }
  console {
    disable;
    insecure;
    log-out-on-disconnect;
    type (ansi | small-xterm | vt100 | xterm);
  }
}
processes {
  802.1x-protocol-daemon {
    command binary-file-path;
    disable;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
adaptive-services {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
alarm-control {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
application-identification {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
application-security {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
audit-process {
   command binary-file-path;
   disable;
}
auto-configuration {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
bootp {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
chassis-control {
   disable;
   failover alternate-media;
}
class-of-service {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
craft-control {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
database-replication {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
datapath-trace-service {
   disable;
```

```
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  dhcp {
    command binary-file-path;
    disable;
  }
  dhcp-service {
    disable;
    failover (alternate-media | other-routing-engine);
    interface-traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  dialer-services {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
      no-remote-trace;
    }
  }
  diameter-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  disk-monitoring {
    command binary-file-path;
    disable;
  }
  dynamic-flow-capture {
    command binary-file-path;
    disable;
  }
  ecc-error-logging {
    command binary-file-path;
    disable;
  }
  ethernet-connectivity-fault-management {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  ethernet-link-fault-management {
    command binary-file-path;
    disable;
  }
  ethernet-switching {
    command binary-file-path;
    disable;
  }
  event-processing {
    command binary-file-path;
    disable;
  }
  fipsd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  firewall {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
```

```
}
firewall-authentication-service {
  disable;
}
forwarding {
  command binary-file-path;
  disable;
}
general-authentication-service {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
gprs-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
group-key-member {
  disable;
}
group-key-server {
  disable;
}
idp-policy {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
ilmi {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
inet-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
init {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
interface-control {
  command binary-file-path;
  disable;
```


Copyright © 2016, Juniper Networks, Inc.

```
    failover (alternate-media | other-routing-engine);
}
ipmi {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
ipsec-key-management {
    (disable | enable);
}
jsrp-service {
    disable;
}
jtasktest {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
kernel-replication {
    command binary-file-path;
    disable;
}
l2-learning {
    command binary-file-path;
    disable;
}
l2cpd-service {
    command binary-file-path;
    disable;
}
lacp {
    command binary-file-path;
    disable;
}
lldpd-service {
    command binary-file-path;
    disable;
}
logical-system-mux {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
logical-system-service {
    disable;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
mib-process {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mobile-ip {
  command binary-file-path;
  disable;
}
mountd-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
mspd {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
multicast-snooping {
  command binary-file-path;
  disable;
}
named-service {
  disable;
  failover (alternate-media | other-routing-engine);
}
neighbor-liveness {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
network-security {
  disable;
}
network-security-trace {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
nfsd-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
ntp {
  disable;
  failover (alternate-media | other-routing-engine);
}
ntpd-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
peer-selection-service {
```


Copyright © 2016, Juniper Networks, Inc.

```
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  periodic-packet-services {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  pgcp-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  pgm {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  pic-services-logging {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  ppp {
    command binary-file-path;
    disable;
  }
  pppoe {
    command binary-file-path;
    disable;
  }
  process-monitor {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error |info |notice | verbose | warning);
      no-remote-trace;
    }
  }
  profilerd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  r2cp {
    command binary-file-path;
    disable;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
redundancy-interface-process {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
remote-operations {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
resource-cleanup {
   disable;
   traceoptions {
      file {
         filename;
         files number;
         match regular-expression;
         size maximum-file-size;
         (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error |info |notice | verbose | warning);
      no-remote-trace;
   }
}
routing {
   disable;
   failover (alternate-media | other-routing-engine);
}
sampling {
   command binary-file-path;
   disable;
   failover (alternate-media | other-routing-engine);
}
sbc-configuration-process {
   disable;
   failover (alternate-media | other-routing-engine);
   traceoptions {
      file {
         filename;
         files number;
         match regular-expression;
         size maximum-file-size;
         (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
   }
}
sdk-service {
   disable;
   traceoptions {
      file {
         filename;
         files number;
         match regular-expression;
```


Copyright © 2016, Juniper Networks, Inc.

```
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error |info |notice | verbose | warning);
      no-remote-trace;
    }
  }
  secure-neighbor-discovery {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  security-log {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  send {
    disable;
  }
  service-deployment {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  shm-rtsdbd {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  simple-mail-client-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  smtpd-service {
    disable;
  }
  snmp {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  static-subscribers {
    disable;
  }
  statistics-service {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
  }
  subscriber-management {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
```

Copyright © 2016, Juniper Networks, Inc.

```
}
subscriber-management-helper {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
system-health-management {
  disable;
}
system-log-vital {
  disable;
}
tunnel-oamd {
  command binary-file-path;
  disable;
}
uac-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
usb-control {
  command binary-file-path;
  disable;
}
virtualization-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
vrrp {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
wan-acceleration {
  disable;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
watchdog {
  enable;
  disable;
  timeout value;
}
web-management {
  disable;
```


Copyright © 2016, Juniper Networks, Inc.

```
    failover (alternate media | other-routing-engine);
  }
  wireless-lan-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  wireless-wan-service {
    disable;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
proxy {
  password password;
  port port-number;
  server url;
  username user-name;
}
radius-options {
  attributes {
    nas-ip-address nas-ip-address;
    }
  password-protocol mschap-v2;
}
radius-server server-address {
  accounting-port number;
  max-outstanding-requests number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
root-authentication {
  encrypted-password password;
  load-key-file url;
  plain-text-password;
  ssh-dsa public-key {
```

```
      <from pattern-list>;
    }
    ssh-rsa public-key {
      <from pattern-list>;
    }
  }
  saved-core-context;
  saved-core-files number;
  scripts {
    commit {
      allow-transients;
      direct-access;
      file filename {
        checksum (md5 | sha-256 | sha1);
        optional;
        refresh;
        refresh-from url;
        source url;
      }
      refresh;
      refresh-from url;
      traceoptions {
        file {
          filename;
          files number;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
    load-scripts-from-flash;
    op {
      file filename {
        arguments name {
          description text;
        }
        checksum (md5 | sha-256 | sha1);
        command filename-alias;
        description cli-help-text;
        refresh;
        refresh-from url;
        source url;
      }
      no-allow-url;
      refresh;
      refresh-from url;
      traceoptions {
        file {
          filename;
          files number;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
        no-remote-trace;
    }
}
security-profile security-profile-name {
    address-book {
        maximum amount;
        reserved amount;
    }
    appfw-profile {
        maximum amount;
        reserved amount;
    }
    appfw-rule {
        maximum amount;
        reserved amount;
    }
    appfw-rule-set {
        maximum amount;
        reserved amount;
    }
    auth-entry {
        maximum amount;
        reserved amount;
    }
    cpu {
        reserved percent;
    }
    dslite-softwire-initiator {
        maximum amount;
        reserved amount;
    }
    flow-gate {
        maximum amount;
        reserved amount;
    }
    flow-session {
        maximum amount;
        reserved amount;
    }
    idp-policy idp-policy-name;
    logical-system logical-system-name;
    nat-cone-binding {
        maximum amount;
        reserved amount;
    }
    nat-destination-pool {
        maximum amount;
        reserved amount;
    }
    nat-destination-rule {
        maximum amount;
        reserved amount;
    }
    nat-interface-port-ol {
        maximum amount;
        reserved amount;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        nat-nopat-address {
            maximum amount;
            reserved amount;
        }
        nat-pat-address {
            maximum amount;
            reserved amount;
        }
        nat-pat-portnum {
            maximum amount
            reserved amount
        }
        nat-port-ol-ipnumber {
            maximum amount;
            reserved amount;
        }
        nat-rule-referenced-prefix {
            maximum amount;
            reserved amount;
        }
        nat-source-pool {
            maximum amount;
            reserved amount;
        }
        nat-source-rule {
            maximum amount;
            reserved amount;
        }
        nat-static-rule {
            maximum amount;
            reserved amount;
        }
        policy {
            maximum amount;
            reserved amount;
        }
        policy-with-count {
            maximum amount;
            reserved amount;
        }
        root-logical-system;
        scheduler {
            maximum amount;
            reserved amount;
        }
        zone {
            maximum amount;
            reserved amount;
        }
    }
    security-profile-resources {
        cpu-control;
        cpu-control-target percent;
    }
    services {
```


Copyright © 2016, Juniper Networks, Inc.

```
database-replication {
  traceoptions {
    file {
      filename ;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
dhcp {
  boot-file filename;
  boot-server (address | hostname);
  default-lease-time (infinite | seconds);
  domain-name domain-name;
  domain-search dns-search-suffix;
  maximum-lease-time (infinite | seconds);
  name-server ip-address;
  next-server ip-address;
  option option-identifier-code array type-name [ type-values ] | byte 8-bit-value | flag
    (false | off | on | true) | integer signed-32-bit-value | ip-address address | short
    signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
    unsigned-short 16-bit-value);
  pool subnet-ip-address/mask {
    address-range {
      high address;
      low address;
    }
    boot-file filename;
    boot-server (address | hostname);
    default-lease-time (infinite | seconds);
    domain-name domain-name;
    domain-search dns-search-suffix;
    exclude-address ip-address;
    maximum-lease-time (infinite | seconds);
    name-server ip-address;
    next-server ip-address;
    option option-identifier-code array type-name [ type-values ] | byte 8-bit-value |
      flag (false | off | on | true) | integer signed-32-bit-value | ip-address address |
      short signed-16-bit-value | string text-string | unsigned-integer 32-bit-value |
      unsigned-short 16-bit-value);
    propagate-ppp-settings interface-name;
    propagate-settings interface-name;
    router ip-address;
    server-identifier dhcp-server;
    sip-server {
      address ip-address;
      name sip-server-name;
    }
    wins-server ip-address;
  }
  propagate-ppp-settings interface-name;
  propagate-settings interface-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
    router ip-address;
    server-identifier dhcp-server;
    sip-server {
      address ip-address;
      name sip-server-name;
    }
    static-binding mac-address;
    traceoptions {
      file {
        filename ;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
    wins-server ip-address;
  }
  dhcp-local-server {
    dhcpv6 {
      authentication {
        password password;
        username-include {
          circuit-type;
          client-id;
          delimiter delimiter-character;
          domain-name domain-name;
          interface-name;
          logical-system-name;
          relay-agent-interface-id;
          relay-agent-remote-id;
          relay-agent-subscriber-id;
          routing-instance-name;
          user-prefix user-prefix;
        }
      }
      dynamic-profile {
        profile-name;
        aggregate-clients {
          merge;
          replace;
        }
        junos-default-profile;
        use-primary dynamic-profile-name;
      }
      group group-name {
        authentication {
          password password;
          username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
      interface-name;
      logical-system-name;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix;
    }
  }
  dynamic-profile {
    profile-name;
    aggregate-clients {
      merge;
      replace;
    }
    junos-default-profile;
    use-primary dynamic-profile;
  }
  interface interface-name {
    dynamic-profile {
      profile-name;
      aggregate-clients {
        merge;
        replace;
      }
      junos-default-profile;
      use-primary dynamic-profile-name;
    }
    exclude;
    overrides {
      delegated-pool pool-name;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      rapid-commit ;
    }
    service-profile service-profile-name
    trace ;
    upto interface-name;
  }
  liveness-detection {
    failure-action {
      clear-binding;
      clear-binding-if-interface-up;
      log-only;
    }
    method {
      bfd {
        detection-time {
          threshold milliseconds;
        }
        holddown-interval interval;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
```

```
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
          version (0 | 1 | automatic);
        }
      }
      overrides {
        delegated-pool pool-name;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
        rapid-commit ;
      }
      reconfigure {
        attempts number;
        clear-on-abort;
        strict;
        timeout number;
        token token-name;
        trigger {
          radius-disconnect;
        }
      }
      service-profile service-profile-name;
    }
    liveness-detection {
      failure-action {
        clear-binding;
        clear-binding-if-interface-up;
        log-only;
      }
      method {
        bfd {
          detection-time {
            threshold milliseconds;
          }
          holddown-interval interval;
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          session-mode (automatic | multihop | single-hop);
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
        version (0 | 1 | automatic);
      }
    }
    overrides {
      delegated-pool pool-name;
      interface-client-limit number;
```


Copyright © 2016, Juniper Networks, Inc.

```
            process-inform {
                pool pool-name;
            }
            rapid-commit ;
        }
        reconfigure {
            attempts number;
            clear-on-abort;
            strict;
            timeout number;
            token token-name;
            trigger {
                radius-disconnect;
            }
        }
        service-profile service-profile-name;
    }
    group group-name {
        interface interface-name {
            exclude;
            upto upto-interface-name;
        }
    }
}
dns {
    dns-proxy {
        cache hostname inet ip-address;
        default-domain domain-name {
            forwarders ip-address;
        }
        interface interface-name;
        propogate-setting (enable | disable);
        view view-name {
            domain domain-name {
                forward-only;
                forwarders ip-address;
            }
            match-clients subnet-address;
        }
    }
}
dnssec {
    disable;
    dlv {
        domain-name domain-name trusted-anchor  trusted-anchor;
        }
secure-domains domain-name;
    trusted-keys (key dns-key | load-key-file url);
forwarders {
    ip-address;
}
max-cache-ttl seconds;
max-ncache-ttl seconds;
traceoptions {
    category {
        category-type;
```

```
        }
        debug-level level;
        file {
          filename;
          files number;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
    }
    dynamic-dns {
      client hostname {
        agent agent-name;
        interface interface-name;
        password server-password;
        server server-name;
        username user-name;
      }
    }
    finger {
      connection-limit number;
      rate-limit number;
    }
    ftp {
      connection-limit number;
      rate-limit number;
    }
    netconf {
      ssh {
        connection-limit number;
        port port-number;
        rate-limit number;
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
        on-demand;
      }
    }
    outbound-ssh {
      client client-id {
        address {
          port port-number;
          retry number;
          timeout value;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        device-id device-id;
        keep-alive {
            retry number;
            time-out value;
        }
        reconnect-strategy (in-order |sticky);
        secret secret;
        services {
            netconf;
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
service-deployment {
    local-certificate certificate-name;
    servers server-address {
        port port-number;
        security-options {
            ssl3;
            tls;
        }
        user user-name;
    }
    source-address source-address;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
ssh {
    ciphers [cipher];
    client-alive-count-max number;
    client-alive-interval seconds;
    connection-limit number;
    hostkey-algorithm {
        (ssh-dss | no-ssh-dss);
        (ssh-ecdsa |no-ssh-ecdsa);
        (ssh-rsa | no-ssh-rsa);
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    key-exchange [algorithm];
    macs [algorithm];
    max-sessions-per-connection number;
    protocol-version {
      v1;
      v2;
    }
    rate-limit number;
    root-login (allow | deny | deny-password);
    (tcp-forwarding | no-tcp-forwarding);
  }
  subscriber-management {
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    maintain-subscriber interface-delete;
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  subscriber-management-helper {
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
  telnet {
    connection-limit number;
    rate-limit number;
  }
  web-management {
    control {
      max-threads number;
    }
    http {
      interface [interface-name];
      port port-number;
    }
    https {
      interface [interface-name];
      local-certificate name;
      pki-local-certificate name;
```


Copyright © 2016, Juniper Networks, Inc.

```
        port port-number;
        system-generated-certificate;
      }
      management-url url;
      session {
        idle-timeout minutes;
        session-limit number;
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
    }
    xnm-clear-text {
      connection-limit number;
      rate-limit number;
    }
    xnm-ssl {
      connection-limit number;
      local-certificate name;
      rate-limit number;
    }
  }
  static-host-mapping hostname {
    alias [host-name-alias];
    inet [ip- address];
    inet6 [ipv6- address];
    sysid system-identifier;
  }
  syslog {
    allow-duplicates;
    archive {
      binary-data;
      files number;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    console {
      (any | facility) severity;
    }
    file filename {
      allow-duplicates;
      archive {
        archive-sites url {
          password password;
        }
        (binary-data| no-binary-data);
        files number;
```

```
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
          }
          structure-data {
            brief;
          }
          (any | facility) severity;
        }
        host (hostname | other-routing-engine) {
          (any | facility) severity;
        }
        log-rotate-frequency minutes;
        source-address source-address;
        time-format {
          millisecond;
          year;
        }
        user (username | *) {
          (any | facility) severity;
        }
      }
      tacplus-options {
        (exclude-cmd-attribute | no-cmd-attribute-value);
        service-name service-name;
      }
      tacplus-server server-address {
        port port-number;
        secret password;
        single-connection;
        source-address source-address;
        timeout seconds;
      }
      time-zone (GMThour-offset | time-zone);
      tracing {
        destination-override {
          syslog {
            host address;
            }
        }
      }
      use-imported-time-zones;
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security firewall-authentication] Hierarchy Level

```
security {
  firewall-authentication {
    traceoptions {
      flag flag;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security policies] Hierarchy Level

```
security {
    policies {
        default-policy (deny-all | permit-all);
        from-zone zone-name to-zone zone-name {
            policy policy-name {
                description description;
                match {
                    application {
                        [application];
                        any;
                    }
                    destination-address {
                        [address];
                        any;
                        any-ipv4;
                        any-ipv6;
                        }
                    destination-address-excluded;
                    source-address {
                        [address];
                        any;
                        any-ipv4;
                        any-ipv6;
                        }
                    source-address-excluded;
                    source-identity {
                        [role-name];
                        any;
                        authenticated-user;
                        unauthenticated-user;
                        unknown-user;
                    }
                }
                scheduler-name scheduler-name;
                then {
                    count {
                        alarm {
                            per-minute-threshold number;
                            per-second-threshold number;
                        }
                    }
                    deny;
                    log {
                        session-close;
                        session-init;
                    }
                    permit {
```

Copyright © 2016, Juniper Networks, Inc.

```
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
          redirect-wx | reverse-redirect-wx;
          ssl-proxy {
            profile-name profile-name;
          }
          uac-policy {
            captive-portal captive-portal;
          }
          utm-policy policy-name;
        }
        destination-address {
          drop-translated;
          drop-untranslated;
        }
        firewall-authentication {
          pass-through {
            access-profile profile-name;
            client-match user-or-group-name;
            ssl-termination-profile profile-name;
            web-redirect;
            web-redirect-to-https;
          }
          user-firewall {
            access-profile profile-name;
            domain domain-name
            ssl-termination-profile profile-name;
          }
          web-authentication {
            client-match user-or-group-name;
          }
        }
        services-offload;
        tcp-options {
          sequence-check-required;
          syn-check-required;
        }
        tunnel {
          ipsec-group-vpn group-vpn;
          ipsec-vpn vpn-name;
          pair-policy pair-policy;
        }
      }
      reject;
    }
  }
}
global {
```


Copyright © 2016, Juniper Networks, Inc.

```
policy policy-name {
  description description;
  match {
    application {
      [application];
      any;
    }
    destination-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    from-zone {
      [zone-name];
      any;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
    to-zone {
      [zone-name];
      any;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
```

```
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            syn-check-required;
          }
        }
        reject;
      }
    }
  }
  policy-rematch;
  policy-stats {
    system-wide (disable | enable);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
      size maximum-file-size;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        flag flag;
        no-remote-trace;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

[edit security user-identification] Hierarchy Level

```
security {
  user-identification {
    authentication-source {
      active-directory-authentication-table priority priority;
      firewall-authentication priority priority;
      local-authentication-table priority priority;
      unified-access-control priority priority;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- authentication-source on page 1234
- Security Configuration Statement Hierarchy on page 603

[edit security zones] Hierarchy Level

```
security {
  zones {
    functional-zone {
      management {
        description text;
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
```

```
        }
      }
      interfaces interface-name  {
        host-inbound-traffic {
          protocols protocol-name {
            except;
          }
          system-services service-name {
            except;
          }
        }
      }
      screen screen-name;
    }
  }
  security-zone zone-name {
    address-book {
      address address-name {
        ip-prefix {
          description text;
        }
        description text;
        dns-name domain-name {
          ipv4-only;
          ipv6-only;
        }
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
      }
      address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
      }
    }
    application-tracking;
    description text;
    host-inbound-traffic {
      protocols protocol-name {
        except;
      }
      system-services service-name {
        except;
      }
    }
    interfaces interface-name {
      host-inbound-traffic {
        protocols protocol-name {
          except;
        }
        system-services service-name {
          except;
        }
      }
    }
    screen screen-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
            tcp-rst;
        }
    }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Logical Systems for SRX Series Services Gateways on page 3205
- Security Zones and Interfaces Overview on page 1049

Copyright © 2016, Juniper Networks, Inc.

actions (Services SSL Proxy)

Syntax

```
actions {
  crl {
    disable;
    if-not-present (allow | drop);
    ignore-hold-instruction-code;
  }
  disable-session-resumption;
  ignore-server-auth-failure;
  logs {
    all;
    errors;
    info;
    sessions-allowed;
    sessions-dropped;
    sessions-ignored;
    sessions-whitelisted;
    warning;
  }
  renegotiation {
    (allow | allow-secure | drop);
  }
}
```

Hierarchy Level [edit services ssl proxy profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the logging and traffic related actions.

Options

- **crl**—Specify the certificate revocation actions.
  - **disable**—Disable CRL verification.
  - **if-not-present**—Specify actions for sessions.
    - **allow**—Allow sessions when CRL information is not available.
    - **drop**—Drop sessions when CRL information is not available.
  - **ignore-hold-instruction-code**—Ignore the unconfirmed (on hold) revocation status, and accept a certificate.
- **disable-session-resumption**—Disable session resumption.
- **ignore-server-auth-failure**—Ignore server authentication failure.
- **log**—Specify the logging actions.
  - **all**—Log all events.
  - **errors**—Log all error events.
  - **info**—Log all information events.
  - **sessions-allowed**—Log SSL session allowed events after an error.

Copyright © 2016, Juniper Networks, Inc.

- **sessions-dropped**—Log only SSL session dropped events.
- **sessions-ignored**—Log session ignored events.
- **sessions-whitelisted**—Log SSL session whitelisted events.
- **warning**—Log all warning events.

- **renegotiation**—Specify the renegotiation options.
  - **allow**—Allow secure and nonsecure renegotiation.
  - **allow-secure**—Allow secure negotiation only.
  - **drop**—Drop session on renegotiation request.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

active-directory-access

Syntax

```
active-directory-access {
  domain domain-name {
    user username;
      password password;
    domain-controller domain-controller-name {
      address domain-controller-address;
    }
    ip-user-mapping {
      discovery-method {
        wmi {
          event-log-scanning-interval seconds;
          initial-event-log-timespan hours;
        }
      }
    }
    user-group-mapping {
      ldap {
        authentication-algorithm {
          simple;
        }
        ssl;
        base base;
        user name {
          password password;
        }
        address ip-address {
          port port;
        }
      }
    }
  }
}
```

Hierarchy Level [edit services user-identification]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Identify the domain and domain controllers where the integrated user firewall feature is implemented; configure the IP address-to-user mapping information and the user-to-group mapping information for accessing the LDAP server.

Options

domain *domain-name*—Required. Name of the domain; the length of the name ranges from 1 through 64 characters. The SRX Series device can have the integrated user firewall feature configured in a maximum of two domains.

user *username*—Required. Active Directory account name.

Range: 1 through 64 characters.

password *password*—Required. Password of the Active Directory account.

Range: 1 through 128 characters.

Copyright © 2016, Juniper Networks, Inc.

domain-controller ***domain-controller-name***—Required. Name of the domain controller; the length of the name can range from 1 through 64 characters. A maximum of 10 domain controllers can be configured.

address ***domain-controller-address***—Required. IP address of the domain controller.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- user-identification (Services) on page 5499
- LDAP Functionality in Integrated User Firewall on page 5317

Copyright © 2016, Juniper Networks, Inc.

active-directory-authentication-table

Syntax

```
active-directory-authentication-table {
    priority priority;
}
```

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description An authentication table is generated by polling Active Directory domain controllers for source identity information about active users. Each entry in the table correlates an authenticated user with an IP address and associated user groups. That information is used for matching in IP-based firewall policies. The user information must be retrieved from the table before policy lookup can proceed and traffic is allowed to pass through the firewall.

Options

priority *priority*—Specify the priority of the Active Directory authentication table. The priority determines the sequence for searching among various other authentication tables to retrieve a user role. The priorities of the following tables are considered: local authentication table, firewall authentication table, Active Directory authentication table, and UAC authentication table.

Each authentication table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. Setting the priority value of a table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Range: A unique value from 0 through 65535.

Default: The default priority of the Active Directory authentication table is 125.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- authentication-source on page 1234
- Overview of Integrated User Firewall on page 5309
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

address (Services)

Syntax address *ip-address*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the IP address of the IC Series device with which the SRX Series devices should communicate.

This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Unified Access Control on page 5269
- Acquiring User Role Information from an Active Directory Authentication Server on page 5269

admin-search

Syntax

```
admin-search {
    distinguished-name distinguished-name;
    password password;
}
```

Hierarchy Level [edit access ldap-options search],
[edit access profile*profile-name* ldap-options search]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that a Lightweight Directory Access Protocol (LDAP) administrator search is performed. By default, the search is an anonymous search. To perform an administrator search, you must specify administrator credentials, which are used in the bind as part of performing the search.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

application (Security Policies)

Syntax

```
application {
  [application];
  any;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP or remote procedure call (RPC) application or set of applications to be used as match criteria.

Options

application-name-or-set—Name of the predefined or custom application or application set used as match criteria.

any—Any predefined or custom applications or application sets.

NOTE: A custom application that does not use a well-known destination port for the application will not be included in the any option, and must be named explicitly.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application-services (Security Policies)

Syntax

```
application-services {
  application-firewall {
    rule-set rule-set-name;
  }
  application-traffic-control {
    rule-set rule-set-name;
  }
  gprs-gtp-profile profile-name;
  gprs-sctp-profile profile-name;
  idp;
  redirect-wx | reverse-redirect-wx;
  ssl-proxy {
    profile-name profile-name;
  }
  uac-policy {
    captive-portal captive-portal;
  }
  utm-policy policy-name;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement modified in Junos OS Release 11.1.

Description Enable application services within a security policy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Application Firewall Overview on page 555

Copyright © 2016, Juniper Networks, Inc.

assemble

Syntax

```
assemble {
  common-name common-name;
}
```

Hierarchy Level [edit access ldap-options],
[edit access profile *profile-name* ldap-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that a user's LDAP distinguished name (DN) is assembled through the use of a common name identifier, the username, and base distinguished name.

Options **common-name** *common-name*—Common name identifier used as a prefix for the username during the assembly of the user's distinguished name. For example, **uid** specifies " user id," and **cn** specifies "common name."

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

authentication-source

Syntax

```
authentication-source {
  active-directory-authentication-table priority priority;
  firewall-authentication priority priority;
  local-authentication-table priority priority;
  unified-access-control priority priority;
}
```

Hierarchy Level [edit security user-identification]

Release Information Statement introduced in Junos OS Release 12.1. Statement updated in Junos OS Release 12.1-X45-D10. Support for the **active-directory-authentication-table priority** command statement added in Junos OS Release 12.1X47-D10.

Description Identifies one or more tables to be used as the source for user role information. Tables are searched in sequence based on lowest to highest priority.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- local-authentication-table on page 1265
- firewall-authentication (User Identification) on page 1252
- unified-access-control (Security) on page 1329
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

banner (Access FTP HTTP Telnet Authentication)

Syntax

```
banner {
  fail string;
  login string;
  success string;
}
```

Hierarchy Level [edit access firewall-authentication pass-through (ftp | http | telnet)]

Release Information Statement introduced in Junos OS Release 8.5.
HTTPS for web-authentication is supported on high-end SRX Series Services Gateways from Junos OS Release 12.1X44-D10 and on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways from Junos OS Release 15.1X49-D40.

Description Configure the banners that appear to users during the FTP, HTTP, HTTPS, and Telnet firewall authentication process. The banners appear during login, after successful authentication, and after failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Understanding Firewall Authentication Banner Customization on page 5251

banner (Access Web Authentication)

Syntax

```
banner {
  success string;
}
```

Hierarchy Level [edit access firewall-authentication web-authentication]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the banner that appears to users during the Web authentication process. The banner appears during login, after successful authentication, and after failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

base-distinguished-name

Syntax base-distinguished-name *base-distinguished-name*;

Hierarchy Level [edit access ldap-options],
[edit access profile*profile-name* ldap-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the base distinguished name (DN), which can be used in one of the following ways:

- If you are using the **assemble** statement so that the user's distinguished name is being assembled, the base distinguished name is appended to a username to generate the user's distinguished name. The resulting distinguished name is used in the LDAP bind call.
- If you are using the **search** statement so that the user's distinguished name is found by a search, the search is restricted to the subtree of the base distinguished name.

Options *base-distinguished-name*—Series of basic properties that define the user. For example in the base distinguished name **o=juniper, c=us**, where **c** stands for country, and **o** for organization.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

ca-profile (Services)

Syntax ca-profile *ca-profile*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the certificate authority (CA) of the certificate that the SRX Series device should use in communications with an Infranet Enforcer. The SRX Series device uses the CA to validate the IC Series UAC Appliance server certificate.

Use this statement if you have loaded certificates from multiple certificate authorities (CAs) onto your SRX Series device and you need to configure the device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series appliance .

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

captive-portal (Services UAC)

Syntax

```
captive-portal redirect-policy-name{
  redirect-traffic (all | unauthenticated);
  redirect-url redirect-url;
}
```

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the preconfigured security policy for captive portal on the Junos OS Enforcer to enable the captive portal feature. The captive portal policy is configured as part of the UAC policy.

By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- *Unified Access Control Design and Implementation Guide for Security Devices*

captive-portal (Services UAC Policy)

Syntax

```
captive-portal captive-portal-policy-name;
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services uac-policy]

Release Information Statement introduced in Junos OS Release 10.2.

Description Create the captive portal policy in the UAC security policy. You use the captive portal policy to configure the captive portal feature on the Junos OS Enforcer. By configuring the captive portal feature, you can redirect traffic destined for protected resources to the IC Series device or to the URL you configure on the Junos OS Enforcer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

certificate-verification

Syntax certificate-verification [optional | required | warning]

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 12.1.

Description This option determines whether server certificate verification is required when initiating a connection between an SRX Series device and a Junos Pulse Access Control Service in a UAC configuration. If no CA profile contains the certificate authority (CA) that signed the configured server certificate for the Access Control Service, this option determines whether the commit check should fail, a warning should be displayed, or the connection should be made without any warning.

By default, an administrator is warned if the CA certificate is not configured in the **ca-profile**.

NOTE: For strict security, this option should be reset to required, and the proper CA certificate should be specified in the CA profile.

Options

- **optional**—Certificate verification is not required. If the CA certificate is not specified in the **ca-profile** option, the commit check passes and no warning is issued.
- **required**—Certificate verification is required. If the CA certificate is not specified in the **ca-profile** option, an error message is displayed, and the commit check fails. Use this option to ensure strict security.
- **warning**—(Default) Certificate verification is not required, however, a warning message is displayed during commit check if the CA certificate is not specified in the **ca-profile** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Communications Between the Junos OS Enforcer and the IC Series UAC Appliance on page 5261
- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

client-group

Syntax client-group [*group-names*];

Hierarchy Level [edit access profile *profile-name* client *client-name*]
[edit access profile *profile-name* session-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a list of client groups that the client belongs to. If the group list is not defined as part of the client profile, the client group configured in the **profile session-options** is used.

Options *group-names*—Names of one or more groups to which the client belongs, separated by spaces—for example **g1 g2 g3**. The total length of the group name string cannot exceed 256 characters.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

client-idle-timeout (Access Profile)

Syntax client-idle-timeout *minutes*;

Hierarchy Level [edit access profile *profile-name* session-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

Options *minutes*—Number of minutes of idle time that elapse before the session is terminated.

Range: 10 through 255 minutes

Default: 10 minutes

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

client-name-filter

Syntax

```
client-name-filter client-name {
    count number;
    domain-name domain-name;
    separator special-character;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define restrictions for named clients. Client whose name matches these restrictions is authenticated on the server.

Options *client-name*—Name of the client.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

client-session-timeout (Access Profile)

Syntax client-session-timeout *minutes*;

Hierarchy Level [edit access profile *profile-name* session-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the amount of time after which a user's sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Options *minutes*—Number of minutes after which a user's sessions are terminated.

Range: 1 through 10,000 minutes

Default: Off

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

configuration-file

Syntax configuration-file *filepath*;

Hierarchy Level [edit access securid-server *server-name*]

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify the name of the SecurID server and the path to the configuration file. The file is copied on the devices in some directory location—for example, **/var/db/securid/sdconf.rec**.

Options
- *server-name*—Name of the SecurID authentication server.
- *filepath*—Path of the SecurID server configuration file.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

count

Syntax count *number*;

Hierarchy Level [edit access profile *profile-name* client-name-filter]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the number of characters to be stripped from a client name, from right to left, until the specified number of characters are deleted. The resulting name is sent to the authentication server.

Options *number*—Number of characters to be stripped from a client name.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

custom-ciphers

Syntax custom-ciphers [rsa-with-rc4-128-md5 RSA | 128bit rc4 | md5 hash rsa-with-rc4-128-sha RSA | 128bit rc4 |sha hash rsa-with-des-cbc-sha RSA | des cbc | sha hash rsa-with-3des-ede-cbc-sha RSA | 3des ede/cbc | sha hash rsa-with-aes-128-cbc-sha RSA | 128 bit aes/cbc | sha hash rsa-with-aes-256-cbc-sha RSA | 256 bit aes/cbc |sha hash rsa-export-with-rc4-40-md5 RSA-export | 40 bit rc4 | md5 hash rsa-export-with-des40-cbc-sha RSA-export | 40 bit des/cbc | sha hash rsa-with-null-md5 RSA | no symmetric cipher | md5 hash rsa-with-null-sha RSA | no symmetric cipher | sha hash];

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Display the custom cipher list.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

default-profile

Syntax default-profile *profile-name*;

Hierarchy Level [edit access firewall-authentication web-authentication]
[edit access firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the authentication profile to use if no profile is specified in a policy.

Options *profile-name*—Name of the default authentication profile.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

distinguished-name (Access)

Syntax distinguished-name *distinguished-name*;

Hierarchy Level [edit access ldap-options search admin-search]
[edit access profile *profile-name* ldap-options search admin-search]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the distinguished name of an administrative user. The distinguished name is used in the bind for performing the LDAP search.

Options *distinguished-name*—Set of properties that define the user. For example, **cn**=admin, **ou**=eng, **o**=juniper, **dc**=net.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

domain-name (Access Profile)

Syntax domain-name *domain-name*;

Hierarchy Level [edit access profile *profile-name* client-name-filter]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a domain name that must be in a client's name during the authentication process.

Options *domain-name*—Domain name that must be in a client name. The name must not exceed 128 characters.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

enable-flow-tracing (Services)

Syntax enable-flow-tracing;

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Enable flow tracing for the profile.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

enable-session-cache

Syntax enable-session-cache;

Hierarchy Level [edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Enable SSL session cache.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- SSL Proxy Overview on page 531
- Configuring SSL Proxy on page 540
- Enabling Debugging and Tracing for SSL Proxy on page 552

Copyright © 2016, Juniper Networks, Inc.

fail

Syntax fail *string*;

Hierarchy Level [edit access firewall-authentication pass-through *profile-name* (ftp | http | telnet) banner]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the banner that a client sees if the authentication process fails.

Options *string*—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text in spaces or special characters, such as quotation marks (" ").

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

file (Services)

Syntax

```
file file-name; {
  files;
  match;
  no-world-readable size;
  world-readable;
}
```

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the trace file information.

Options

- **files**—Specify the maximum number of trace files. Range: 2 through 1000.
- **match**—Specify the regular expression for lines to be logged.
- **no-world-readable size**—Do not allow any user to read the log file.
- **size**—Specify the maximum trace file size. Range: 10,240 to 1,073,741,824.
- **world-readable**—Allow any user to read the log file.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

files (Services)

Syntax files *files*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the maximum number of trace files.

Options **files**—Specify the maximum number of trace files.

Range: 2 to 1000

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

file (System Logging)

Syntax

```
file filename {
  allow-duplicates;
  any (alert | any | critical | emergency | error | info | none | notice | warning);
  archive {
    archive-sites {
      url password;
    }
    (binary-data | no-binary-data);
    files number;
    size size;
    start-time start-time;
    transfer-interval transfer-interval;
    (world-readable | no-world-readable);
  }
  authorization (alert | any | critical | emergency | error | info | none | notice | warning);
  change-log (alert | any | critical | emergency | error | info | none | notice | warning);
  conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
  daemon (alert | any | critical | emergency | error | info | none | notice | warning);
  dfc (alert | any | critical | emergency | error | info | none | notice | warning);
  explicit-priority;
  external (alert | any | critical | emergency | error | info | none | notice | warning);
  firewall (alert | any | critical | emergency | error | info | none | notice | warning);
  ftp (alert | any | critical | emergency | error | info | none | notice | warning);
  interactive-commands (alert | any | critical | emergency | error | info | none | notice | warning);
  kernel (alert | any | critical | emergency | error | info | none | notice | warning);
  match "regular-expression";
  ntp (alert | any | critical | emergency | error | info | none | notice | warning);
  pfe (alert | any | critical | emergency | error | info | none | notice | warning);
  security (alert | any | critical | emergency | error | info | none | notice | warning);
  structured-data {
    brief;
  }
  user (alert | any | critical | emergency | error | info | none | notice | warning);
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 12.1X47 for SRX Series.

Description Specify the file in which to log data.

Options

- *filename*—Specify the name of the file in which to log data.
- *allow-duplicates*—Do not suppress the repeated messages.
- *any*—Specify all facilities information.
  - *alert*—Specify the conditions that should be corrected immediately.
  - *critical*—Specify the critical conditions.
  - *emergency*—Specify the conditions that cause security functions to stop.
  - *error*—Specify the general error conditions.

Copyright © 2016, Juniper Networks, Inc.

    - *info*—Specify the information about normal security operations.
    - *none*—Do not specify any messages.
    - *notice*—Specify the conditions that should be handled specifically.
    - *warning*—Specify the general warning conditions.
- *archive*—Specify the archive file information.
    - *archive-sites*—Specify a list of destination URLs for the archived log files.
        - *url*—Specify the primary and failover URLs to receive archive files.
    - *binary-data*—Mark file such that it contains binary data.
    - *no-binary-data*—Do not mark the file such that it contains binary data.
    - *files*—Specify the number of files to be archived. Range: 1 through 1000 files.
    - *size*—Specify the size of files to be archived. Range: 65,536 through 1,073,741,824 bytes.
    - *world-readable*—Allow any user to read the log file.
    - *no-world-readable*—Do not allow any user to read the log file.
    - *start-time*—Specify the start time for file transmission. Enter the start time in the yyyy-mm-dd.hh:mm format.
    - *transfer-interval*—Specify the frequency at which to transfer the files to archive sites.
- *authorization*—Specify the authorization system.
- *change-log*—Specify the configuration change log.
- *conflict-log*—Specify the configuration conflict log.
- *daemon*—Specify the various system processes.
- *dfc*—Specify the dynamic flow capture.
- *explicit-priority*—Include the priority and facility in messages.
- *external*—Specify the local external applications.
- *firewall*—Specify the firewall filtering system.
- *ftp*—Specify the FTP process.
- *interactive-commands*—Specify the commands executed by the UI.
- *kernel*—Specify the kernel information.
- *match*—Specify the regular expression for lines to be logged.
- *ntp*—Specify the NTP process.
- *pfe*—Specify the Packet Forwarding Engine.
- *security*—Specify the security-related information.

Copyright © 2016, Juniper Networks, Inc.

- *structured-data*—Log the messages in structured log format.
  - *brief*—Omit English language text from the end of the logged message.
- *user*—Specify the user processes.
  - *info*—Specify the informational messages.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Junos OS System Log Overview*
- *syslog (System)*

Copyright © 2016, Juniper Networks, Inc.

firewall-authentication

Syntax

```
firewall-authentication {
  pass-through {
    default-profile profile-name;
    ftp {
      banner {
        fail string;
        login string;
        success string;
      }
    }
    http {
      banner {
        fail string;
        login string;
        success string;
      }
    telnet {
      banner {
        fail string;
        login string;
        success string;
      }
  }
  traceoptions {
    file {
      filename;
      files number;
      flag flag;
      match regular-expression;
      no-remote-trace;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
  }
  web-authentication {
    banner {
      success string;
    }
    default-profile profile-name;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure default firewall authentication settings used by firewall authentication policies that restrict and permit access of firewall users to protected resources behind a firewall.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- *Dynamic VPN Overview*
- Firewall User Authentication Overview on page 5179

firewall-authentication (Security)

Syntax

```
firewall-authentication {
  traceoptions {
    flag flag;
  }
}
```

Hierarchy Level
[edit security]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Define data-plane firewall authentication tracing options.

Options
- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple flag statements.
  - **all**—Enable all tracing operations.
  - **authentication**—Trace data-plane firewall authentication events.
  - **proxy**—Trace data-plane firewall authentication proxy events.
- **detail**—Display moderate amount of data.
- **extensive**—Display extensive amount of data.
- **terse**—Display minimum amount of data.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Understanding Logical System Firewall Authentication on page 3267

Copyright © 2016, Juniper Networks, Inc.

firewall-authentication (Security Policies)

Syntax

```
firewall-authentication {
  pass-through {
    access-profile profile-name;
    client-match user-or-group-name;
    ssl-termination-profile profile-name;
    web-redirect;
    web-redirect-to-https;
  }
  user-firewall {
    access-profile profile-name;
    domain domain-name
    ssl-termination-profile profile-name;
  }
  web-authentication {
    client-match user-or-group-name;
  }
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit]

Release Information Statement introduced in Junos OS Release 8.5.
Support added for the **user-firewall** option in Junos OS Release 12.1X45-D10.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Configure firewall authentication methods.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

firewall-authentication (User Identification)

Syntax firewall-authentication priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Junos OS Release 12.1X45-D10. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description Enables the firewall authentication table as an authentication source. The priority of this table among other authentication tables establishes the search sequence used to identify user and role values.

Options priority *priority*—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the firewall authentication table is 150.

Setting the priority value of the firewall authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127

firewall-authentication-service

Syntax firewall-authentication-service (enable | disable);

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable or disable the firewall authentication service process.

Options
- **enable**—Start the firewall authentication service process.
- **disable**—Stop the firewall authentication service process.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

firewall-user

Syntax

```
firewall-user {
    password password;
}
```

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a client as a firewall user and the associated password (encrypted).

Options **password *password*** —Password used by the firewall user during local authentication.

Range: 1 through 128 characters

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

flag (Services)

Syntax flag (*all* | *cli-configuration* | *initiation* | *proxy* | *selected-profile* | *termination*);

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the tracing flag parameters.

Options

- *all*—Trace all the parameters.
- *cli-configuration*—Trace CLI configuration events.
- *initiation*—Trace initiation service events.
- *proxy*—Trace proxy service events.
- *selected-profile*—Trace events for profiles with enable-flow-tracing set.
- *termination*—Trace termination service events.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

from-zone (Security Policies)

Syntax

```
from-zone zone-name to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```

Copyright © 2016, Juniper Networks, Inc.

```
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

Hierarchy Level [edit security policies]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **description** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Specify a source zone and destination zone to be associated with the security policy.

Options

- **from-zone** *zone-name*—Name of the source zone.
- **to-zone** *zone-name*—Name of the destination zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

ftp (Access)

Syntax

```
ftp {
  banner {
    fail string;
    login string;
    success string;
  }
}
```

Hierarchy Level [edit access firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure banners for the FTP login prompt, successful authentication, and failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Understanding Pass-Through Authentication on page 5205
- Example: Configuring Pass-Through Authentication on page 5207

Copyright © 2016, Juniper Networks, Inc.

group-profile (Access)

Syntax

```
group-profile profile-name {
  ppp {
    cell-overhead;
    encapsulated-overhead encapsulated-overhead-value;
    framed-pool address-pool-name;
    idle-timeout seconds;
    interface-id interface-identifier;
    keepalive seconds;
    ppp-options {
      chap;
      pap;
    }
    primary-dns name;
    primary-wins name;
    secondary-dns name;
    secondary-wins name;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure a group profile to define Point-to-Point Protocol (PPP) attributes. Any client referencing the configured group profile inherits all the group profile attributes.

Options

- **ppp**—Configure Point-to-Point Protocol (PPP) attributes.
  - **cell-overhead**—Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping.
  - **encapsulated-overhead *encapsulated-overhead-value***—Configure the encapsulation overhead for class-of-service calculations.
  - **framed-pool *pool-name***—Configure a framed-pool.
  - **idle-timeout**—Configure the idle timeout for a user.
  - **interface-id**—Configure the interface identifier.
  - **keep-alive**—Configure the keepalive interval for an L2TP tunnel.
- **ppp-options**—Configure PPP authentication.
  - **pap**—Specify Password Authentication Protocol.
  - **chap**—Specify Challenge Handshake Authentication Protocol.
- **primary-dns**—Specify the primary-dns IP address.
- **secondary-dns**—Specify the secondary-dns IP address.
- **primary-wins**—Specify the primary-wins IP address.
- **secondary-wins**—Specify the secondary-wins IP address.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

http (Access)

Syntax

```
http {
  banner {
    fail string;
    login string;
    success string;
  }
}
```

Hierarchy Level
[edit access firewall-authentication pass-through]

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Configure banners for the HTTP login prompt, successful authentication, and failed authentication.

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

Copyright © 2016, Juniper Networks, Inc.

http (Services)

Syntax

```
http {
  interfaces [ interface-names ];
  port port;
}
```

Hierarchy Level [edit system services web-management]

Release Information Statement introduced in Junos OS Release 8.5.

Description Unencrypted HTTP connection setting.

Options **interfaces [*interface-names*]**—Name of one or more interfaces on which to allow the HTTP service.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring HTTPS Traffic to Trigger Pass-Through Authentication on page 5213
- Example: Configuring HTTPS Traffic to Trigger Web Authentication on page 5230
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

https (Services)

Syntax

```
https {
  interfaces [ interface-names ];
  local-certificate name;
  pki-local-certificate name;
  port port;
  system-generated-local-certificate name;
}
```

Hierarchy Level [edit system services web-management]

Release Information Statement introduced on high-end SRX Series Services Gateways, from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Encrypted HTTPS connections.

Options

interfaces [*interface-names*]—Name of one or more interfaces on which to allow the HTTPS service.

local-certificate *name*—Name of the X.509 certificate for a Secure Sockets Layer (SSL) connection. An SSL connection is configured at the **[edit security certificates local]** hierarchy.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring HTTPS Traffic to Trigger Pass-Through Authentication on page 5213
- Example: Configuring HTTPS Traffic to Trigger Web Authentication on page 5230
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

infranet-controller

Syntax

```
infranet-contoller host-name {
  address ip-address;
  ca-profile [ca-profile];
  interface interface-name;
  password password;
  port port-number;
  server-certificate-subject subject;
}
```

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description To configure an Infranet Controller, specify the hostname of the IC Series device with which the SRX Series device should communicate. Possible values for this statement range from 1 to 31 characters.

This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

One or more IC Series devices can be configured as Infranet Controllers on the SRX Series device. There is no maximum number of IC Series devices that can be configured. However, only one IC Series device can be active at any time. The others are failover devices. A round-robin algorithm determines which of the configured IC Series devices is the active Infranet Controller. If the active Infranet Controller becomes inoperative, the algorithm is reapplied to the remaining IC Series devices that are configured to establish the new active Infranet Controller.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

interface (Services)

Syntax interface *interface-name*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the SRX Series interface through which the IC Series device should connect.

This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- port (Services) on page 5456
- password (Services) on page 5446

interval (Services)

Syntax interval *seconds*;

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the value in seconds that the SRX Series device should expect to receive a heartbeat signal from the IC Series device (default 30). This configuration statement is used in conjunction with the **timeout** statement to test active communications with the IC Series device. The value of the **interval** statement must be smaller than the value of **timeout** statement.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- timeout (Services) on page 5481
- timeout-action on page 5482

Copyright © 2016, Juniper Networks, Inc.

ip-address (Access Profile)

Syntax ip-address *address*;

Hierarchy Level [edit access profile *name* client *name* xauth]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address for the client.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

ip-user-mapping

Syntax

```
ip-user-mapping {
  discovery-method {
    wmi {
      event-log-scanning-interval seconds;
      initial-event-log-timespan hours;
    }
  }
}
```

Hierarchy Level [edit services user-identification active-directory domain]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Control how the SRX Series device accesses a domain controller in order to monitor and scan security event logs on the domain controller. By parsing the event log, the SRX Series gets IP address-to-user mappings. This process is part of the integrated user firewall feature. The **ip-user-mapping** statement is optional because WMI is the default discovery method and its properties have default values.

The other available method the SRX Series uses to retrieve address-to-user mapping information is manual (on-demand) probing of a domain PC.

Options

discovery-method—Method of discover IP address-to-user mappings.

wmi—Windows Management Instrumentation (WMI) is the discovery method used to access the domain controller.

event-log-scanning-interval *seconds*—Optional. Interval at which the SRX Series scans the event log on the domain controller.

Range: 5 through 60 seconds

Default: 10 seconds

initial-event-log-timespan *hours*—Optional. Time of the earliest event log on the domain controller that the SRX Series will initially scan. This argument applies to the initial deployment only. After WMIC and the user identification start working, the SRX Series scans only the latest event log.

Range: 1 through 168 hours

Default: 1 hour

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- active-directory-access on page 5400
- clear services user-identification active-directory-access on page 5524
- request services user-identification active-directory-access ip-user-probe on page 5528
- user-identification (Services) on page 5499
- show services user-identification active-directory-access statistics on page 5572

Copyright © 2016, Juniper Networks, Inc.

- traceoptions (Active Directory Access) on page 5488

ldap-options

Syntax

```
ldap-options {
  assemble {
    common-name common-name;
  }
  base-distinguished-name base-distinguished-name;
  revert-interval seconds;
  search {
    admin-search {
      distinguished-name distinguished-name;
      password password;
    }
    search-filter filter-name;
  }
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure LDAP authentication options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

Copyright © 2016, Juniper Networks, Inc.

ldap-server

Syntax

```
ldap-server hostname-or-address; {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify that the device use an LDAP server for authentication.

Options

server-address—Address of the LDAP authentication server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

level (Services)

Syntax

level [*brief* | *detail* | *extensive* | *verbose*];

Hierarchy Level

[edit services ssl traceoptions]

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

Description

Specify the level of debugging the output.

Options

- *brief*—Specify brief debugging output.
- *detail*—Specify detailed debugging output.
- *extensive*—Specify extensive debugging output.
- *verbose*—Specify verbose debugging output.

Required Privilege Level

services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

lifetime-seconds (Security IKE)

Syntax lifetime-seconds *seconds*;

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Default value modified in Junos OS Release 10.2.

Description Specify the lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

Options ***seconds***—Lifetime of the IKE SA.
Range: 180 through 86,400 seconds
Default: 28,800 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041
- *Understanding User Authentication Methods*

link (Access)

Syntax link *pool-name*;

Hierarchy Level [edit access address-assignment pool *pool-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the name of the secondary address-assignment pool that is linked to a primary address-assignment pool. The secondary pool provides a backup pool for local address assignment.

Options *pool-name*—Name of the address assignment pool.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

Copyright © 2016, Juniper Networks, Inc.

local-authentication-table

Syntax local-authentication-table priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Release 12.1 of Junos OS. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description An authentication table created on the SRX Series device using the **request security user-identification local-authentication-table add** command.

Options

priority ***priority***—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the local authentication table is 100.

Setting the priority value of the local authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

log (Services)

Syntax

```
log {
    all;
    errors;
    info;
    sessions-allowed;
    sessions-dropped;
    sessions-ignored;
    sessions-whitelisted;
    warning;
}
```

Hierarchy Level [edit services ssl proxy profile *profile-name* actions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the logging actions.

Options

- **all**—Log all events.
- **errors**—Log all error events.
- **info**—Log all information events.
- **sessions-allowed**—Log SSL session allowed events after an error.
- **sessions-dropped**—Log only SSL session dropped events.
- **sessions-ignored**—Log session ignored events.
- **sessions-whitelisted**—Log SSL session whitelisted events.
- **warning**—Log all warning events.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

login (Access)

Syntax login *string*;

Hierarchy Level [edit access firewall-authentication pass-through default-profile *profile-name* (ftp | http | telnet) banner]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the login banner for users using FTP, HTTP, and Telnet during the authentication process.

Options ***string***—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text within spaces or special characters—for example quotation marks (" ").

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

match (Services)

Syntax match *match*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the regular expression for lines to be logged.

Options **match**—Specify the regular expression for lines to be logged.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

network (Access)

Syntax network

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6)]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IPv4 network address for the pool. This attribute is mandatory. For an IPv6 pool, you will set the IPv6 network prefix.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

no-remote-trace (Services)

Syntax no-remote-trace;

Hierarchy Level [edit services ssl traceoptions]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Disable remote tracing.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540

Copyright © 2016, Juniper Networks, Inc.

pass-through

Syntax

```
pass-through {
  default-profile profile-name;
  ftp {
    banner {
      fail string;
      login string;
      success string;
    }
  }
  http {
    banner {
      fail string;
      login string;
      success string;
    }
  telnet {
    banner {
      fail string;
      login string;
      success string;
    }
}
```

Hierarchy Level [edit access firewall-authentication]

Release Information Statement introduced in Junos OS Release 8.5.
HTTPS for pass-through authentication is supported on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Configure pass-through firewall user authentication, when a host or user from one zone needs to access a protected resource in another zone. A user must use an FTP, Telnet, or HTTP client to access the IP address of the protected resource and get authenticated by the firewall. The device uses FTP, Telnet, and HTTP to collect username and password information. Subsequent traffic from the user or host is allowed or denied based on the result of this authentication. After the user is authenticated, the firewall proxies the connection.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

Copyright © 2016, Juniper Networks, Inc.

password (Access)

Syntax password *password*;

Hierarchy Level [edit access ldap-options search admin-search],
[edit access profile *profile-name* ldap-options search admin-search]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the plain-text password for the administrative user. This password is used in the bind for performing the LDAP search.

Options ***password***—Administrative user password.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

password (Services)

Syntax password *password*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the password that the SRX Series device should send to the IC Series device to establish communications. The SRX Series device sends the password in its first message to the IC Series device.

This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- ca-profile (Services) on page 5408
- server-certificate-subject on page 5470

Copyright © 2016, Juniper Networks, Inc.

permit (Security Policies)

Syntax

```
permit {
  application-services {
    application-firewall {
      rule-set rule-set-name;
    }
    application-traffic-control {
      rule-set rule-set-name;
    }
    gprs-gtp-profile profile-name;
    gprs-sctp-profile profile-name;
    idp;
    redirect-wx | reverse-redirect-wx;
    ssl-proxy {
      profile-name profile-name;
    }
    uac-policy {
      captive-portal captive-portal;
    }
    utm-policy policy-name;
  }
  destination-address {
    drop-translated;
    drop-untranslated;
  }
  firewall-authentication {
    pass-through {
      access-profile profile-name;
      client-match user-or-group-name;
      ssl-termination-profile profile-name;
      web-redirect;
      web-redirect-to-https;
    }
    user-firewall {
      access-profile profile-name;
      domain domain-name
      ssl-termination-profile profile-name;
    }
    web-authentication {
      client-match user-or-group-name;
    }
  }
  services-offload;
  tcp-options {
    sequence-check-required;
    syn-check-required;
  }
  tunnel {
    ipsec-group-vpn group-vpn;
    ipsec-vpn vpn-name;
    pair-policy pair-policy;
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **tcp-options** added in Junos OS Release 10.4. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.

Description Specify the policy action to perform when packets match the defined criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

 Copyright © 2016, Juniper Networks, Inc.

policies

Syntax

```
policies {
  default-policy (deny-all | permit-all);
  from-zone zone-name to-zone zone-name {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
        deny;
        log {
          session-close;
          session-init;
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
}
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
```


Copyright © 2016, Juniper Networks, Inc.

```
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    from-zone {
      [zone-name];
      any;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
    to-zone {
      [zone-name];
      any;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
```

```
                captive-portal captive-portal;
              }
              utm-policy policy-name;
            }
            destination-address {
              drop-translated;
              drop-untranslated;
            }
            firewall-authentication {
              pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
              }
              web-authentication {
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable) ;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Copyright © 2016, Juniper Networks, Inc.

Release Information

Statement introduced in Junos OS Release 8.5.
Support for the **services-offload** option added in Junos OS Release 11.4.
Support for the **source-identitiy** option added in Junos OS Release 12.1.
Support for the **description** option added in Junos OS Release 12.1.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.
Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.
Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description

Configure network security policies.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

Copyright © 2016, Juniper Networks, Inc.

pool (Access)

Syntax

```
pool pool-name {
  family {
    inet {
      dhcp-attributes {
        boot-file boot file name;
        boot-server boot server name;
        domain-name domain name;
        grace-period seconds;
        maximum-lease-time (seconds | infinite);
        name-server ipv4-address;
        name-server address;
        netbios-node-type (b-node | h-node | m-node | p-node);
        option dhcp option-identifier-code;
        option-match {
          option-82 {
            circuit-id match-value;
            remote-id match-value;
          }
        }
        router IPv4 address;
        server-identifier IP address;
        tftp-server server name;
        wins-server IPv4 address;
      }
      host hostname;
      network network address;
      range range-name {
        high upper-limit;
        low lower-limit;
      }
      xauth-attributes {
        primary-dns IP address;
        primary-wins IP address;
        secondary-dns IP address;
        secondary-wins IP address;
      }
    }
    inet6 {
      dhcp-attributes {
        dns-server IPv6-address;
        grace-period seconds;
        maximum-lease-time seconds;
        option dhcp-option-identifier-code;
        sip-server-address IPv6-address;
        sip-server-domain-name domain-name;
      }
      prefix IPv6-network-prefix;
      range range-name {
        high upper-limit;
        low lower-limit;
        prefix-length delegated-prefix-length;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
    link pool-name;
}
```

Hierarchy Level [edit access address-assignment]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the name of an address assignment pool.

The remaining statements are explained separately. See CLI Explorer.

Options *pool-name*—Name assigned to the address-assignment pool.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- Obtaining Username and Role Information Through Firewall Authentication on page 1136

port (Access LDAP)

Syntax port *port-number*;

Hierarchy Level [edit access ldap-server *server-address*],
[edit access profile *profile-name* ldap-server *server-address*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the port number on which to contact the LDAP server.

Options *port-number*—Port number on which to contact the LDAP server.

Default: 389

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

port (Services)

Syntax port *port-number*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the port on the IC Series device through which the SRX Series device should establish connections (default 11123). Possible values for this statement range from 1 through 65,535.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- interface (Services) on page 5435
- password (Services) on page 5446

preferred-ciphers

Syntax preferred-ciphers (custom | medium | strong | weak);

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Select the preferred ciphers.

Options
- **custom**—Configure custom cipher suite and order of preference.
- **medium**—Use ciphers with key strength of 128 bits or greater.
- **strong**—Use ciphers with key strength of 168 bits or greater.
- **weak**—Use ciphers with key strength of 40 bits or greater.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- SSL Proxy Overview on page 531

Copyright © 2016, Juniper Networks, Inc.

prefix (Access IPv6)

Syntax prefix *ipv6-network-prefix*;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet6]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.

Options *ipv6-network-prefix*—IPv6 prefix.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

protocol-version

Syntax protocol-version (all | tls1 | tls11 | tls12);

Hierarchy Level [edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10. The **tls11** and **tls12** options are introduced in 15.1X49-D30.

Description Specify the accepted SSL protocol version.

Options
- **all**—Accept all versions of TLS.
- **TLS version 1.0**—Accept TLS version 1.0. It provides secure communication over networks by providing privacy and data integrity between communicating applications
- **TLS version 1.1**—Accept TLS version 1.1. This enhanced version of TLS provides protection against cipher-block chaining (CBC) attacks.
- **TLS version 1.2** —Accept TLS version 1.2. This enhanced version of TLS provides improved flexibility for negotiation of cryptographic algorithms.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- SSL Proxy Overview on page 531

Copyright © 2016, Juniper Networks, Inc.

radius-options (Access)

Syntax

```
radius-options {
  revert-interval seconds;
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Configure RADIUS options.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

radius-server (Access)

Syntax

```
radius-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-name*]

Release Information

Statement modified in Junos OS Release 8.5.

Description

Configure RADIUS for Layer 2 Tunneling Protocol (L2TP) or Point-to-Point Protocol (PPP) authentication.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—Address of the RADIUS authentication server.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

range (Access)

Syntax

```
range range-name {
    high upper-limit;
    low lower-limit;
    prefix-length delegated-prefix-length;
}
```

Hierarchy Level

[edit access address-assignment pool pool-name family inet6]
[edit access address-assignment pool pool-name family inet]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Configure an IP name range used within an address-assignment pool. For IPv4, you do not create a prefix-length.

Options

- *range-name*—Name of the range.
- **high***upper-limit*—Upper limit of IPv6 address range.
- **low***lower-limit*—Lower limit of IPv6 address range.
- **prefix-length***delegated-prefix-length*—IPv6 delegated prefix length.

Required Privilege Level

access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

redirect-traffic

Syntax redirect-traffic (all | unauthenticated);

Hierarchy Level [edit services unified-access-control captive-portal *policy*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify to redirect traffic destined for protected sources to the IC Series device. You can choose to redirect all traffic or only unauthenticated traffic.

Options

- **all**—Redirect all traffic destined for the protected sources to the IC Series device. Specify this option if you want to redirect all traffic (IPsec or source IP) to the currently connected IC Series device or to an IP address or domain name that you specify in a redirect URL.
- **unauthenticated**—Redirect unauthenticated traffic destined for the protected sources to the IC Series device. Select this option if your deployment uses source IP only or a combination of source IP and IPsec. The Junos OS Enforcer redirects clear-text traffic from unauthenticated users to the currently connected IC Series device or to an IP address or domain name that you specify in a redirect URL.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

redirect-url

Syntax redirect-url *url*;

Hierarchy Level [edit services unified-access-control captive-portal *policy*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify to redirect traffic destined for protected sources to a specified URL.

You can configure the following options in the redirect URL string:

- ***%dest-url%***—Specifies the protected resource which the user is trying to access.
- ***%enforcer-id%***—Specifies the ID assigned to the Junos OS Enforcer by the IC Series device.
- ***%policy-id%***—Specifies the encrypted policy ID for the security policy that redirected the traffic.
- ***%dest-ip%***—Specifies the IP address or hostname of the protected resource that the user is trying to access.
- ***%ic-ip%***—Specifies the IP address or hostname of the IC Series device to which the Junos OS Enforcer is currently connected.

If you do not specify the redirect URL, the Junos OS Enforcer uses the following default configuration:

```
https://%ic-ip%/?target = %dest-url%  &enforcer = %enforcer-id%  &policy =
%policy-id%  &dest-ip = %dest-ip%
```

NOTE: The maximum size of a redirect payload is 1450 bytes. The size of the redirect URL is restricted to 1407 bytes (excluding a few HTTP headers). If a user accesses a destination URL that is larger than 1407 bytes, the Infranet Controller authenticates the payload, calculates the exact length of the redirect URL, and trims the destination URL so that it can fit into the redirect URL. The destination URL can be fewer than 1407 bytes based on what else is present in the redirect URL (for example, policy ID). The destination URL in the default redirect URL is trimmed so that the redirect packet payload size is limited to 1450 bytes. If the length of the payload is larger than 1450 bytes, the excess length is trimmed and the user is directed to the destination URL that has been resized to 1450 bytes.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

 Copyright © 2016, Juniper Networks, Inc.

retry (Access LDAP)

Syntax retry *attempts*;

Hierarchy Level [edit access ldap-server *server-address*],
[edit access profile *profile-name* ldap-server *server-address*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the number of retries that a device can attempt to contact an LDAP server.

Options ***attempts***—Number of retries that the device is allowed to attempt to contact an LDAP server.

Range: 1 through 10

Default: 3

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

retry (Access RADIUS)

Syntax retry *attempts*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the number of retries that a device can attempt to contact a RADIUS authentication server.

Options ***attempts***—Number of retries that the device is allowed to attempt to contact a RADIUS server.

Range: 1 through 10

Default: 3

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

revert-interval (Access LDAP)

Syntax revert-interval *seconds*;

Hierarchy Level [edit access ldap-options],
[edit access profile *profile-name* ldap-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the amount of time that elapses before the primary server is contacted if a backup server is being used.

Options ***seconds***—Number of seconds that elapse before the primary server is contacted.

Range: 60 through 4,294,967,295 seconds

Default: 600 seconds

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

revert-interval (Access RADIUS)

Syntax revert-interval *seconds*;

Hierarchy Level [edit access radius-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the amount of time that elapses before the primary server is contacted if a backup server is being used.

Options ***seconds***—Number of seconds that elapse before the primary server is contacted.

Range: 60 through 4,294,967,295 seconds

Default: 600 seconds

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

root-ca (Services)

Syntax root-ca *root-certificate*;

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Root certificate for interdicting server certificates in proxy mode.

Options *root-ca-name*—Specify root certificate for interdicting server certificates in proxy mode.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

routing-instance (Access LDAP)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access ldap-server *server-address*],
[edit access profile *profile-name* ldap-server *server-address*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the routing instance used to send LDAP packets to the LDAP server. A routing instance is a collection of routing tables, the interfaces contained in the routing tables, and the routing protocol parameters that control the information in the routing tables.

Options ***routing-instance-name***—Name of the routing instance.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

routing-instance (Access RADIUS)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure the routing instance used to send RADIUS packets to the RADIUS server. A routing instance is a collection of routing tables, the interfaces contained in the routing tables, and the routing protocol parameters that control the information in the routing tables.

Options *routing-instance-name*—Name of the routing instance.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

search

Syntax
```
search {
  admin-search {
    distinguished-name distinguished-name;
    password password;
  }
  search-filter filter-name;
}
```

Hierarchy Level [edit access ldap-options],
[edit access profile *profile-name* ldap-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that a search is used to get a user's LDAP distinguished name (DN). The search is performed based on the search filter and the part typed in by the user during authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

search-filter

Syntax search-filter *filter-name*;

Hierarchy Level [edit access ldap-options search],
[edit access profile *profile-name* ldap-options search]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that a search filter is used to find the user's LDAP distinguished name (DN). For example, a filter of **cn** specifies that the search matches a user whose common name is the username.

Options ***filter-name***—Name of the filter used to find the user's distinguished name.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

secret (Access Profile)

Syntax secret *password*;

Hierarchy Level [edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify the RADIUS secret password, which is shared between the router and the RADIUS server. The device uses this secret to encrypt the user's password that is sent to the RADIUS server.

Options ***password***—RADIUS secret. Maximum length is 256 characters.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

securid-server

Syntax

```
securid-server server-name {
  configuration-file filepath;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 9.1.

Description Configure SecurID server for SecurID authentication type.

Options The remaining statements are explained separately. See CLI Explorer.

NOTE: You can configure only one SecurID server. SecurID challenges are not yet supported.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

separator

Syntax separator *special-character*;

Hierarchy Level [edit access profile *profile-name* client-name-filter *client-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a character to identify where stripping of characters occurs in a client name. Stripping removes characters to the right of each instance of the specified character, plus the character itself. The stripping begins with the rightmost separator character.

Use the **separator** statement with the **count** statement to determine which characters in a client name are stripped. If the specified number of separator characters (count) exceeds the actual number of separator characters in the client name, stripping stops at the last available separator character.

Options ***special-character***—Character used to identify where to start the stripping of characters in a client name.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

server-certificate (Services)

Syntax server-certificate *server-certificate*;

Hierarchy Level [edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the local certificate identifier.

Options **server-certificate**—Specify the name of the local certificate identifier.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

server-certificate-subject

Syntax server-certificate-subject *subject*;

Hierarchy Level [edit services unified-access-control infranet-controller *hostname*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Optionally specify the full subject name of the certificate that the SRX Series device should use to validate the IC Series device's server certificate.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- ca-profile (Services) on page 5408
- password (Services) on page 5446

session-options (Access Profile)

Syntax
```
session-options {
  client-group [group-names];
  client-idle-timeout minutes;
  client-session-timeout minutes;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define options that control a user's session after successful authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

size (Services)

Syntax size *size*;

Hierarchy Level [edit services ssl traceoptions file *file-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the maximum trace file size.

Options **size**—Specify the maximum trace file size.

Range: 10,240 to 1,073,741,824.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

source-address (Access LDAP)

Syntax source-address *source-address*;

Hierarchy Level [edit access ldap-server server-address],
[edit access profile *profile-name* ldap-server *server-address*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure a source address for each configured LDAP server. Each LDAP request sent to a LDAP server uses the specified source address.

Options ***source-address***—Valid IP address configured on one of the device interfaces.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

source-address (Access RADIUS)

Syntax source-address *source-address*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options *source-address*—Valid IP address configured on one of the device interfaces.

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

ssl (Services)

Syntax

```
ssl {
  initiation {
    profile profile-name {
      actions {
        ignore-server-auth-failure;
      }
      client-certificate;
      custom-ciphers [cipher];
      enable-flow-tracing;
      enable-session-cache;
      preferred-ciphers (custom | medium | strong | weak);
      protocol-version (all | tls1 | tls11 | tls12);
      trusted-ca (all | [ca-profile] );
    }
  }
  proxy {
    global-config {
      session-cache-timeout seconds;
    }
    profile profile-name {
      actions {
        crl {
          disable;
          if-not-present (allow | drop);
          ignore-hold-instruction-code;
        }
        disable-session-resumption;
        ignore-server-auth-failure;
        log {
          all;
          errors;
          info;
          sessions-allowed;
          sessions-dropped;
          sessions-ignored;
          sessions-whitelisted;
          warning;
        }
        renegotiation {
          (allow | allow-secure | drop);
        }
      }
      custom-ciphers [cipher];
      enable-flow-tracing;
      preferred-ciphers (custom | medium | strong | weak);
      root-ca root-certificate;
      trusted-ca (all | [ca-profile] );
      whitelist [global-address-book-addresses];
    }
  }
  termination {
    profile profile-name {
```

Copyright © 2016, Juniper Networks, Inc.

```
        custom-ciphers [cipher];
        enable-flow-tracing;
        enable-session-cache;
        preferred-ciphers (custom | medium | strong | weak);
        protocol-version (all | tls1 | tls11 | tls12);
        server-certificate certificate-identifier;
      }
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      level [brief | detail | extensive | verbose];
      no-remote-trace;
    }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

ssl-termination-profile

Syntax ssl-termination-profile *profile-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the SSL termination profile used for SSL offloading.

Options ***profile-name***—Specify the name of the SSL termination profile used to the SSL offload.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

success

Syntax success *string*;

Hierarchy Level [edit access firewall-authentication pass-through default-profile *name* (ftp | http | telnet) banner],
[edit access firewall-authentication web-authentication]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the banner (message) that users see when trying to connect using FTP, HTTP, or Telnet after successful authentication.

Options *string*—Banner text. Maximum length of the message text is 250 characters. Enclose the banner text within spaces or special characters, such as quotation marks (“ “).

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

telnet (Access)

Syntax

```
telnet {
  banner {
    fail string;
    login string;
    success string;
  }
}
```

Hierarchy Level [edit access firewall-authentication pass-through]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure banners for Telnet login prompt, successful authentication, and failed authentication.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

termination (Services)

Syntax

```
termination {
  profile profile-name {
    custom-ciphers [cipher];
    enable-flow-tracing;
    enable-session-cache;
    preferred-ciphers (custom | medium | strong | weak);
    protocol-version (all | tls1 | tls11 | tls12);
    server-certificate certificate-identifier;
  }
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the configuration for Secure Socket Layer (SSL) termination support service.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

test-only-mode

Syntax test-only-mode (true | false):

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the device in test-only mode to log access decisions from the IC Series device without actually enforcing the decisions. When configured in test-only mode, the SRX Series device enables all UAC traffic to go through so you can test the implementation without impeding traffic.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

then (Security Policies)

Syntax

```
then {
    count {
        alarm {
            per-minute-threshold number;
            per-second-threshold number;
        }
    }
    deny;
    log {
        session-close;
        session-init;
    }
    permit {
        application-services {
            application-firewall {
                rule-set rule-set-name;
            }
            application-traffic-control {
                rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
                profile-name profile-name;
            }
            uac-policy {
                captive-portal captive-portal;
            }
            utm-policy policy-name;
        }
        destination-address {
            drop-translated;
            drop-untranslated;
        }
        firewall-authentication {
            pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
            }
            user-firewall {
                access-profile profile-name;
                domain domain-name
                ssl-termination-profile profile-name;
            }
            web-authentication {
                client-match user-or-group-name;
            }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    services-offload;
    tcp-options {
      initial-tcp-mss mss-value;
      reverse-tcp-mss mss-value;
      sequence-check-required;
      syn-check-required;
    }
    tunnel {
      ipsec-group-vpn group-vpn;
      ipsec-vpn vpn-name;
      pair-policy pair-policy;
    }
  }
  reject;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Specify the policy action to be performed when packets match the defined criteria.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

timeout (Access LDAP)

Syntax timeout *seconds*;

Hierarchy Level [edit access ldap-server *server-address*]
[edit access profile *profile-name* ldap-server *server-address*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the amount of time that the local device waits to receive a response from an LDAP server.

Options *seconds*—Amount of time to wait.

Range: 1 through 90 seconds

Default: 3 seconds

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

timeout (Access RADIUS)

Syntax timeout *seconds*;

Hierarchy Level [edit access radius-server *server-address*]
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure the amount of time that the local device waits to receive a response from a RADIUS server.

Options *seconds*—Amount of time to wait.

Range: 1 through 90 seconds

Default: 3 seconds

Required Privilege Level secret—To view this statement in the configuration.
secret-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

timeout (Services)

Syntax timeout *seconds*;

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the value, in seconds, that the SRX Series device should wait to get a heartbeat response from an IC Series UAC Appliance (default is 300). If the SRX Series device does not receive it in the specified time, it takes the action specified by the **timeout-action** configuration statement. It also tries again to make a connection to the IC Series appliance. After the second failed attempt, the SRX Series device fails over to the next IC Series appliance in the cluster. The SRX Series device continues trying to reach IC Series appliances in the cluster until a connection is established.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series appliance. When working with a cluster of IC Series appliances, the Junos OS Enforcer connects to one at a time, failing over to other IC Series appliances in the cluster as required.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- interval (Services) on page 5435
- timeout-action on page 5482

Copyright © 2016, Juniper Networks, Inc.

timeout-action

Syntax timeout-action (close | no-change | open):

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify what the SRX Series device should do when a timeout occurs and the device cannot connect to an Infranet Enforcer.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series UAC Appliance.

Options

- **close**—Close existing sessions and block any further traffic. This is the default option.
- **no-change**—Preserve existing sessions and require authentication for new sessions.
- **open**—Preserve existing sessions and allow new sessions access.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- interval (Services) on page 5435
- timeout (Services) on page 5481

 Copyright © 2016, Juniper Networks, Inc.

to-zone (Security Policies)

Syntax

```
to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```

Copyright © 2016, Juniper Networks, Inc.

```
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

Hierarchy Level [edit security policies from-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** and **junos-host** options added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10.

Description Specify a destination zone to be associated with the security policy.

Options

- **zone-name**—Name of the destination zone object.
- **junos-host**—Default security zone for self-traffic of the device.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Access)

Syntax

```
traceoptions {
  file filename {
    files number;
    match regular-expression;
    size maximum-file-size;
    <world-readable | no-world-readable>;
  }
  flag flag;
}
```

Hierarchy Level [edit access firewall-authentication]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define Routing Engine firewall authentication tracing options.

Options

- file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.
- files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed to *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.
- If you specify a maximum number of files, you also must specify a maximum file size with the size option and a filename.

  Range: 2 through 1000 files
  Default: 10 files

- match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.
- size *maximum-file-size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named**trace-file** reaches this size, it is renamed **trace-file.0**. When the file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1**, and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
- If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.

  Syntax: *x***k** to specify KB, *x***m** to specify MB, or *x***g** to specify GB.
  Range: 10 KB through 1 GB
  Default: 128 KB

- **world-readable | no-world-readable**—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option

Copyright © 2016, Juniper Networks, Inc.

enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag** *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags.
- **all**—All tracing operations
- **authentication**—Trace authentication events
- **configuration**—Trace configuration events
- **setup**—Trace setup of firewall authentication service

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Active Directory Access)

Syntax

```
traceoptions {
  file filename ;
  flag {
    active-directory-authentication;
    all;
    configuration;
    db;
    ip-user-mapping;
    ip-user-probe;
    ipc;
    user-group-mapping;
    wmic;
  }
  level {
    all
    error
    info
    notice
    verbose
    warning
  }
  no-remote-trace;
}
```

Hierarchy Level [edit services user-identification active-directory-access]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Define Active Directory trace options for the integrated user firewall feature.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

flag—Trace the operation or operations to perform on the integrated user firewall. To specify more than one trace operation, include multiple flag statements.

active-directory-authentication—Trace the building of and modifications to the Active Directory authentication table.

all—Trace everything.

configuration—Trace configuration events.

db—Trace the database.

ip-user-mapping—Trace the ip-user-mapping module.

ip-user-probe—Trace PC client probing.

ipc—Trace communication events with the Packet Forwarding Engine.

user-group-mapping—Trace the process of getting user-to-group-mapping.

Copyright © 2016, Juniper Networks, Inc.

wmic—Trace the Windows Management Instrumentation Client process.

level—Level of trace operation to perform.

all—Match all levels.

error—Match error conditions.

info—Match informational messages.

notice—Match conditions that should be handled specially.

verbose—Match verbose messages.

warning—Match warning messages.

no-remote-trace—Disallow tracing from a remote device.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- active-directory-access on page 5400
- user-identification (Services) on page 5499
- Overview of Integrated User Firewall on page 5309

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Firewall Authentication)

Syntax

```
traceoptions {
  flag {
    all <detail | extensive | terse>;
    authentication <detail | extensive | terse>;
    proxy <detail | extensive | terse>;
  }
}
```

Hierarchy Level [edit security firewall-authentication]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define data-plane firewall authentication tracing options.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Enable all tracing operations
  - **authentication**—Trace data-plane firewall authentication events
  - **proxy**—Trace data-plane firewall authentication proxy events
- **detail**—Display moderate amount of data in trace.
- **extensive**—Display extensive amount of data in trace.
- **terse**—Display minimum amount of data in trace.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Services SSL)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  level [brief | detail | extensive | verbose];
  no-remote-trace;
}
```

Hierarchy Level [edit services ssl]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the trace file information.

Options

- *file-name*—Specify the name of file in which to write trace information.
- **files**—Specify the maximum number of trace files. Range: 2 to 1000.
- **match**—Specify the regular expression for lines to be logged.
- **no-world-readable size**—Do not allow any user to read the log file.
- **size**—Specify the maximum trace file size. Range: 10,240 to 1,073,741,824.
- **world-readable**—Allow any user to read the log file.

Required Privilege Level
services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Services UAC)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit services unified-access-control]

Release Information Statement introduced in Junos OS Release 9.4.

Description Define Unified Access Control (UAC) tracing options.

Use this statement when you are configuring the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series UAC Appliance.

Options **flag**—Trace operation to perform. To specify more than one trace option, include multiple flag statements.

- **all**—Trace with all flags enabled
- **config**—Trace configuration information for all UAC-related configurations. This includes all configuration controlled through the **unified-access-control** statements at the **edit services** hierarchy level. It also includes other standard Junos OS configurations required for UAC enforcement such as zones, policies, and interfaces.
- **connect**—Trace communications between the Junos OS Enforcer and the IC Series appliance, including SSL handshakes and timeouts.
- **ipc**—Trace interprocess communications. Use this option to trace communications between the Routing Engine (RE) and the UACD enforcement plugin inside the Packet Forwarding Engine (PFE).

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- Understanding Unified Access Control on page 5269
- Acquiring User Role Information from an Active Directory Authentication Server on page 5269

Copyright © 2016, Juniper Networks, Inc.

trusted-ca (Services)

Syntax trusted-ca (all | [*ca-profile*]);

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]
[edit services ssl initiation profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the list of trusted certificate authority profiles.

Options
- *trusted-ca-name*—Specify the certificate authority profile name.
- **all**—Select all certificate authority profiles.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

uac-policy (Application Services)

Syntax
```
uac-policy {
  captive-portal captive-portal;
}
```

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement modified in Junos OS Release 9.4.

Description Enable Unified Access Control (UAC) for the security policy. This statement is required when you are configuring the SRX Series device to act as a Junos OS Enforcer in a UAC deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series UAC Appliance .

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding User Role Firewalls on page 1127
- Example: Configuring a User Role Firewall on an SRX Series Device on page 1138

Copyright © 2016, Juniper Networks, Inc.

uac-service

Syntax

```
uac-service {
  command binary-file-path;
  disable;
  failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the unified access control daemon process.

Options

- **command *binary-file-path***—Path to the binary process.
- **disable**—Disable the unified access control daemon process.
- **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179
- System Configuration Statement Hierarchy on page 612

Copyright © 2016, Juniper Networks, Inc.

unified-access-control (Security)

Syntax unified-access-control priority *priority*;

Hierarchy Level [edit security user-identification authentication-source]

Release Information Statement introduced in Release 12.1 of Junos OS. Support for **disable** option dropped in Junos OS Release 12.1X47-D10.

Description An authentication table pushed from a configured authentication device, such as the Junos Pulse Access Control Service.

Options priority ***priority***—A unique value between 0 and 65535 that determines the sequence for searching multiple tables to retrieve a user role. Each authentication table is given a unique priority value. The lower the value, the higher the priority. A table with priority 120 is searched before a table with priority 200. The default priority value of the unified-access-control authentication table is 200.

Setting the priority value of the unified-access-control authentication table to 0 is equivalent to disabling the table and eliminating it from the search sequence.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- authentication-source on page 1234
- Understanding User Role Firewalls on page 1127
- Understanding the User Identification Table on page 1130

Copyright © 2016, Juniper Networks, Inc.

unified-access-control (Services)

Syntax

```
unified-access-control {
  captive-portal redirect-policy-name{
    redirect-traffic (all | unauthenticated);
    redirect-url redirect-url;
  }
  certificate-verification [ optional | required | warning ];
  infranet-contoller host-name {
    address ip-address;
    ca-profile [ca-profile];
    interface interface-name;
    password password;
    port port-number;
    server-certificate-subject subject;
  }
  interval seconds;
  test-only-mode;
  timeout seconds;
  timeout-action (close | no-change | open);
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      (no-world-readable | world-readable);
      size maximum-file-size;
    }
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 9.4.

Description Use this statement to configure the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series device.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

user-group-mapping

Syntax

```
user-group-mapping {
  ldap {
    address ip-address {
      port port;
    }
    authentication-algorithm {
      simple;
    }
    base base;
    ssl;
    user username {
      password password;
    }
  }
}
```

Hierarchy Level [edit services user-identification active-directory-access domain]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Configure the SRX Series device to connect to an LDAP server, so that the server can provide the SRX Series with user-to-group mappings. These mappings are used to implement the integrated user firewall feature. The domain controller acts as the LDAP server in typical customer scenarios.

Most of this statement is optional, because the default communication method is LDAP and most arguments have default values. Only the LDAP keyword and the base are required.

Options **ldap**—Required. LDAP is the protocol used to access the LDAP server to get user-to-group mappings.

address *ip-address*—Optional. Specify the IP address of the LDAP server. If no address is specified, the system uses one of the configured Active Directory domain controllers.

port *port*—Optional. Specify the port number of the LDAP server. If no port number is specified, the system uses port 389 for plaintext or port 636 for encrypted text.

authentication-algorithm—Optional. Specify the algorithm used while the SRX Series communicates with the LDAP server. The default method is Kerberos.

simple—Configure simple (plaintext) authentication method.

base *base*—Required. LDAP base distinguished name (DN).

ssl—Optional. Enable Secure Sockets Layer (SSL) to ensure secure transmission with the LDAP server. Disabled by default, which means that the password is sent in plaintext.

Copyright © 2016, Juniper Networks, Inc.

user *username*—Optional. Username of the LDAP account. If no username is specified, the system will use the configured domain controller's username.

password *password*—Optional. Specify the password for the account. If no password is specified, the system uses the configured domain controller's password.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- active-directory-access on page 5400
- clear services user-identification active-directory-access on page 5524
- show services user-identification active-directory-access statistics on page 5572
- show services user-identification active-directory-access user-group-mapping on page 5575
- traceoptions (Active Directory Access) on page 5488
- user-identification (Services) on page 5499
- LDAP Functionality in Integrated User Firewall on page 5317

Copyright © 2016, Juniper Networks, Inc.

user-identification (Services)

Syntax

```
user-identification {
  active-directory-access {
    domain domain-name {
      user username;
        password password;
      domain-controller domain-controller-name {
        address domain-controller-address;
      }
      ip-user-mapping {
        discovery-method {
          wmi {
            event-log-scanning-interval seconds;
            initial-event-log-timespan hours;
          }
        }
      }
      user-group-mapping {
        ldap {
          address ip-address {
            port port;
          }
          authentication-algorithm {
            simple;
          }
          base base;
          ssl;
          user username {
            password password;
          }
        }
      }
    }
    authentication-entry-timeout minutes;
    filter {
      include address;
      exclude address;
    }
    no-on-demand-probe;
    wmi-timeout seconds;
    traceoptions {
      file file;
      flag {
        active-directory-authentication;
        all;
        configuration;
        db;
        ip-user-mapping;
        ip-user-probe;
        ipc;
        user-group-mapping;
        wmic;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
        level {
          all;
          error;
          info;
          notice;
          verbose;
          warning;
        }
        no-remote-trace;
      }
    }
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 12.1X47-D10.

Description Configure the integrated user firewall feature, including access to the Active Directory domain and domain controller, IP address-to-user mapping, and user-to-group mapping. One or two Active Directories are allowed under one domain. The IP address-to-user mapping and user-to-group mapping are configured per domain.

Options

authentication-entry-timeout *minutes*—Timeout interval starting from the Active Directory/domain controller login time, the last active session, or the last successful probe. A setting of 0 means the authentication does not need a timeout. We recommend that you configure a setting of 0 when you disable on-demand-probe to prevent someone from accessing the Internet without logging in again.

Range: 10 through 1440 minutes

Default: 30 minutes

filter—Optional. Range of IP addresses that needs to be monitored or not monitored.

include *address*—Include IP address or range. Maximum of 20 addresses.

exclude *address*—Exclude IP address or range. Maximum of 20 addresses.

no-on-demand-probe—Do not use traffic to discover user. Default is disabled.

wmi-timeout *seconds*—Optional. Configures the number of seconds that the domain PC has to respond to the SRX Series device's query through WMI/DCOM.

- If the PC responds within that timeframe to the WMI query, the SRX creates an authentication entry for this PC.
- If the PC does not respond within that timeframe, the WMI query failed. In the case of a failed query, if the SRX had an authentication entry about the queried PC before the WMI query, that authentication entry is deleted. If the SRX had no authentication entry before the WMI query, the SRX does not create an authentication entry.

Range: 3 through 120 seconds

Default: 10 seconds

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- active-directory-access on page 5400
- traceoptions (Active Directory Access) on page 5488

web-authentication

Syntax
```
web-authentication {
  client-match user-or-group-name;
}
```

Hierarchy Level
[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication]

Release Information
Statement introduced in Junos OS Release 8.5.
HTTPS for Web authentication is supported on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description
Specify that the policy allows access to users who have previously been authenticated by Web authentication. Web authentication must be enabled on one of the addresses on the interface to which the HTTP or HTTPS request is redirected.

Options
client-match ***user-or-group*** —(Optional) Username or user group name.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

web-authentication (Access)

Syntax

```
web-authentication {
  banner {
    success string;
  }
  default-profile profile-name;
}
```

Hierarchy Level [edit access firewall-authentication]

Release Information Statement introduced in Junos OS Release 8.5.
HTTPS for Web authentication is supported on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Specify that users go through the Web authentication process. The user uses HTTP or HTTPS to access an IP address on the device that is enabled for Web authentication. In this scenario, the user does not use HTTP or HTTPS to access the IP address of the protected resource. The user is prompted for a username and password, which are verified by the device. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the results of this authentication. This method of authentication differs from pass-through authentication in that users need to access the protected resource directly after accessing the Web authentication IP address and being authenticated.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

 Copyright © 2016, Juniper Networks, Inc.

web-authentication (Interfaces)

Syntax

```
web-authentication {
  http;
  https;
  redirect-to-https;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family-name* address *address*]

Release Information Statement introduced in Junos OS Release 9.2.
Support for **https** and **redirect-to-https** introduced for high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Enable the Web authentication process for firewall user authentication.

Options

http—Enable HTTP service.

https—Enable authentication through HTTPS.

redirect-to-https—Redirect Web authentication to HTTPS.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Understanding Interfaces on page 2441

Copyright © 2016, Juniper Networks, Inc.

web-management (System Services)

Syntax

```
web-management {
  http {
    interfaces interface-names ;
    port port;
  }
  https {
    interfaces interface-names;
    local-certificate name;
    pki-local-certificate name;
    system-generated-certificate name;
    port port;
  }
  management url management url;
  session {
    idle-timout minutes;
    session-limit number;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (no-world-readable | world-readable);
    }
    flag flag;
    level level;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 9.0.
Support for **https** introduced for high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Configure settings for HTTP or HTTPS access. HTTP access allows management of the device using the J-Web interface. HTTPS access allows secure management of the device using the J-Web interface. With HTTPS access, communication is encrypted between your browser and the webserver for your device.

Options **control**—Disable the SBC process.

- **max-threads**—Maximum simultaneous threads to handle requests.

  Range: 0 through 16

http—Configure HTTP.

- **interface [*value*]**—Interface value that accepts HTTP access.

Copyright © 2016, Juniper Networks, Inc.

- **port *number***—TCP port for incoming HTTP connections.

  Range: 1 through 65,535

https—Configure HTTPS.

- **interface [*value*]**—Interface value that accept HTTP access.
- **port *number***—TCP port for incoming HTTP connections.

  Range: 1 through 65,535
- **local-certificate**—X.509 certificate to use from the configuration.
- **pki-local-certificate**—X.509 certificate to use from the PKI local store.
- **system-generated-certificate**—X.509 certificate generated automatically by the system.

management url *management url*—URL path for Web management access.

session—Configure the Web-management session.

- **idle-timout *minutes***—Default timeout of Web-management sessions in minutes.
- **session-limit *number***—Maximum number of Web-management sessions to allow.

Copyright © 2016, Juniper Networks, Inc.

traceoptions—Set the trace options.

- **file**—Configure the trace file information.
  - *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files** ***number***— Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** ***maximum file-size*** option.

    Range: 2 through 1000 files

    Default: 10 files
- **match** ***regular-expression***—Refine the output to include lines that contain the regular expression.
- **size** ***maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

  Range: 10 KB through 1 GB

  Default: 128 KB

  If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** ***number*** option.
- **(world-readable | no-world-readable)**— By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.
- **flag** ***flag***—Specify which tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags.
  - **all**—Trace all areas.
  - **configuration**—Trace configuration.
  - **dynamic-vpn**—Trace dynamic VPN events.
  - **init**—Trace the daemon init process.
  - **mgd**—Trace MGD requests.
  - **webauth**—Trace Web authentication requests.
- **level** ***level*** —Specify the level of debugging output.
  - **all**—Match all levels.
  - **error**—Match error conditions.

Copyright © 2016, Juniper Networks, Inc.

- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

- **no-remote-trace**—Disable remote tracing.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Firewall User Authentication Overview on page 5179
- *Dynamic VPN Overview*

web-redirect

Syntax
```
web-redirect;
```

Hierarchy Level
```
[edit security policies from-zone zone-name to-zone zone-name  policy policy-name then
  permit firewall-authentication pass-through]
```

Release Information
Statement introduced in Junos OS Release 8.5.

Description
Optionally, redirect HTTP requests to the device's internal webserver by sending a redirect HTTP response to the client system to reconnect to the webserver for user authentication. The interface on which the client's request arrived is the interface to which the request is redirected.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding User Role Firewalls on page 1127

Copyright © 2016, Juniper Networks, Inc.

web-redirect-to-https

Syntax web-redirect-to-https;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit firewall-authentication pass-through]

Release Information Statement introduced on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Redirect unauthenticated HTTP requests to the internal HTTPS webserver of the device.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- *UTM Feature Guide for Security Devices*
- Firewall User Authentication Overview on page 5179

whitelist (Services)

Syntax whitelist [global-address-book-addresses];

Hierarchy Level [edit services ssl proxy profile *profile-name*]
[edit services ssl termination profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10.

Description Specify the addresses exempted from the SSL proxy.

Options
- *whitelist-address*—Specify address from the global address book.

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation
- Configuring SSL Proxy on page 540
- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

wins-server (Access)

Syntax wins-server *address*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) xauth-attributes]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the wins-server IP address.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- Firewall User Authentication Overview on page 5179

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 229

Operational Commands

- clear network-access requests pending
- clear network-access requests statistics
- clear network-access securid-node-secret-file
- clear security firewall-authentication history
- clear security firewall-authentication history address
- clear security firewall-authentication history identifier
- clear security firewall-authentication users
- clear security firewall-authentication users address
- clear security firewall-authentication users identifier
- clear security user-identification local-authentication-table
- clear services user-identification active-directory-access
- request services user-identification active-directory-access active-directory-authentication-table delete
- request services user-identification active-directory-access domain-controller
- request services user-identification active-directory-access ip-user-probe
- show network-access requests pending
- show network-access requests statistics
- show network-access securid-node-secret-file
- show security firewall-authentication history
- show security firewall-authentication history address
- show security firewall-authentication history identifier
- show security firewall-authentication users
- show security firewall-authentication users address
- show security firewall-authentication users identifier
- show security policies
- show services unified-access-control authentication-table
- show services unified-access-control counters
- show services unified-access-control policies

Copyright © 2016, Juniper Networks, Inc.

- show services unified-access-control roles
- show services unified-access-control status
- show services user-identification active-directory-access active-directory-authentication-table
- show services user-identification active-directory-access domain-controller status
- show services user-identification active-directory-access statistics
- show services user-identification active-directory-access user-group-mapping

Copyright © 2016, Juniper Networks, Inc.

clear network-access requests pending

Syntax clear network-access requests pending
<index *index-number* >

Release Information Command introduced in Junos OS Release 8.5.

Description Clear or cancel all pending authentication requests.

Options
- none—Clear all network access requests pending.
- **index** *index-number* —Clear the specified authentication request. To display index numbers, use the **show network-access requests pending** command.

Required Privilege Level clear

Related Documentation
- Firewall User Authentication Overview on page 5179
- show network-access requests pending on page 5529

List of Sample Output clear network-access requests pending on page 5513

Output Fields This command produces no output.

Sample Output

The following example displays the network access requests that are pending, clears the requests, and displays the results of the clear operation:

clear network-access requests pending

```
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                  Status
1     Sun                   Processing
2     Sam                   Processed

user@host> clear network-access requests pending
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                  Status
1     Sun                   Cancelled by Admin
2     Sam                   Cancelled by Admin
```

Copyright © 2016, Juniper Networks, Inc.

clear network-access requests statistics

Syntax clear network-access requests statistics

Release Information Command introduced in Junos OS Release 8.5.

Description Clear general authentication statistics for the configured authentication type.

Required Privilege Level clear

Related Documentation
- Firewall User Authentication Overview on page 5179
- authentication-order (Access Profile) on page 3040
- show network-access requests statistics on page 5531

List of Sample Output clear network-access requests statistics on page 5514

Output Fields This command produces no output.

Sample Output

clear network-access requests statistics

```
user@host> clear network-access requests statistics
```


Copyright © 2016, Juniper Networks, Inc.

clear network-access securid-node-secret-file

Syntax clear network-access securid-node-secret-file

Release Information Command introduced in Junos OS Release 9.1.

Description Delete the node secret file for the SecurID authentication type.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179
- configuration-file on page 5413
- securid-server on page 5468
- show network-access securid-node-secret-file on page 5532

List of Sample Output clear network-access securid-node-secret-file on page 5515

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access securid-node-secret-file

```
user@host> clear network-access securid-node-secret-file
```

Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication history

Syntax clear security firewall-authentication history
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OSRelease 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear all firewall authentication history information.

Options

- **node**—(Optional) For chassis cluster configurations, clear all firewall authentication history on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
clear security firewall-authentication history on page 5516
clear security firewall-authentication history node 1 on page 5516

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication history

```
user@host> clear security firewall-authentication history
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history node 1

```
user@host> clear security firewall-authentication history node 1
node1:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication history address

Syntax clear security firewall-authentication history address *address*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear firewall authentication history for this source IP address.

Options

- **address** *address* —Source IP address for which to clear firewall authentication history.
- none—Clear all firewall authentication history for this address.
- **node**—(Optional) For chassis cluster configurations, clear firewall authentication history for this address on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output

clear security firewall-authentication history address 100.0.0.1 on page 5517
clear security firewall-authentication history address 100.0.0.1 node 1 on page 5517

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication history address 100.0.0.1

```
user@host> clear security firewall-authentication history address 10.0.0.1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history address 100.0.0.1 node 1

```
user@host> clear security firewall-authentication history address 10.0.0.1 node 1
node1:
--------------------------------------------------------------------------
```

Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication history identifier

Syntax clear security firewall-authentication history identifier *identifier*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear firewall authentication history information for the authentication with this identifier.

Options

- **identifier** *identifier* —Identification number of the authentication for which to clear authentication history.
- none—Clear all firewall authentication history information for the authentication with this identifier.
- **node**—(Optional) For chassis cluster configurations, clear firewall authentication history on a specific node for the authentication with this identifier.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
clear security firewall-authentication history identifier 2 on page 5518
clear security firewall-authentication history identifier 2 node 1 on page 5518

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication history identifier 2

```
user@host> clear security firewall-authentication history identifier 2
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication history identifier 2 node 1

```
user@host> clear security firewall-authentication history identifier 2 node 1
node1:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication users

Syntax clear security firewall-authentication users
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear firewall authentication tables for all users.

Options

- **node**—(Optional) For chassis cluster configurations, clear firewall authentication details for all users on a specific node.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179
- show security firewall-authentication users on page 3664

List of Sample Output

clear security firewall-authentication users on page 5520
clear security firewall-authentication users node 1 on page 5520

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication users

```
user@host> clear security firewall-authentication users node 1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users node 1

```
user@host> clear security firewall-authentication users node 1
node1:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication users address

Syntax clear security firewall-authentication users address *address*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear information about the users at the specified IP address that are currently authenticated.

Options

- **address** *address* —IP address for which to clear user firewall authentication information.
- none—Clear all the firewall authentication information for users at this IP address.
- **node**—(Optional) For chassis cluster configurations, clear user firewall authentication entries on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
clear security firewall-authentication users address 10.0.0.1 on page 5521
clear security firewall-authentication users address 10.0.0.1 node 1 on page 5521

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication users address 10.0.0.1

```
user@host> clear security firewall-authentication users address 10.0.0.1
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users address 10.0.0.1 node 1

```
user@host> clear security firewall-authentication users address 10.0.0.1 node 1
node1:
--------------------------------------------------------------------------
```

Copyright © 2016, Juniper Networks, Inc.

clear security firewall-authentication users identifier

Syntax clear security firewall-authentication users identifier *identifier*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Clear firewall authentication details about the user with this identification number.

Options

- none—Identification number of the user for which to clear authentication details.
- **node**—(Optional) For chassis cluster configurations, clear the firewall authentication details on a specific node (device) in the cluster for the user with this identification number.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all** —Clear all nodes.
  - **local** —Clear the local node.
  - **primary**—Clear the primary node.

Required Privilege Level clear

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
clear security firewall-authentication users identifier 2 on page 5522
clear security firewall-authentication users identifier 2 node 1 on page 5522

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security firewall-authentication users identifier 2

```
user@host> clear security firewall-authentication users identifier 2
node0:
--------------------------------------------------------------------------
node1:
--------------------------------------------------------------------------
```

Sample Output

clear security firewall-authentication users identifier 2 node 1

```
user@host> clear security firewall-authentication users identifier 2 node 1
node1:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

clear security user-identification local-authentication-table

Syntax clear security user-identification local-authentication-table

Release Information Command introduced in Junos OS Release 12.1.

Description This command removes all entries from the local authentication table.

Required Privilege Level clear

Related Documentation
- Understanding AppSecure Services on page 483
- Firewall User Authentication Overview on page 5179

List of Sample Output clear security user-identification local-authentication-table on page 5523

Output Fields When you enter this command, all entries are cleared from the local authentication table.

Sample Output

clear security user-identification local-authentication-table

```
user@host> clear security user-identification local-authentication-table
user@host> show security user-identification local-authentication-table all
Total entries: 0
```

Copyright © 2016, Juniper Networks, Inc.

clear services user-identification active-directory-access

Syntax

```
clear services user-identification active-directory-access
  (active-directory-authentication-table |
statistics (ip-user-mapping | ip-user-probe | user-group-mapping))
```

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Delete entries from the Active Directory authentication table or statistics related to integrated user firewall mappings.

Options

- **active-directory-authentication-table**—Remove all entries from the Active Directory authentication table.
- **statistics**—Remove the specified type of statistics:
  - **ip-user-mapping**—IP address-to-user mappings
  - **ip-user-probe**—PC probe statistics
  - **user-group-mapping**—User-to-group mappings

Required Privilege Level clear

Related Documentation

- ip-user-mapping on page 5437
- request services user-identification active-directory-access ip-user-probe on page 5528
- show services user-identification active-directory-access statistics on page 5572
- show services user-identification active-directory-access user-group-mapping on page 5575
- user-group-mapping on page 5497
- user-identification (Services) on page 5499

List of Sample Output clear services user-identification active-directory-access active-directory-authentication-table on page 5524

Output Fields This command produces no output.

Sample Output

The following example clears all entries in the Active Directory authentication table:

clear services user-identification active-directory-access active-directory-authentication-table

```
user@host> clear services user-identification active-directory-access
active-directory-authentication-table
```


Copyright © 2016, Juniper Networks, Inc.

request services user-identification active-directory-access active-directory-authentication-table delete

Syntax

request services user-identification active-directory-access
active-directory-authentication-table delete
(domain *name* | ip-address *ip-address* | group *group-name* <domain *name*> | user *name*
<domain *name*>

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Delete entries from the active directory authentication table by domain, address, group, or user. This command provides the network administrator with flexibility and control over the table entries beyond what is automatically added to or deleted from the table. For example, if a person leaves the company, the corresponding username can be deleted; after a department reorganization, a group can be deleted.

Options

- **domain** *name*—Delete the entries from the authentication table for the specified domain.
- **ip-address** *ip-address*—Delete the entry from the authentication table for the specified IP address.
- **group** *group-name*—Delete the entries from the authentication table for the specified group.
  - **domain** *name*—Delete the group only from the specified domain.
- **user** *name*—Delete the entries from the authentication table for the specified username.
  - **domain** *name*—Delete the user only from the specified domain.

Required Privilege Level

maintenance

Related Documentation

- show services user-identification active-directory-access active-directory-authentication-table on page 5565
- user-identification (Services) on page 5499
- Understanding Active Directory Authentication Tables on page 5312

List of Sample Output

request services user-identification active-directory-access active-directory-authentication-table delete user jsmith on page 5526

Output Fields

This command produces no output.

Sample Output

The following example clears all entries in the Active Directory authentication table for the specified user:

Copyright © 2016, Juniper Networks, Inc.

request services user-identification active-directory-access active-directory-authentication-table delete user jsmith

```
user@host> request services user-identification active-directory-access
active-directory-authentication-table delete user jsmith
```


Copyright © 2016, Juniper Networks, Inc.

request services user-identification active-directory-access domain-controller

Syntax request services user-identification active-directory-access domain-controller discovery domain *name*

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Discover and display the name and address of all domain controllers in the specified domain.

Options
- domain *name*—Name of the domain for which to get and display domain controller names and addresses.

Required Privilege Level maintenance

Related Documentation
- active-directory-access on page 5400
- show services user-identification active-directory-access domain-controller status on page 5569
- user-identification (Services) on page 5499

List of Sample Output request services user-identification active-directory-access domain-controller discovery domain <domain-name> on page 5527

Output Fields This command displays the discovered domain controllers.

Sample Output

request services user-identification active-directory-access domain-controller discovery domain <domain-name>

```
user@host> request services user-identification active-directory-access domain-controller
discovery domain example.net
Domain: example.net
 Domain controller: dc-example.net
   Address: 192.0.2.2
```

request services user-identification active-directory-access ip-user-probe

Syntax request services user-identification active-directory-access ip-user-probe address *ip-address* <domain *name*>

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Probe the PC at the specified IP address to get an authentication entry, which is used for the integrated user firewall feature. You can display the authentication table to see the results. If the probe succeeded, there will be a valid authentication entry. If the probe failed, there will be an invalid authentication entry.

Options

- address *ip-address*—Probe the PC at this IP address.

  domain *name*—Probe the IP address in the specified domain.

Required Privilege Level maintenance

Related Documentation

- clear services user-identification active-directory-access on page 5524
- show services user-identification active-directory-access active-directory-authentication-table on page 5565
- show services user-identification active-directory-access statistics on page 5572
- user-identification (Services) on page 5499

List of Sample Output show services user-identification active-directory-access active-directory-authentication-table address <ip-address> on page 5528

Output Fields The following command displays the results of the IP address probe:

Sample Output

show services user-identification active-directory-access active-directory-authentication-table address <ip-address>

```
user@host> show services user-identification active-directory-access
active-directory-authentication-table address 192.0.2.3
Domain: ad02.net
 Source-ip: 192.0.2.3
  Username: lesjay
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2014-03-10
  Access start time: 13:59:56
  Age time: 1437
```


Copyright © 2016, Juniper Networks, Inc.

show network-access requests pending

Syntax
show network-access requests pending
<**detail**>
<**index** ***number*** >

Release Information Command introduced in Release 8.5 of Junos OS.

Description Display the status of pending authentication requests.

Options
- none—Show pending authentication requests.
- **detail**—Display detailed information about all pending requests.
- **index** ***number*** —(Optional) Display detailed information about the request specified by this index number. Use the command without options to obtain a list of requests and index numbers.

Required Privilege Level view

Related Documentation
- clear network-access requests pending on page 5513

List of Sample Output
show network-access requests pending on page 5530
show network-access requests pending detail on page 5530
show network-access requests pending index 1 on page 5530

Output Fields Table 474 lists the output fields for the show network-access requests pending command. Output fields are listed in the approximate order in which they appear.

Table 474: show network-access requests pending Output Fields

Field Name	Field Description
Index	Internal number identifying the pending request. Use this number to obtain more information on the record.
User	Originator of authentication request.
Status	The pending requests are requests and responses that are not yet sent back to the respective clients. The pending requests can be in one of the following states: • **Processing**: This request is being processed by the device. The authentication process has started but is not complete. • **Waiting on Auth Server**: The request is sent to an external authentication server, and the device is waiting for the response. • **Processed:** This request has completed authentication (success or failure). The results are not yet forwarded back to the client. • **Request cancelled by Admin**: This request was cancelled by the Admin. The reply with cancel code is not yet sent back to the client.

Copyright © 2016, Juniper Networks, Inc.

Table 474: show network-access requests pending Output Fields *(continued)*

Field Name	Field Description
Profile	The profile determines how the user is authenticated. Local clients defined with the statement **access profile client** are authenticated with the password authentication. Clients configured external to the device, on a RADIUS or LDAP server are authenticated with RADIUS or LDAP authentication.

Sample Output

show network-access requests pending

```
user@host> show network-access requests pending
Information about pending authentication entries
  Total pending authentication requests: 2
Index User                 Status
1     Sun                  Processing
2     Sam                  Processed
```

Sample Output

show network-access requests pending detail

```
user@host> show network-access requests pending detail
Information about pending authentication entries
  Total pending authentication requests: 2
Index: 1  User: Sun
  Status: Processing
  Profile: Sunnyvale-firewall-users
Index: 2  User: Sam
  Status: Processed
  Profile: Westford-profile
```

Sample Output

show network-access requests pending index 1

```
user@host> show network-access requests pending index 1
Index: 1  User: Sun
  Status: Processing
  Profile: Sunnyvale-firewall-users
```


Copyright © 2016, Juniper Networks, Inc.

show network-access requests statistics

Syntax **show network-access requests statistics**

Release Information Command modified in Release 9.1 of Junos OS.

Description Display authentication statistics for the configured authentication type.

Required Privilege Level view

Related Documentation
- authentication-order (Access Profile) on page 3040
- clear network-access requests statistics on page 5514

Output Fields Table 475 lists the output fields for the network-access requests statistics command. Output fields are listed in the approximate order in which they appear.

Table 475: show network-access requests statistics Output Fields

Field Name	Field Description
Total requests received	Total number of authentication requests that the device received from clients.
Total responses sent	Total number of authentication responses that the device sent to the clients.
Success responses	Total number of clients that authenticated successfully.
Failure responses	Total number of clients that failed to authenticate.

show network-access requests statistics

```
user@host> show network-access requests statistics
General authentication statistics
  Total requests received: 100
  Total responses sent: 70
Radius authentication statistics
  Total requests received: 40
  Success responses: 20
  Failure responses: 20
LDAP authentication statistics
  Total requests received: 30
  Success responses: 15
  Failure responses: 15
Local authentication statistics
  Total requests received: 5
  Success responses: 2
  Failure responses: 3
Securid authentication statistics
  Total requests received: 15
  Success responses: 3
  Failure responses: 12
```

Copyright © 2016, Juniper Networks, Inc.

show network-access securid-node-secret-file

Syntax **show network-access securid-node-secret-file**

Release Information Command introduced in Junos OS Release 9.1.

Description Display the path to the node secret file for the SecurID authentication type.

Required Privilege Level view

Related Documentation
- Firewall User Authentication Overview on page 5179

List of Sample Output show network-access securid-node-secret-file on page 5532

Output Fields Table 476 lists the output fields for the network-access securid-node-secret-file command. Output fields are listed in the approximate order in which they appear.

Table 476: show network-access securid-node-secret-file Output Fields

Field Name	Field Description
SecurID Server	Name of the SecurID authentication server.
Node Secret File	Path to the node secret file.

Sample Output

show network-access securid-node-secret-file

```
user@host> show network-access securid-node-secret-file
SecurID server node secret file:
SecurID Server       Node Secret File
ace-server1          /var/db/securid/ace-server1/node-secret
```


Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication history

Syntax **show security firewall-authentication history**
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display security firewall authentication history information.

Options
- none—Display history of firewall authentication information.
- **node**—(Optional) For chassis cluster configurations, display all firewall authentication history on a specific node (device) in the cluster.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation
- Understanding Logical System Firewall Authentication on page 3267
- Firewall User Authentication Overview on page 5179

List of Sample Output
show security firewall-authentication history on page 5534
show security firewall-authentication history node all on page 5534

Output Fields Table 358 lists the output fields for the **show security firewall-authentication history** command. Output fields are listed in the approximate order in which they appear.

Table 477: show security firewall-authentication history Output Fields

Field Name	Field Description
Authentications	Number of authentications.
Id	Identification number.
Source IP	IP address of the authentication source.
Date	Authentication date.
Time	Authentication time.
Duration	Authentication duration.
Status	Authentication status success or failure.

Copyright © 2016, Juniper Networks, Inc.

Table 477: show security firewall-authentication history Output Fields *(continued)*

Field Name	Field Description
User	Name of the user.

Sample Output

show security firewall-authentication history

```
user@host> show security firewall-authentication history
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           1 211.0.0.6       2007-04-03 11:43:06 00:00:45   Success  hello
```

Sample Output

show security firewall-authentication history node all

```
user@host> show security firewall-authentication history node all
node0:
--------------------------------------------------------------------------
History of firewall authentication data:
Authentications: 2
Id Source Ip       Date       Time     Duration   Status   User
1 100.0.0.1       2008-01-04 12:00:10 0:05:49    Success  local1
2 100.0.0.1       2008-01-04 14:36:52 0:01:03    Success  local1
node1:
--------------------------------------------------------------------------
History of firewall authentication data:
  Authentications: 1
          Id Source Ip       Date       Time     Duration   Status   User
           1 100.0.0.1       2008-01-04 14:59:43 1193046:06: Success local1
```


Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication history address

Syntax **show security firewall-authentication history address** *ip-address* <node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display security firewall authentication history for this source IP address.

Options

- **address** *ip-address* —IP address of the authentication source.
- none—Display all firewall authentication history for this address.
- **node**—(Optional) For chassis cluster configurations, display firewall authentication history for this address on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output

show security firewall-authentication history address 4.4.4.2 on page 5536
show security firewall-authentication history address 100.0.0.1 node local on page 5536

Output Fields Table 478 lists the output fields for the **show security firewall-authentication history address** command. Output fields are listed in the approximate order in which they appear.

Table 478: show security firewall-authentication history address Output Fields

Field Name	Field Description
Username	User ID.
Source IP	IP address of the authentication source.
Authentication state	Status of authentication (success or failure).
Authentication method	Path chosen for authentication.
Access start date	Date when user authenticated.
Access start time	Time when user authenticated.

Copyright © 2016, Juniper Networks, Inc.

Table 478: show security firewall-authentication history address Output Fields *(continued)*

Field Name	Field Description
Duration of user access	Time duration of the accessing firewall.
Policy name	Name of the policy.
Source zone	User traffic received from the zone.
Destination zone	User traffic destined to the zone.
Access profile	Name of profile used for authentication.
Bytes sent by this user	Number of bytes sent by the user.
Bytes received by this user	Number of bytes received by the user.

Sample Output

show security firewall-authentication history address 4.4.4.2

```
user@host> show security firewall-authentication history address 4.4.4.2
Username: u1
Source IP: 4.4.4.2
Authentication state: Success
Authentication method: Pass-through using HTTP
Access start date: 2007-09-12
Access start time: 15:33:29
Duration of user access: 0:00:48
Policy name: Z1-Z2
Source zone: Z1
Destination zone: Z2
Access profile: profile-local
Bytes sent by this user: 0
Bytes received by this user: 449
```

Sample Output

show security firewall-authentication history address 100.0.0.1 node local

```
user@host> show security firewall-authentication history address 100.0.0.1 node local
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 100.0.0.1
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2008-01-04
Access start time: 12:00:10
Duration of user access: 0:05:49
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
Bytes sent by this user: 0
Bytes received by this user: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Username: local1
Source IP: 100.0.0.1
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2008-01-04
Access start time: 14:36:52
Duration of user access: 0:01:03
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
Bytes sent by this user: 2178
Bytes received by this user: 4172
```

Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication history identifier

Syntax **show security firewall-authentication history identifier** *identifier*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display security firewall authentication history information for the authentication with this identifier.

Options

- **identifier** *identifier*—Identifying number of the authentication process.
- none—Display all firewall authentication history information for the authentication with this identifier.
- **node**—(Optional) For chassis cluster configurations, display firewall authentication history on a specific node for the authentication with this identifier.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
show security firewall-authentication history identifier 1 on page 5539
show security firewall-authentication identifier 1 node primary on page 5539

Output Fields Table 479 lists the output fields for the **show security firewall-authentication history identifier** command. Output fields are listed in the approximate order in which they appear.

Table 479: show security firewall-authentication history identifier Output Fields

Field Name	Field Description
Username	User ID.
Source IP	IP address of the authentication source.
Authentication state	Status of authentication (success or failure).
Authentication method	Path chosen for authentication.
Access start date	Date when user authenticated.

Copyright © 2016, Juniper Networks, Inc.

Table 479: show security firewall-authentication history identifier Output Fields *(continued)*

Field Name	Field Description
Access start time	Time when user authenticated.
Duration of user access	Time duration of the accessing firewall.
Policy index	Identification number of the policy.
Policy name	Name of the policy.
Source zone	User traffic received from the zone.
Destination zone	User traffic destined to the zone.
Access profile	Name of profile used for authentication.
Bytes sent by this user	Number of bytes sent by the user.
Bytes received by this user	Number of bytes received by the user.
Client-groups	Name of the client group.

Sample Output

show security firewall-authentication history identifier 1

```
user@host> show security firewall-authentication history identifier 1
Username: hello
Source IP: 211.0.0.6
Authentication state: Success
Authentication method: Pass-through using Telnet
Access start date: 2007-04-03
Access start time: 11:43:06
Duration of user access: 00:00:45
Policy index: 4
Source zone: z2
Destination zone: z1
Access profile: profile1
Bytes sent by this user: 0
Bytes received by this user: 1050
Client-groups: Sunnyvale Bangalore
```

Sample Output

show security firewall-authentication identifier 1 node primary

```
user@host> show security firewall-authentication history identifier 1 node primary
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 100.0.0.1
Authentication state: Success
Authentication method: Pass-through using Telnet
```

Copyright © 2016, Juniper Networks, Inc.

```
Access start date: 2008-01-04
Access start time: 12:00:10
Duration of user access: 0:05:49
Policy name: POL1
Source zone: z1
Destination zone: z2
Access profile: p1
Bytes sent by this user: 0
Bytes received by this user: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users

Syntax
show security firewall-authentication users
<node (*node-id* | all | local | primary)>

Release Information
Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description
Display firewall authentication details about all users.

Options
- none—Display details about all firewall authentication users.
- **node**—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
  - *node-id*—Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level
view

Related Documentation
- Firewall User Authentication Overview on page 5179

List of Sample Output
show security firewall-authentication users on page 5542
show security firewall-authentication users node 0 on page 5542
show security firewall-authentication users node all on page 5542

Output Fields
Table 359 lists the output fields for the **show security firewall-authentication users** command. Output fields are listed in the approximate order in which they appear.

Table 480: show security firewall-authentication users Output Fields

Field Name	Field Description
Total users in table	Gives count of how many entries/users the command will display.
Id	Identification number.
Source IP	IP address of the authentication source.
Src zone	User traffic received from the zone.
Dst zone	User traffic destined to the zone.
Profile	Name of profile used for authentication.
Age	Idle timeout for the user.

Copyright © 2016, Juniper Networks, Inc.

Table 480: show security firewall-authentication users Output Fields *(continued)*

Field Name	Field Description
Status	Authentication status success or failure.
User	Name of the user.

Sample Output

show security firewall-authentication users

```
user@host> show security firewall-authentication users
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           1 1111:1212/64        z1       z2       p1       0 Success  local1
```

Sample Output

show security firewall-authentication users node 0

```
user@host> show security firewall-authentication users node 0
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User

           3 100.0.0.1       z1       z2       p1           1 Success  local1
```

Sample Output

show security firewall-authentication users node all

```
user@host> show security firewall-authentication users node all
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 100.0.0.1       z1       z2       p1           1 Success  local1

node1:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           2 100.0.0.1       z1       z2       p1           1 Success  local1
```


Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users address

Syntax **show security firewall-authentication users address** *ip-address*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display information about the users at the specified IP address that are currently authenticated.

Options

- **address** *ip-address*—IP address of the authentication source.
- none—Display all the firewall authentication information for users at this IP address.
- **node**—(Optional) For chassis cluster configurations, display user firewall authentication entries on a specific node.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Understanding User Role Firewalls on page 1127

List of Sample Output

show security firewall-authentication users address 192.0.2.9 on page 5544
show security firewall-authentication users address 192.0.2.9 node local on page 5544
show security firewall-authentication users address 192.0.2.9 on page 5545

Output Fields Table 99 lists the output fields for the **show security firewall-authentication users address** command. Output fields are listed in the approximate order in which they appear.

Table 481: show security firewall-authentication users address Output Fields

Field Name	Field Description
Username	User ID.
Source IP	IP address of the authentication source.
Authentication state	Status of authentication (success or failure).
Authentication method	Path chosen for authentication.
Access time remaining	Duration for which the connection exists.

Copyright © 2016, Juniper Networks, Inc.

Table 481: show security firewall-authentication users address Output Fields *(continued)*

Field Name	Field Description
Lsys	The logical system where the traffic was received.
Source zone	User traffic received from the zone.
Destination zone	User traffic destined to the zone.
Policy index	Identification number of the policy.
Policy name	Name of the policy.
Access profile	Name of profile used for authentication.
Interface Name	Name of the interface.
Bytes sent by this user	Number of bytes sent by the user.
Bytes received by this user	Number of bytes received by the user.
Client-groups	Name of the client group.

Sample Output

show security firewall-authentication users address 192.0.2.9

```
user@host>show security firewall-authentication users address 192.0.2.9
Username: hello
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: Pass-through using Telnet
Access time remaining: 0
Source zone: z2
Destination zone: z1
Policy index: 5
Access profile: profile1
Interface Name: ge-0/0/2.0
Bytes sent by this user: 0
Bytes received by this user: 0
Client-groups: Sunnyvale Bangalore
```

Sample Output

show security firewall-authentication users address 192.0.2.9 node local

```
user@host> show security firewall-authentication users address 192.0.2.9 node local
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: Pass-through using Telnet
Age: 2
```


Copyright © 2016, Juniper Networks, Inc.

```
Access time remaining: 4
Source zone: z1
Destination zone: z2
Policy name: POL1
Access profile: p1
Interface Name: reth1.0
Bytes sent by this user: 614
Bytes received by this user: 1880
```

show security firewall-authentication users address 192.0.2.9

```
user@host> show security firewall-authentication users address 192.0.2.9
Username: hello
Source IP: 192.0.2.9
Authentication state: Success
Authentication method: User-firewall
Age: 0
Access time remaining: 10
Lsys: root-logical-system
Source zone: N/A
Destination zone: N/A
Access profile: test
```

Copyright © 2016, Juniper Networks, Inc.

show security firewall-authentication users identifier

Syntax **show security firewall-authentication users identifier** *identifier*
<node (*node-id* | all | local | primary)>

Release Information Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.

Description Display firewall authentication details about the user with this identification number.

Options

- **identifier** *identifier*—Identification number of the user for which to display authentication details.
- **node**—(Optional) For chassis cluster configurations, display the firewall authentication details security firewall authentication entry on a specific node (device) in the cluster for the user with this identification number.
  - ***node-id*** —Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.

Required Privilege Level view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output
show security firewall-authentication users identifier 3 on page 5547
show security firewall-authentication users identifier 3 node primary on page 5547

Output Fields Table 482 lists the output fields for the **show security firewall-authentication users identifier** command. Output fields are listed in the approximate order in which they appear.

Table 482: show security firewall-authentication users identifier Output Fields

Field Name	Field Description
Username	User ID.
Source IP	IP address of the authentication source.
Authentication state	Status of authentication (success or failure).
Authentication method	Path chosen for authentication.
Age	Idle timeout for the user.
Access time remaining	Duration for which the connection exists.

Copyright © 2016, Juniper Networks, Inc.

Table 482: show security firewall-authentication users identifier Output Fields *(continued)*

Field Name	Field Description
Source zone	User traffic received from the zone.
Destination Zone	User traffic destined to the zone.
Policy Name	Name of the policy.
Access profile	Name of profile used for authentication.
Interface Name	Name of the interface
Bytes sent by this user	Number of bytes sent by the user.
Bytes received by this user	Number of bytes received by the user.

Sample Output

show security firewall-authentication users identifier 3

```
user@host> show security firewall-authentication users identifier 3
Username: u1
Source IP: 4.4.4.2
Authentication state: Success
Authentication method: Pass-through using HTTP
Age: 1
Access time remaining: 254
Source zone: Z1
Destination zone: Z2
Policy name: Z1-Z2
Access profile: profile-local
Interface Name: ge-0/0/1.0
Bytes sent by this user: 0
Bytes received by this user: 449
```

Sample Output

show security firewall-authentication users identifier 3 node primary

```
user@host> show security firewall-authentication users identifier 3 node primary
node0:
--------------------------------------------------------------------------
Username: local1
Source IP: 100.0.0.1
Authentication state: Success
Authentication method: Pass-through using Telnet
Age: 1
Access time remaining: 5
Source zone: z1
Destination zone: z2
Policy name: POL1
Access profile: p1
Interface Name: reth1.0
Bytes sent by this user: 614
Bytes received by this user: 1880
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 5552
show security policies policy-name p1 detail on page 5553
show security policies (services-offload) on page 5554
show security policies detail on page 5554
show security policies detail (TCP Options) on page 5555
show security policies policy-name p1 (Negated Address) on page 5555
show security policies policy-name p1 detail (Negated Address) on page 5556
show security policies global on page 5556

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 483: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 483: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 483: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```


Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                  545 bps
      Initial direction:                9072                 272 bps
      Reply direction  :                9072                  272 bps
    Output bytes       :               18144                  545 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
      Initial direction:                9072                  272 bps
      Reply direction  :                9072                  272 bps
    Input  packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Output packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Session rate       :                 108                    3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```


Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control authentication-table

Syntax show services unified-access-control authentication-table

Release Information Command introduced in Junos OS Release 9.4. Options updated in Junos OS Release 12.1.

Description Display a summary of the authentication table entries configured from the IC Series UAC Appliance. Authentication tables store mappings between traffic sessions and Unified Access Control (UAC) roles. The IC Series appliance uses the roles specified in the mappings to help determine which UAC policies to apply to a session.

Use this command when you have configured the SRX Series device to act as a Junos OS Enforcer in a UAC deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series appliance.

You can also use this command to display the content of the authentication table in a user role firewall implementation. The table, pushed from a supporting UAC device, provides the user roles associated with incoming traffic.

Options

- **detail**—Display a detailed view of all authentication table entries.
- **extended**—Display a view of all authentication table entries with the user roles listed.
- **identifier** *id*—Display all authentication table entries with the specified identifier number.
- **ip** *source-ip-address*—Display any authentication table entry for the specified IP address.
- **role** *role-name*—Display all authentication table entries for the specified role name.
- **user** *username*—Display all authentication table entries for the specified user.

Required Privilege Level view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output

show services unified-access-control authentication-table on page 5557
show services unified-access-control authentication-table detail on page 5558
show services unified-access-control authentication-table extended on page 5558
show services unified-access-control authentication-table identifier id on page 5558
show services unified-access-control authentication-table ip on page 5558
show services unified-access-control authentication-table role on page 5558
show services unified-access-control authentication-table user username on page 5558

Sample Output

show services unified-access-control authentication-table

```
user@host>show services unified-access-control authentication-table
Id    Source IP       Username     Age      Role identifier
1     198.51.100.22    atsang       0        0000000001.000005.0
Total: 1
```

Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control authentication-table detail

```
user@host>show services unified-access-control authentication-table detail
Identifier: 1
  Source IP: 198.51.100.22
  Username: atsang
  Age: 0
  Role identifier       Role name
    0000000001.000005.0 Users
    1113249951.100616.0 PersonalFirewall
    1183670148.427197.0 UAC
Total: 1
```

show services unified-access-control authentication-table extended

```
user@host>show services unified-access-control authentication-table extended
Id    Source IP       Username     Age   Role name
3     10.214.161.195  prasanta     60    Users, PersonalFirewall
6     10.214.161.183  june         60    role-1
Total: 2
```

show services unified-access-control authentication-table identifier id

```
user@host>show services unified-access-control authentication-table identifier 1
Identifier: 1
  Source IP: 198.51.100.22
  Username: john
  Age: 0
  Role identifier       Role name
    0000000001.000005.0 Users
    1113249951.100616.0 PersonalFirewall
    1183670148.427197.0 UAC
Total: 1
```

show services unified-access-control authentication-table ip

```
user@host>show services unified-access-control authentication-table ip 10.214.161.183
Id    Source IP       Username     Age      Role identifier
8     10.214.161.183  june          0        1420298444.225667.0
Total: 1
```

show services unified-access-control authentication-table role

```
user@host>show services unified-access-control authentication-table role role-1
Id    Source IP       Username     Age      Role identifier
6     10.214.161.183  june         60       1420298444.225667.0
Total: 1
```

show services unified-access-control authentication-table user username

```
user@host>show services unified-access-control authentication-table user prasanta
Id    Source IP       Username     Age      Role identifier
7     10.214.161.195  prasanta     0        0000000001.000005.0
Total: 1
```


Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control counters

Syntax show services unified-access-control counters

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display the number of sessions allowed, denied, and terminated by the Unified Access Control (UAC) service when invoked by a firewall policy with the uac-policy action. Counts are reported for each action taken by UAC. Sessions that were allowed, denied, or terminated by other firewall policy actions are not included in these statistics.

On high-end SRX Series devices, UAC counts are grouped and displayed for each PIC on the device. On branch SRX Series devices, UAC counts are accumulated by device only. There is no PIC specification on these devices.

Required Privilege Level view

Related Documentation
- Firewall User Authentication Overview on page 5179

List of Sample Output show services unified-access-control counters on page 5560

Output Fields Table 484 lists the output fields for the **show services unified-access-control counters** command. Output fields are listed in the approximate order in which they appear.

Table 484: show services unified-access-control counters Output Fields

Field Name	Field Description
PIC	If applicable, the number of each PIC implementing UAC. UAC statistics are grouped by PIC.
Sessions allowed	The sessions permitted by UAC when invoked by a user role firewall policy.
Policy action	Number of sessions permitted by UAC based on the UAC policy action.
Timeout action	Number of sessions permitted by the timeout action while the SRX was disconnected from the UAC device.
Sessions denied	The sessions denied by UAC when invoked by a user role firewall policy.
Unauthenticated	Number of sessions denied by UAC because the user was not authenticated.
Policy action	Number of sessions denied by UAC based on the UAC policy action.
Policy not matched	Number of sessions denied because no UAC policy match was found.
Timeout action	Number of sessions denied by the timeout action while the SRX was disconnected from the access control device.
Sessions terminated	The sessions originally permitted that were later terminated.

Copyright © 2016, Juniper Networks, Inc.

Table 484: show services unified-access-control counters Output Fields *(continued)*

Field Name	Field Description
Reevaluation	Number of sessions terminated due to a change in the UAC user roles associated with the session.
Signout	Number of sessions terminated due to the user signing out.

Sample Output

show services unified-access-control counters

```
user@host> show services unified-access-control counters
PIC: fpc2.pic0
  Sessions allowed
    Policy action: 0
    Timeout action: 0
  Sessions denied
    Unauthenticated: 0
    Policy action: 0
    Policy not matched: 0
    Timeout action: 0
  Sessions terminated
    Reevaluation: 0
    Signout: 0
```

Statistics on branch devices are accumulated by device only. There is no PIC specification on these devices.

```
user@host> show services unified-access-control counters
Sessions allowed
    Policy action: 0
    Timeout action: 0
  Sessions denied
    Unauthenticated: 0
    Policy action: 0
    Policy not matched: 0
    Timeout action: 0
  Sessions terminated
    Reevaluation: 0
    Signout: 0
```


Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control policies

Syntax show services unified-access-control policies

Release Information Command introduced in Junos OS Release 9.4.

Description Display a summary of resource access policies configured from the IC Series UAC Appliance.

Use this command when you have configured the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series appliance.

Options

- **detail**—Display a detailed view of all policies.
- **identifier *id***—Display information about a specific policy by identification number.

Required Privilege Level view

Related Documentation

- Firewall User Authentication Overview on page 5179

List of Sample Output

show services unified-access-control policies on page 5561
show services unified-access-control policies detail on page 5561
show services unified-access-control policies identifier 1 on page 5562

Sample Output

show services unified-access-control policies

```
user@host> services unified-access-control policies
Id    Resource               Action Apply        Role identifier
1     10.100.15.0/24:*       allow  selected     1113249951.100616.0
2     10.100.17.0/24:*       deny   all
```

Sample Output

show services unified-access-control policies detail

```
user@host> services unified-access-control policies detail
Identifier: 1
  Resource: 10.100.15.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Action: allow
  Apply: selected
  Role identifier       Role name
    1113249951.100616.0 Personal Firewall
    1112927873.881659.0 Antivirus
    1183670148.427197.0 UAC
Identifier: 2
  Resource: 10.100.17.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Resource: 10.100.18.0/24:*
```

Copyright © 2016, Juniper Networks, Inc.

```
  Action: deny
  Apply: all
```

Sample Output

show services unified-access-control policies identifier 1

```
user@host> show services unified-access-control policies identifier 1
Identifier: 1
  Resource: 10.100.15.0/24:*
  Resource: 10.100.16.23-10.100.16.60:*
  Action: allow
  Apply: selected
  Role identifier       Role name
    1113249951.100616.0 Personal Firewall
    1112927873.881659.0 Antivirus
    1183670148.427197.0 UAC
```


Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control roles

Syntax show services unified-access-control roles

Release Information Command introduced in Junos OS Release 12.1.

Description When implementing user role firewall, display a summary of the roles that have been pushed to the SRX Series device from the access control service.

Required Privilege Level view

Related Documentation
- *Building Blocks Feature Guide for Security Devices*
- Firewall User Authentication Overview on page 5179

List of Sample Output show services unified-access-control roles on page 5563

Output Fields Table 485 lists the output fields for the **show services unified-access-control roles** command. Output fields are listed in the approximate order in which they appear.

Table 485: show services unified-access-control roles Output Fields

Field Name	Field Description
Name	Name of the user role.
Identifier	Unique identifier associated with the specified user role.
Total	Total number of user roles specified in the table.

Sample Output

show services unified-access-control roles

```
user@host> show services unified-access-control roles
Name                                          Identifier
Users                                         0000000001.000005.0
admin-1                                       1420298444.225667.0
Total: 2
```

Copyright © 2016, Juniper Networks, Inc.

show services unified-access-control status

Syntax show services unified-access-control status

Release Information Command introduced in Junos OS Release 9.4.

Description Display the status of the connection between the SRX Series device and the IC Series UAC Appliance as well as statistics to help debug connections to the IC Series appliance.

Use this command when you have configured the SRX Series device to act as a Junos OS Enforcer in a Unified Access Control (UAC) deployment. When deployed as a Junos OS Enforcer, the SRX Series device enforces the policies that are defined on the UAC's IC Series appliance.

Required Privilege Level view

Related Documentation
- Firewall User Authentication Overview on page 5179

List of Sample Output show services unified-access-control status on page 5564

Sample Output

show services unified-access-control status

```
user@host> show services unified-access-control status
Host           Address         Port   Interface     State
dev106vm26     10.64.11.106    11123  ge-0/0/0.0    connected
dev107vm26     10.64.11.106    11123  ge-0/0/0.0    closed
```


Copyright © 2016, Juniper Networks, Inc.

show services user-identification active-directory-access active-directory-authentication-table

Syntax

show services user-identification active-directory-access
active-directory-authentication-table
(all | group *name* | ip-address *ip-address* | user *name*
<domain *name*> <node (***node-id*** | all | local | primary)> <brief | extensive>

Release Information

Command introduced in Junos OS Release 12.1X47-D10.

Description

Display information about all entries in the Active Directory authentication table used in the integrated user firewall feature, or for a specific group, IP address or user.

Options

- **all**—Summary of the authentication entry information.
- **group *group-name***—Display the entries from the authentication table for the specified group.
- **ip-address *ip-address***—Display the entries from the authentication table for the specified IP address.
- **user *name***—Display the entries from the authentication table for the specified username.
- **domain *name***—(Optional) Display the summary, group, or user entries for the specified domain.
- **node**—(Optional) For chassis cluster configurations, display the summary, IP address, or user entries for a specific node.
  - ***node-id***—Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.
- **brief | extensive**—Display the specified level of output (the default is brief).

Required Privilege Level

view

Related Documentation

- clear services user-identification active-directory-access on page 5524
- show services user-identification active-directory-access domain-controller status on page 5569

List of Sample Output

show services user-identification active-directory-access active-directory-authentication-table ip-address <ip-address> on page 5566
show services user-identification active-directory-access active-directory-authentication-table all on page 5567
show services user-identification active-directory-access active-directory-authentication-table all domain on page 5567

Copyright © 2016, Juniper Networks, Inc.

show services user-identification active-directory-access active-directory-authentication-table all extensive on page 5567

Output Fields Table 486 lists the output fields for the **show services user-identification active-directory-access active-directory-authentication-table all extensive** command.

Table 486: show services user-identification active-directory-access active-directory-authentication-table all extensive Output Fields

Field Name	Field Description
Source IP	IP address for user who is logged in through the domain controller.
Username	ID of the user who is logged in through the domain controller.
Groups	Groups to which the user is associated in the domain controller.
State	States include the following: Pending—This IP address is being probed. Initial—The authentication entry is only received from the WMIC daemon, not pushed to the Packet Forwarding Engine. Valid—The authentication entry is pushed to the Packet Forwarding Engine. Invalid—The PC probe failed.
Access start date	Date that the authentication entry was created.
Access start time	Time that the authentication entry was created.
Age time	Number of minutes after which the authentication entry will time out.

Sample Output

show services user-identification active-directory-access active-directory-authentication-table ip-address <ip-address>

```
user@host> show services user-identification active-directory-access
active-directory-authentication-table ip-address 192.0.2.3
Domain: ad02.net
 Source-ip: 192.0.2.3
  Username: lesjay
  Groups:group1
  State: Valid
  Source: wmic
  Access start date: 2014-03-10
  Access start time: 13:59:56
  Age time: 1437
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show services user-identification active-directory-access active-directory-authentication-table all

```
user@host> show services user-identification active-directory-access
active-directory-authentication-table all
Domain: www.engineering-example.net
Total count: 2
Source IP       Username      Groups          State
 1.1.1.2         u2           r1, r3, r4      initial
 1.1.1.3         u3           r5, r6, r4      pending

Domain: www.hr-example.net
Total count: 2
Source IP       Username      Groups          State
 10.1.1.2         u2           r1, r3, r4      initial
 10.1.1.3         u3           r5, r6, r4      pending
```

Sample Output

show services user-identification active-directory-access active-directory-authentication-table all domain

```
user@host> show services user-identification active-directory-access
active-directory-authentication-table all domain www.engineering-example.net
Domain: www.engineering-example.net
Total count: 2
Source IP       Username      Groups          State
 1.1.1.2         u2           r1, r3, r4      initial
 1.1.1.3         u3           r5, r6, r4      pending
```

Sample Output

show services user-identification active-directory-access active-directory-authentication-table all extensive

```
user@host> show services user-identification active-directory-access
active-directory-authentication-table all extensive
Domain: www.engineering-example.net
Total entries: 2

Source IP: 1.1.1.2
Username: u2
Groups: r1, r3, r4
State: inital
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Source IP: 1.1.1.3
Username: u3
Groups: r5, r6, r4
State: pending
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Domain: www.hr-example.net
Total entries: 2

Source IP: 10.1.1.2
Username: u2
```

Copyright © 2016, Juniper Networks, Inc.

```
Groups: r1, r3, r4
State: inital
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20 min

Source IP: 10.1.1.3
Username: u3
Groups: r5, r6, r4
State: pending
Access start date: 2013-05-22
Access start time: 10:56:58
Age time: 20
```


Copyright © 2016, Juniper Networks, Inc.

show services user-identification active-directory-access domain-controller status

Syntax show services user-identification active-directory-access domain-controller status <domain *name*> <node (***node-id*** | all | local | primary)> <brief | extensive>

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Display status information for the Active Directory domain controllers configured for the integrated user firewall feature.

Options

- **domain *name***—(Optional) Display the status of the domain controllers for a specific domain.
- **node**—(Optional) For chassis cluster configurations, display the status of the domain controllers for a specific node.
  - ***node-id***—Identification number of the node. It can be 0 or 1.
  - **all**—Display information about all nodes.
  - **local**—Display information about the local node.
  - **primary**—Display information about the primary node.
- **brief | extensive**—Display the specified level of output (the default is brief).

Required Privilege Level view

Related Documentation

- active-directory-access on page 5400
- show services user-identification active-directory-access active-directory-authentication-table on page 5565

List of Sample Output

show services user-identification active-directory-access domain-controller status on page 5570
show services user-identification active-directory-access domain-controller status brief domain on page 5570
show services user-identification active-directory-access domain-controller status extensive domain on page 5570

Output Fields Table 487 lists the output fields for the **show services user-identification active-directory-access domain-controller status** command.

Table 487: show services user-identification active-directory-access domain-controller Output Fields

Field Name	Field Description
Domain controller	Domain controller name.
Address	IP address of the domain controller.

Table 487: show services user-identification active-directory-access domain-controller Output Fields *(continued)*

Field Name	Field Description
Status	Connection status of the domain controller: connected or disconnected.
Reason	Reason for a disconnected status: network issue, authentication failed, or host unreachable.

Sample Output

show services user-identification active-directory-access domain-controller status

Displays brief information for domain controllers in all configured domains.

```
user@host> show services user-identification active-directory-access domain-controller status
Domain: example-domain-controller.com
    Domain controller       Address         Status
    DC1                   203.0.113.51     Connected
    DC2                   203.0.113.12     Connected
    DC3                   203.0.113.6      Connected
    DC4                   203.0.113.11     Disconnected
    DC5                   203.0.113.7      Disconnected

Domain: example-domain
    Domain controller       Address          Status
    example-domain10          10.1.1.1       Disconnected
    example-domain20          10.2.2.2       Disconnected
    example-domain30          10.3.3.3       Disconnected
```

Sample Output

show services user-identification active-directory-access domain-controller status brief domain

```
user@host> show services user-identification active-directory-access domain-controller status
brief domain example-domain-controller.com
Domain: example-domain
  Domain controller: example-domain10
    Address: 10.1.1.1
    Status: Disconnected
    Reason: Network issue
  Domain controller: example-domain20
    Address: 10.2.2.2
    Status: Disconnected
    Reason: Authentication failed
  Domain controller: example-domain30
    Address: 10.3.3.3
    Status: Disconnected
    Reason: Host unreachable
```

Sample Output

show services user-identification active-directory-access domain-controller status extensive domain

```
user@host> show services user-identification active-directory-access domain-controller status
extensive domain Mars
```


Copyright © 2016, Juniper Networks, Inc.

```
Domain: Mars
  Domain controller: Mars10
    Address: 10.1.1.1
    Status: Disconnected
    Reason: Network issue
  Domain controller: Mars20
    Address: 10.2.2.2
    Status: Disconnected
    Reason: Authentication failed
  Domain controller: Mars30
    Address: 10.3.3.3
    Status: Disconnected
    Reason: Host unreachable
```

Copyright © 2016, Juniper Networks, Inc.

show services user-identification active-directory-access statistics

Syntax show services user-identification active-directory-access statistics (ip-user-mapping | ip-user-probe | user-group-mapping) <domain *name*>

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Display statistics about IP address-to-user mapping, user-to-group mapping, and IP user probes used for the integrated user firewall feature. If two domains are configured, output is provided per domain.

Options

- **ip-user-mapping**—Number of total queries and failed queries to the event log on the domain controller for address-to-user mappings. Includes additional information, such as the log scan interval and the timestamp of the last event read.
- **ip-user-probe**—Number of total PC probes and failed probes.
- **user-group-mapping**—Number of total queries and failed queries to the LDAP server for user-to-group mappings
- **domain *name***—(Optional) Display the statistics for the specified domain.

Required Privilege Level view

Related Documentation

- clear services user-identification active-directory-access on page 5524
- ip-user-mapping on page 5437
- request services user-identification active-directory-access ip-user-probe on page 5528
- user-group-mapping on page 5497

List of Sample Output

show services user-identification active-directory-access statistics ip-user-mapping on page 5573
show services user-identification active-directory-access statistics ip-user-probe on page 5574
show services user-identification active-directory-access statistics user-group-mapping on page 5574

Output Fields Table 488 lists the output fields for the **show services user-identification active-directory-access statistics ip-user-mapping** command.

Table 488: show services user-identification active-directory-access statistics ip-user-mapping Output Fields

Field Name	Field Description
Host	IP address of the domain controller.
Initial event log timespan	When the feature is first deployed, the number of previous hours for which the event log on the domain controller is read. A one means the last hour of the event log is read.

Copyright © 2016, Juniper Networks, Inc.

Table 488: show services user-identification active-directory-access statistics ip-user-mapping Output Fields *(continued)*

Field Name	Field Description
Eventlog scan interval	Number of seconds between event log scans.
Total log query number	Count of the queries on the event log.
Failed log query number	Count of the failed queries on the event log.
Log read number	Count of the times the event log was read.
Latest timestamp	Year:month:date:hours:minutes:seconds is the timestamp taken from the event log.

Table 489 lists the output fields for the **show services user-identification active-directory-access statistics ip-user-probe** command.

Table 489: show services user-identification active-directory-access statistics ip-user-probe Output Fields

Field Name	Field Description
Total user probe number	Count of the probes of IP addresses to get IP address-to-user mappings.
Failed user probe number	Count of failed probe attempts.

Table 490 lists the output fields for the **show services user-identification active-directory-access statistics user-group-mapping** command.

Table 490: show services user-identification active-directory-access statistics user-group-mapping Output Fields

Field Name	Field Description
Host	IP address and port being queried.
Total query number	Count of queries.
Failed query number	Count of failed query attempts.

Sample Output

show services user-identification active-directory-access statistics ip-user-mapping

```
user@host> show services user-identification active-directory-access statistics ip-user-mapping
Domain: ad03.net
     Host: 192.0.2.192
     Initial event log timespan : 1
     Eventlog scan interval : 60
     Total log query number : 240
     Failed log query number : 0
```

Copyright © 2016, Juniper Networks, Inc.

```
        Log read number : 838
        Latest timestamp :2013-10-11:15:11:54
        Host: 192.0.2.50
        Initial event log timespan : 1
        Eventlog scan interval : 60
        Total log query number : 273
        Failed log query number : 0
        Log read number : 2012
        Latest timestamp :2013-10-11:15:11:23
    Domain: acme.net
        Host: 192.0.2.39
        Initial event log timespan : 1
        Eventlog scan interval : 10
        Total log query number : 1596
        Failed log query number : 0
        Log read number : 6691
        Latest timestamp :2013-10-11:15:25:03
        Host: 192.0.2.1
        Initial event log timespan : 1
        Eventlog scan interval : 10
        Total log query number : 2628
        Failed log query number : 0
        Log read number : 114953
        Latest timestamp :2013-10-11:15:24:01
```

Sample Output

show services user-identification active-directory-access statistics ip-user-probe

```
user@host> show services user-identification active-directory-access statistics ip-user-probe
Domain: www.apac-acme.net
Total user probe number : 176116
Failed user probe number : 916
Domain: www.usa-acme.net
Total user probe number : 17632
Failed user probe number : 342
```

Sample Output

show services user-identification active-directory-access statistics user-group-mapping

```
user@host> show services user-identification active-directory-access statistics
user-group-mapping
Domain: www.apac-acme.net
Host: 192.0.2.1 Port  389
Total query number : 176116
Failed query number : 916
Domain: www.usa-acme.net
Host: 192.0.2.5 Port  389
Total query number : 8965
```


Copyright © 2016, Juniper Networks, Inc.

show services user-identification active-directory-access user-group-mapping

Syntax show services user-identification active-directory-access user-group-mapping (group *name* | status | user *name*) <domain *name*>

Release Information Command introduced in Junos OS Release 12.1X47-D10.

Description Display user-to-group mapping information used in the integrated user firewall feature. Note that the LDAP server is often the domain controller.

Options

- **group** ***group-name***—Display the users mapped to the specified group.
- **status**—Display the status of the last query to the LDAP server for user-group mapping.
- **user** ***name***—Display the groups for the specified username.
- **domain** ***name***—(Optional) Display the group, status, or user information for the specified domain.

Required Privilege Level view

Related Documentation

- LDAP Functionality in Integrated User Firewall on page 5317
- user-group-mapping on page 5497

List of Sample Output

show services user-identification active-directory-access user-group-mapping group <name> domain <name> on page 5576
show services user-identification active-directory-access user-group-mapping status on page 5576
show services user-identification active-directory-access user-group-mapping user <name> on page 5577

Output Fields Table 491 lists the output fields for the **show services user-identification active-directory-access user-group-mapping group** command.

Table 491: show services user-identification active-directory-access user-group-mapping group Output Fields

Field Name	Field Description
Domain	Domain of the specified group.
Users	Usernames mapped to the specified group.

Table 492 lists the output fields for the **show services user-identification active-directory-access user-group-mapping status** command.

Copyright © 2016, Juniper Networks, Inc.

Table 492: show services user-identification active-directory-access user-group-mapping status Output Fields

Field Name	Field Description
Domain	Domain for which the status is displayed.
LDAP server	IP address of the LDAP server.
Port	Port number on the LDAP server.
Last-query-status	Status of the last query from the SRX Series device.
Last-query-time	Year-month-date:hour:minutes:seconds when the SRX device last queried the LDAP server.

Table 493 lists the output fields for the **show services user-identification active-directory-access user-group-mapping user** command.

Table 493: show services user-identification active-directory-access user-group-mapping user Output Fields

Field Name	Field Description
Domain controller	Domain controller about which the user information is displayed.
Groups	Groups to which the user belongs.
Referenced by policy	Groups to which the user belongs and that are referenced by a firewall policy.

Sample Output

show services user-identification active-directory-access user-group-mapping group <name> domain <name>

```
user@host> show services user-identification active-directory-access user-group-mapping group
Finance domain www.apac-acme.net
show services user-identification active-directory-access user-group-mapping group
 Finance-group
Domain: www.apac-acme.net
Users: Davidlmu, Kevinabc

Domain: www.usa-acme.net
Users: Lilyqxz
```

Sample Output

show services user-identification active-directory-access user-group-mapping status

```
user@host> show services user-identification active-directory-access user-group-mapping status
Domain:tudor.test.com
LDAP server    Port      Last-query-status    Last-query-time
192.0.2.87      389       Query success        2014-02-07:15:50:52

Domain: second.net
```


Copyright © 2016, Juniper Networks, Inc.

```
LDAP server   Port      Last-query-status    Last-query-time
192.0.2.144    389       Idle                0
```

Sample Output

show services user-identification active-directory-access user-group-mapping user <name>

```
user@host> show services user-identification active-directory-access user-group-mapping user
Kwang
Domain www.apac-acme.net
Groups: Dev, NAT, SBU
Referenced by policy: SBU
Domain: www.usa-acme.net
Groups: HR, USA
```

Copyright © 2016, Juniper Networks, Inc.

UTM Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 71

Overview

- Understanding Unified Threat Management on page 5583
- Managing UTM Licensing on page 5587
- Configuring WELF Logging on page 5589
- Configuring UTM for Chassis Cluster on page 5593

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 230

Understanding Unified Threat Management

- Unified Threat Management Overview on page 5583
- Understanding UTM Custom Objects on page 5585

Unified Threat Management Overview

Unified Threat Management (UTM) is a term used to describe the consolidation of several security features into one device, protecting against multiple threat types. The advantage of UTM is streamlined installation and management of these multiple security capabilities.

The security features provided as part of the UTM solution are:

- Antispam Filtering—E-mail spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify e-mail spam. When the device detects an e-mail message deemed to be spam, it either drops the message or tags the message header or subject field with a preprogrammed string. The antispam feature uses a constantly updated spam block list (SBL). Sophos updates and maintains the IP-based SBL. The antispam feature is a separately licensed subscription service.
- Content Filtering—Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, protocol command, and embedded object type. Content filtering does not require a separate license.
- Web Filtering—Web filtering lets you manage Internet usage by preventing access to inappropriate Web content. There are three types of Web filtering solutions. In the case of the integrated Web filtering solution, the decision-making for blocking or permitting Web access is done on the device after it identifies the category for a URL either from user-defined categories or from a category server (Websense provides the CPA Server). The integrated Web filtering feature is a separately licensed subscription service. The redirect Web filtering solution intercepts HTTP requests and forwards the server URL to an external URL filtering server provided by Websense to determine whether to block or permit the requested Web access. Redirect Web filtering does not require a separate license. With Juniper Local Web Filtering, the decision-making for blocking or permitting Web access is done on the device after it identifies the category for a URL from user-defined categories stored on the device. With Local filtering, there is no additional Juniper license or remote category server required.

Copyright © 2016, Juniper Networks, Inc.

- Full File-Based Antivirus—A virus is executable code that infects or attaches itself to other executable code to reproduce itself. Some malicious viruses erase files or lock up systems. Other viruses merely infect files and overwhelm the target host or network with bogus data. The full file-based antivirus feature provides file-based scanning on specific Application Layer traffic checking for viruses against a virus signature database. It collects the received data packets until it has reconstructed the original application content, such as an e-mail file attachment, and then scans this content. Kaspersky Lab provides the internal scan engine. The full file-based antivirus scanning feature is a separately licensed subscription service.
- Express Antivirus—Express antivirus scanning is offered as a less CPU intensive alternative to the full file-based antivirus feature. The express antivirus feature, like the full antivirus feature, scans specific Application Layer traffic for viruses against a virus signature database. However, unlike full antivirus, express antivirus does not reconstruct the original application content. Rather, it just sends (streams) the received data packets, as is, to the scan engine. With express antivirus, the virus scanning is executed by a hardware pattern matching engine. This improves performance while scanning is occurring, but the level of security provided is lessened. Juniper Networks provides the scan engine. The express antivirus scanning feature is a separately licensed subscription service.
- Sophos Antivirus—Sophos antivirus scanning is offered as a less CPU-intensive alternative to the full file-based antivirus feature. Sophos supports the same protocols as full antivirus and functions in much the same manner; however, it has a smaller memory footprint and is compatible with lower end devices that have less memory. Sophos antivirus is as an in-the-cloud antivirus solution. The virus pattern and malware database is located on external servers maintained by Sophos (Sophos Extensible List) servers, thus there is no need to download and maintain large pattern databases on the Juniper device. The Sophos antivirus scanner also uses a local internal cache to maintain query responses from the external list server to improve lookup performance.

NOTE: **The sessions-per-client limit CLI command, which imposes a session throttle to prevent a malicious user from generating large amounts of traffic simultaneously, supports the antispam, content filtering, and antivirus UTM features. It does not support Web filtering.**

Related Documentation

- Understanding UTM Custom Objects on page 5585
- Understanding UTM Licensing on page 5587
- Updating UTM Licenses (CLI Procedure) on page 5588
- Understanding WELF Logging for UTM Features on page 5589
- Example: Configuring WELF Logging for UTM Features on page 5590

Copyright © 2016, Juniper Networks, Inc.

Understanding UTM Custom Objects

Before you can configure most UTM features, you must first configure the custom objects for the feature in question. Custom objects are global parameters for UTM features. This means that configured custom objects can be applied to all UTM policies where applicable, rather than only to individual policies.

The following UTM features make use of certain custom objects:

- Anti-Virus (see "Full Antivirus Pattern Update Configuration Overview" on page 5649)
- Web Filtering (see "Example: Configuring Integrated Web Filtering" on page 5789)
- Anti-Spam (see "Server-Based Antispam Filtering Configuration Overview" on page 5602)
- Content Filtering (see "Content Filtering Configuration Overview" on page 5748)

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding UTM Licensing on page 5587
- Updating UTM Licenses (CLI Procedure) on page 5588
- Understanding WELF Logging for UTM Features on page 5589
- Example: Configuring WELF Logging for UTM Features on page 5590

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 231

Managing UTM Licensing

- Understanding UTM Licensing on page 5587
- Updating UTM Licenses (CLI Procedure) on page 5588

Understanding UTM Licensing

The majority of UTM features function as a subscription service requiring a license. You can redeem this license once you have purchased your subscription license SKUs. You redeem your license by entering your authorization code and chassis serial number into the Customer Service License Management System (LMS) interface. Once your entitlement is generated, you can use the CLI from your device to send a license update request to the LMS server. The LMS server then sends your subscription license directly to the device.

NOTE: UTM requires 1 GB of memory.

Table 494: UTM Feature Subscription Service License Requirements

UTM Feature	Requires License
Antispam	Yes
Antivirus: full	Yes
Antivirus: express	Yes
Antivirus: sophos	Yes
Content Filtering	No
Web Filtering: integrated	Yes
Web Filtering: redirect	No
Web Filtering: local	No
Web Filtering: enhanced	Yes

Copyright © 2016, Juniper Networks, Inc.

NOTE: License enforcement is supported on all high-end SRX Series devices. Licensed features including anti-virus or Enhanced Web Filtering will not function until a license has been installed. The license must be installed after installing or upgrading to a new Junos OS Release version. Unlicensed features such as UTM blacklists and whitelists will continue to function without a license.

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding UTM Custom Objects on page 5585
- Updating UTM Licenses (CLI Procedure) on page 5588
- Understanding WELF Logging for UTM Features on page 5589
- Example: Configuring WELF Logging for UTM Features on page 5590

Updating UTM Licenses (CLI Procedure)

To apply your UTM subscription license to the device, use the following CLI command:

```
user@host> request system license update
```

After you install the license and reboot the device, the device reserves more memory for UTM features, and hence decreases the session capacity. Use the **set security forwarding-process application-services enable-utm-memory** command to manually reallocate the memory for UTM features. You must reboot the device for the configuration to take effect.

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding UTM Custom Objects on page 5585
- Understanding UTM Licensing on page 5587
- Understanding WELF Logging for UTM Features on page 5589
- Example: Configuring WELF Logging for UTM Features on page 5590

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 232

Configuring WELF Logging

- Understanding WELF Logging for UTM Features on page 5589
- Example: Configuring WELF Logging for UTM Features on page 5590

Understanding WELF Logging for UTM Features

UTM features support the WELF standard. The WELF Reference defines the WebTrends industry standard log file exchange format. Any system logging to this format is compatible with Firewall Suite 2.0 and later, Firewall Reporting Center 1.0 and later, and Security Reporting Center 2.0 and later.

A WELF log file is composed of records. Each record is a single line in the file. Records are always in chronological order. The earliest record is the first record in the file; the most recent record is the last record in the file. WELF places no restrictions on log filenames or log file rotation policies.

NOTE: Each WELF record is composed of fields. The record identifier field (id=) must be the first field in a record. All other fields can appear in any order.

The following is a sample WELF record:

```
id=firewall time="2000-2-4 12:01:01" fw=192.168.0.238 pri=6 rule=3 proto=http
src=192.168.0.23 dst=6.1.0.36 rg=www.example.com/index.html op=GET result=0
rcvd=1426
```

The fields from the example WELF record include the following required elements (all other fields are optional):

- **id** (Record identifier)
- **time** (Date/time)
- **fw** (Firewall IP address or name)
- **pri** (Priority of the record)

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding UTM Custom Objects on page 5585

Copyright © 2016, Juniper Networks, Inc.

- Understanding UTM Licensing on page 5587
- Updating UTM Licenses (CLI Procedure) on page 5588
- Example: Configuring WELF Logging for UTM Features on page 5590

Example: Configuring WELF Logging for UTM Features

This example shows how to configure WELF logging for UTM features.

- Requirements on page 5590
- Overview on page 5590
- Configuration on page 5590
- Verification on page 5591

Requirements

Before you begin, review the fields used to create a WELF log file and record. See "Understanding WELF Logging for UTM Features" on page 5589.

Overview

A WELF log file is composed of records. Each record is a single line in the file. Records are always in chronological order. The earliest record is the first record in the file; the most recent record is the last record in the file. WELF places no restrictions on log filenames or log file rotation policies. In this example, the severity level is emergency and the name of the security log stream is **utm-welf**.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security log source-address 1.2.3.4 stream utm-welf
set security log source-address 1.2.3.4 stream utm-welf format welf
set security log source-address 1.2.3.4 stream utm-welf format welf category
  content-security
set security log source-address 1.2.3.4 stream utm-welf format welf category
  content-security severity emergency
set security log source-address 1.2.3.4 stream utm-welf format welf category
  content-security severity emergency host 5.6.7.8
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure WELF logging for UTM features:

1. Set the security log source IP address.

   [edit security log]

 Copyright © 2016, Juniper Networks, Inc.

```
user@host# set source-address 1.2.3.4
```

NOTE: You must save the WELF logging messages to a dedicated WebTrends server.

2. Name the security log stream.

```
[edit security log]
user@host# set source-address 1.2.3.4 stream utm-welf
```

3. Set the format for the log messages.

```
[edit security log]
user@host# set source-address 1.2.3.4 stream utm-welf format welf
```

4. Set the category of log messages that are sent.

```
[edit security log]
user@host# set source-address 1.2.3.4 stream utm-welf format welf category
  content-security
```

5. Set the severity level of log messages that are sent.

```
[edit security log]
user@host# set source-address 1.2.3.4 stream utm-welf format welf category
  content-security severity emergency
```

6. Enter the host address of the dedicated WebTrends server to which the log messages are to be sent.

```
[edit security log]
user@host# set source-address 1.2.3.4 stream utm-welf format welf category
  content-security severity emergency host 5.6.7.8
```

Results From configuration mode, confirm your configuration by entering the **show security log** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security log
stream utm-welf {
   severity emergency;
   format welf;
   category content-security;
   host {
       5.6.7.8;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Security Log

Purpose Verify that the WELF log for UTM features is complete.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security utm status** command to verify if the UTM service is running or not.

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding UTM Custom Objects on page 5585
- Understanding UTM Licensing on page 5587
- Updating UTM Licenses (CLI Procedure) on page 5588

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 233

Configuring UTM for Chassis Cluster

- Understanding UTM Support for Active/Active Chassis Cluster on page 5593
- Understanding Chassis Cluster Support for UTM Modules on page 5595

Understanding UTM Support for Active/Active Chassis Cluster

A chassis cluster environment supports UTM with:

- Packet Forwarding Engine in active/backup chassis cluster configurations with the Packet Forwarding Engine and the Routing Engine being active in the same node (On SRX Series devices.
- On SRX550 devices, the UTM functionality is supported in both active/active and active/backup chassis cluster configurations where the Packet Forwarding Engine can be active on both the cluster nodes and the Routing Engine and the Packet Forwarding Engine can be active in different nodes.

NOTE: **UTM does not require a separate license for chassis cluster mode. The usual UTM license is sufficient and should be available on both of the nodes in the chassis cluster.**

UTM supports stateless (that is, no state regarding UTM is synchronized between the cluster nodes) the Packet Forwarding Engine active/active chassis cluster configurations. All the UTM sessions anchored on the redundancy group being failed over will be aborted and new sessions are set up with the new primary redundancy group.

Stateful active/active cluster mode is not supported. Stateful objects like UTM sessions will not be synchronized; that is, no UTM module runtime objects (RTOs) are synchronized between the cluster nodes. You need to install UTM licenses in both the nodes independently.

UTM is supported in the following chassis cluster modes:

- **Active/active mode**—In this mode, the redundancy groups can be active on both of the cluster nodes. The transit traffic can be processed by both nodes. Any traffic between nodes transits through the fabric link.

Copyright © 2016, Juniper Networks, Inc.

The transit traffic includes:

- Traffic forwarded between interfaces for redundancy groups 1 and up across nodes
- Traffic forwarded between interfaces for redundancy groups 1 and up that are part of the same node but have one or more redundancy groups active on both of the nodes
- Traffic forwarded between RG0-controlled interfaces across devices (traffic from the secondary RG0 is sent to the primary RG0 over the fabric link for routing decisions)

- **Active/backup mode**—In this mode, all the redundancy groups are active in one cluster node. All the transit traffic is processed by this single node.

  The transit traffic includes:

  - Traffic forwarded between interfaces for redundancy groups 1 and up that are part of the same node
  - Traffic forwarded between RG0-controlled interfaces for redundancy groups 1 and up that are in the same node

UTM is supported for the following chassis cluster failover types:

- **Manual failover**—Supports manual failover through the **set chassis cluster failover** command. Both RG0 and redundancy groups 1 and up can fail over using this command.
- **RG0 automatic failover**—This failover is supported through control link failure, monitoring objects (IP address, interface monitoring), or preempt/priority configuration.
- **Redundancy groups 1 and up automatic failover**—This failover is supported through monitoring objects (IP address, interface monitoring) or preempt/priority configuration. This failover leads to active link changes and can result in active/active mode.
- **Failover through reboot**—A primary node can be changed to a secondary node by rebooting the node. All redundancy groups in the node that's is rebooted will no longer be primary nodes.
- **Failover through flowd restart**—Redundancy groups 1 and up will be changed to secondary nodes when the flowd restarts.

The following UTM features are supported in chassis cluster:

- Content filtering
- URL (Web) filtering
- Antispam filtering
- Express antivirus scanning
- Full file-based antivirus scanning
- Sophos antivirus scanning

All the UTM configurations are either maintained in the Routing Engine or pushed to the Packet Forwarding Engine from the Routing Engine. The configuration synchronization

Copyright © 2016, Juniper Networks, Inc.

between the two nodes is taken care of by the chassis cluster infrastructure. This holds true for all the UTM modules too. You can configure UTM either from the primary or secondary node, and the same configuration will be reflected in the other node once you commit the first configuration.

There is a dependency on ACL support on control links. The time taken to spawn the processes depends on the device. There will be a small delay for the Unified Threat Management daemon (utmd) to come up operationally, even though utmd daemon is running in the secondary Routing Engine, because there can be a startup delay for all the dependant daemons.

Related Documentation

- *Chassis Cluster Overview*
- *Preparing Your Equipment for Chassis Cluster Formation*
- *Understanding Chassis Cluster Redundancy Groups*
- *Understanding Chassis Cluster Redundant Ethernet Interfaces*
- Unified Threat Management Overview on page 5583
- Understanding Chassis Cluster Support for UTM Modules on page 5595

Understanding Chassis Cluster Support for UTM Modules

- **Content filtering**—Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, and protocol command. The content filter controls file transfers across the gateway by checking traffic against configured filter lists.

  In content filtering, the user configuration (mime-pattern/filename-extension/protocol-command/content-type) is pushed from the Routing Engine to the Packet Forwarding Engine real-time (PFE-RT). The filtering decision is entirely based on the user configuration and is done on the Packet Forwarding Engine real-time (PFE-RT) side. For the transit traffic, the configuration lookup (for the block/permit decision) and the entire UTM processing occurs in the Packet Forwarding Engine itself and does not go to the Routing Engine (that is, the complete UTM session resides in the Packet Forwarding Engine).

- **URL (Web) filtering**—Web filtering lookups takes place in the primary Routing Engine and both the Packet Forwarding Engines send the lookup request to the primary Routing Engine.

- **Antispam filtering**—Antispam filtering pushes the user configuration (whitelist and blacklist) from the Routing Engine to the PFE-RT.

- **Express antivirus scanning**—In express antivirus scanning, the antivirus detection functionality is performed by the Pattern Matching Engine (PME) in the Packet Forwarding Engine of the node where the UTM traffic is anchored. The signature database is downloaded by the primary Routing Engine and synchronized to the secondary Routing Engine to be loaded in its local PME. If configured, the primary

Copyright © 2016, Juniper Networks, Inc.

Routing Engine does the periodic signature database updates and synchronizes them to the secondary Routing Engine.

- **Full file-based antivirus scanning**—In full antivirus scanning, the Kaspersky Lab engine is responsible for scanning all the data it receives. The signature database is downloaded from external Kaspersky Lab servers and used by the scan engine in the Routing Engine. Full AV antivirus scanning is done in the Routing Engine of each node where the UTM traffic is anchored. The signature database files are downloaded by the primary Routing Engine and synchronized to the secondary Routing Engine. If configured, the primary Routing Engine performs the periodic signature database update and synchronizes it to the secondary Routing Engine.

  Apart from the signature database lookup, full antivirus scanning uses the following configuration that is maintained in the RT side to determine if full antivirus scanning needs to be performed:

  - Mime whitelist—A hit bypasses antivirus scanning
  - Mime exception list—An exception to the whitelist
  - URL whitelist—a hit bypasses antivirus scanning
  - Filename extension—Only these extensions are sent for antivirus scanning

  The packet processing in full antivirus scanning might occur within RT side based on the result of the user-configured lists. Otherwise, the UTM session spans across the RT and Routing Engine side if full virus scanning needs to be performed.

Related Documentation

- *Chassis Cluster Overview*
- *Preparing Your Equipment for Chassis Cluster Formation*
- *Understanding Chassis Cluster Redundancy Groups*
- *Understanding Chassis Cluster Redundant Ethernet Interfaces*
- Unified Threat Management Overview on page 5583
- Understanding UTM Support for Active/Active Chassis Cluster on page 5593

Copyright © 2016, Juniper Networks, Inc.

PART 72

Configuring Antispam Filtering

- Understanding Antispam Filtering on page 5599
- Configuring Server-Based Antispam Filtering on page 5601
- Configuring Local List Antispam Filtering on page 5609

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 234

Understanding Antispam Filtering

- Antispam Filtering Overview on page 5599
- Handling Spam Messages on page 5599

Antispam Filtering Overview

Spam consists of unwanted e-mail messages, usually sent by commercial, malicious, or fraudulent entities. The antispam feature examines transmitted e-mail messages to identify spam. When the device detects a message deemed to be spam, it blocks the e-mail message or tags the e-mail message header or subject with a preprogrammed string.

Antispam filtering allows you to use both a third-party server-based spam block list (SBL) and to optionally create your own local whitelists (benign) and blacklists (malicious) for filtering against e-mail messages. The antispam feature is not meant to replace your antispam server, but to complement it.

Related Documentation

- Understanding Server-Based Antispam Filtering on page 5601
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Understanding Local List Antispam Filtering on page 5609
- Local List Antispam Filtering Configuration Overview on page 5610
- Handling Spam Messages on page 5599

Handling Spam Messages

There are two possible actions the device can take when spam is detected. It can perform a drop action or a tag action.

- Blocking Detected Spam on page 5599
- Tagging Detected Spam on page 5600

Blocking Detected Spam

The device can block and drop detected spam at either the connection level or the e-mail level:

Copyright © 2016, Juniper Networks, Inc.

- Blocking spam at the connection level

  When the SMTP sender is identified as a spam sender based on its IP address, the SMTP connection is rejected and dropped. An error message with a proper error code from the firewall is sent out on behalf of the SMTP server. An example of such an error message is:

  554 Transaction failed due to anti spam setting

- Blocking spam at the e-mail level

  When a particular e-mail sender is identified as spam sender based on its sender address, the e-mail is rejected and dropped. An error message with a proper error code from the firewall is sent back to the sender on behalf of the server. An example of such an error message is:

  550 Requested action not taken: mailbox unavailable

Tagging Detected Spam

The device can allow and tag the e-mail if the message sender is detected as a spammer. This tagging can occur at the connection level so that all the e-mails for the connection in question are tagged. Otherwise, you can tag only an individual e-mail. Two tagging methods are supported:

- Tag the subject: A user-defined string is added at the beginning of the subject of the e-mail.
- Tag the header: A user-defined string is added to the e-mail header.

Related Documentation

- Antispam Filtering Overview on page 5599
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Local List Antispam Filtering Configuration Overview on page 5610

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 235

Configuring Server-Based Antispam Filtering

- Understanding Server-Based Antispam Filtering on page 5601
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Example: Configuring Server-Based Antispam Filtering on page 5603

Understanding Server-Based Antispam Filtering

Server-based antispam filtering requires Internet connectivity with the spam block list (SBL) server. Domain Name Service (DNS) is required to access the SBL server. The firewall performs SBL lookups through the DNS protocol. The lookups are against the IP address of the sender (or relaying agent) of the e-mail, adding the name of the SBL server as the authoritative domain. The DNS server then forwards each request to the SBL server, which returns a DNS response to the device. The device then interprets the DNS response to determine if the e-mail sender is a spammer.

IP addresses that are included in the block lists are generally considered to be invalid addresses for mail servers or easily compromised addresses. Criteria for listing an IP address as a spammer on the SBL can include:

- Running an SMTP open relay service
- Running open proxy servers (of various kinds)
- Being a zombie host possibly compromised by a virus, worm, Trojan, or spyware
- Using a dynamic IP range
- Being a confirmed spam source with a known IP address

By default, the device first checks incoming e-mail against local whitelists and blacklists. If there are no local lists, or if the sender is not found on local lists, the device proceeds to query the SBL server over the Internet. When both server-based spam filtering and local list spam filtering are enabled, checks are done in the following order:

1. The local whitelist is checked. If there is a match, no further checking is done. If there is no match...
2. The local blacklist is checked. If there is a match, no further checking is done. If there is no match...

Copyright © 2016, Juniper Networks, Inc.

3. The SBL server list is checked.

NOTE:

- SBL server matching stops when the antispam license key is expired.
- Server-based spam filtering supports only IP-based spam block list blacklist lookup. Sophos updates and maintains the IP-based spam block list. Server-based antispam filtering is a separately licensed subscription service. When your antispam license key expires, you can continue to use locally defined blacklists and whitelists.

Related Documentation

- Antispam Filtering Overview on page 5599
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Example: Configuring Server-Based Antispam Filtering on page 5603
- Local List Antispam Filtering Configuration Overview on page 5610
- Understanding Local List Antispam Filtering on page 5609
- Handling Spam Messages on page 5599

Server-Based Antispam Filtering Configuration Overview

For each UTM feature, configure feature parameters in the following order:

1. Configure UTM custom objects for the feature:

   user@host# set security utm custom-objects

2. Configure the main feature parameters, using feature profiles.

   user@host# set security utm feature-profile anti-spam

3. Configure a UTM policy for each protocol, and attach this policy to a profile.

   user@host# set security utm utm-policy utmp1 anti-spam smtp-profile smtp1

NOTE: Antispam filtering is only supported for the SMTP protocol.

4. Attach the UTM policy to a security policy.

   user@host# set security policies from-zone trust to-zone untrust policy p1 then permit application-services utm-policy utmp1

Related Documentation

- Antispam Filtering Overview on page 5599
- Understanding Server-Based Antispam Filtering on page 5601
- Example: Configuring Server-Based Antispam Filtering on page 5603
- Understanding Local List Antispam Filtering on page 5609

Copyright © 2016, Juniper Networks, Inc.

- Local List Antispam Filtering Configuration Overview on page 5610
- Handling Spam Messages on page 5599

Example: Configuring Server-Based Antispam Filtering

This example shows how to configure server-based antispam filtering.

- Requirements on page 5603
- Overview on page 5603
- Configuration on page 5603
- Verification on page 5608

Requirements

Before you begin, review how to configure the feature parameters for each UTM feature. See "Server-Based Antispam Filtering Configuration Overview" on page 5602.

Overview

Server-based antispam filtering requires Internet connectivity with the spam block list (SBL) server. Domain Name Service (DNS) is required to access the SBL server.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile anti-spam sbl profile sblprofile1 sbl-default-server
set security utm feature-profile anti-spam sbl profile sblprofile1 sbl-default-server
    spam-action block
set security utm feature-profile anti-spam sbl profile sblprofile1 sbl-default-server
    custom-tag-string ***spam***
set security utm utm-policy spampolicy1 anti-spam smtp-profile sblprofile1
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy1 match
    source-address any
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy1 match
    destination-address any
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy1 match
    application junos-smtp
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy1 then permit
    application-services utm-policy spampolicy1
```

Copyright © 2016, Juniper Networks, Inc.

GUI Step-by-Step Procedure

To configure server-based antispam filtering:

1. Configure a profile and enable/disable the SBL server lookup. Select **Configure>Security>UTM>Anti-Spam**.
   a. In the Anti-Spam profiles configuration window, click **Add** to configure a profile for the SBL server, or click **Edit** to modify an existing item.
   b. In the Profile name box, enter a unique name for the antispam profile that you are creating.
   c. If you are using the default server, select **Yes** next to Default SBL server. If you are not using the default server, select **No**.

      NOTE: The SBL server is predefined on the device. The device comes preconfigured with the name and address of the SBL server. If you do not select Yes, you are disabling server-based spam filtering. You should disable it only if you are using only local lists or if you do not have a license for server-based spam filtering.

   d. In the Custom tag string box, enter a custom string for identifying a message as spam. By default, the devices uses ***SPAM***.
   e. From the antispam action list, select the action that the device should take when it detects spam. Options include Tag subject, Block email, and Tag header.
2. Configure a UTM policy for SMTP to which you attach the antispam profile.
   a. Select **Configure>Security>Policy>UTM Policies**.
   b. In the UTM policy configuration window, click **Add**.
   c. In the policy configuration window, select the **Main** tab.
   d. In the Policy name box, type a unique name for the UTM policy.
   e. In the Session per client limit box, type a session per client limit. Valid values range from 0 to 2000.
   f. From the Session per client over limit list, select the action that the device should take when the session per client limit for this UTM policy is exceeded. Options include Log and permit and Block.
   g. Select the **Anti-Spam profiles** tab in the pop-up window.
   h. From the SMTP profile list, select an antispam profile to attach to this UTM policy.

Copyright © 2016, Juniper Networks, Inc.

3. Attach the UTM policy to a security policy.
   a. Select **Configure>Security>Policy>FW Policies**.
   b. In the Security Policy window, click **Add** to configure a security policy with UTM or click **Edit** to modify an existing policy.
   c. In the Policy tab, type a name in the **Policy Name** box.
   d. Next to From Zone, select a zone from the list.
   e. Next to To Zone, select a zone from the list.
   f. Choose a source address.
   g. Choose a destination address.
   h. Choose an application by selecting **junos-smtp** (for antispam) in the Application Sets box and move it to the Matched box.
   i. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

   j. Select the **Application Services** tab.
   k. Next to UTM Policy, select the appropriate policy from the list. This attaches your UTM policy to the security policy.
   l. Click **OK** to check your configuration and save it as a candidate configuration.
   m. If the policy is saved successfully, you receive a confirmation, and you must click **OK** again. If the profile is not saved successfully, click **Details** in the pop-up window to discover why.

NOTE:

- You must activate your new policy to apply it.
- In SRX Series devices the confirmation window that notifies you that the policy is saved successfully disappears automatically.

   n. If you are done configuring the device, click **Commit Options>Commit**.

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure server-based antispam filtering:

1. Create a profile.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile sblprofile1
```

2. Enable or disable the default SBL server lookup.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile sblprofile1
  sbl-default-server
```

NOTE: If you are using server-based antispam filtering, you should type **sbl-default-server** to enable the default SBL server. (The SBL server is predefined on the device. The device comes preconfigured with the name and address of the SBL server.) You should disable server-based antispam filtering using the **no-sbl-default-server** option only if you are using only local lists or if you do not have a license for server-based spam filtering.

3. Configure the action to be taken by the device when spam is detected (block, tag-header, or tag-subject).

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile sblprofile1sbl-default-server
  spam-action block
```

4. Configure a custom string for identifying a message as spam.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile sblprofile1
  sbl-default-server custom-tag-string ***spam***
```

5. Attach the spam feature profile to the UTM policy.

```
[edit security]
user@host# set utm utm-policy spampolicy1 anti-spam smtp-profile sblprofile1
```

6. Configure a security policy for UTM to which to attach the UTM policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match source-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 match application junos-smtp
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy1 then permit application-services utm-policy spampolicy1
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: The device comes preconfigured with a default antispam policy. The policy is called junos-as-defaults. It contains the following configuration parameters:

```
anti-spam {
   sbl {
   profile junos-as-defaults {
      sbl-default-server;
      spam-action block;
   custom-tag-string "***SPAM***";
      }
   }
}
```

Results From configuration mode, confirm your configuration by entering the **show security utm** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm
feature-profile {
   anti-spam {
      sbl {
         profile sblprofile1 {
            sbl-default-server;
            spam-action block;
            custom-tag-string ***spam***;
         }
      }
   }
utm-policy spampolicy1 {
   anti-spam {
      smtp-profile sblprofile1;
   }
}
```

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
   policy utmsecuritypolicy1 {
      match {
         source-address any;
         destination-address any;
         application junos-smtp;
      }
      then {
         permit {
            application-services {
               utm-policy spampolicy1;
            }
         }
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Antispam Statistics

Purpose Verify the antispam statistics.

Action From operational mode, enter the **show security utm anti-spam status** and **show security utm anti-spam statistics** commands.

The following information appears:

```
SBL Whitelist Server:
SBL Blacklist Server:
msgsecurity.example.net
DNS Server:
Primary : 1.2.3.4, Src Interface: ge-0/0/0
Secondary: 2.3.4.5, Src Interface: ge-0/0/1
Ternary : 0.0.0.0, Src Interface: fe-0/0/2

Total connections: #
Denied connections: #
Total greetings: #
Denied greetings: #
Total e-mail scanned: #
White list hit: #
Black list hit: #
Spam total: #
Spam tagged: #
Spam dropped: #
DNS errors: #
Timeout errors: #
Return errors: #
Invalid parameter errors: #
Statistics start time:
Statistics for the last 10 days.
```

Related Documentation

- Antispam Filtering Overview on page 5599
- Understanding Server-Based Antispam Filtering on page 5601
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Understanding Local List Antispam Filtering on page 5609
- Local List Antispam Filtering Configuration Overview on page 5610
- Handling Spam Messages on page 5599
- **spam-action on page 5960**

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 236

Configuring Local List Antispam Filtering

- Understanding Local List Antispam Filtering on page 5609
- Local List Antispam Filtering Configuration Overview on page 5610
- Example: Configuring Local List Antispam Filtering on page 5610

Understanding Local List Antispam Filtering

When creating your own local whitelist and blacklist for antispam filtering, you can filter against domain names, e-mail addresses, and/or IP addresses. Pattern matching works a bit differently depending upon the type of matching in question. For example, pattern matching for domain names uses a longest suffix match algorithm. If the sender e-mail address has a domain name of aaa.bbb.ccc, the device tries to match "aaa.bbb.ccc" in the list. If no match is found, it tries to match "bbb.ccc", and then "ccc". IP address matching, however, does not allow for partial matches.

Antispam filtering uses local lists for matching in the following manner:

1. Sender IP: The sender IP is checked against the local whitelist, then the local blacklist, and then the SBL IP-based server (if enabled).
2. Sender Domain: The domain name is checked against the local whitelist and then against the local blacklist.
3. Sender E-mail Address: The sender e-mail address is checked against the local whitelist and then against the local blacklist.

By default, the device first checks incoming e-mail against the local whitelist and blacklist. If the sender is not found on either list, the device proceeds to query the SBL server over the Internet. When both server-based antispam filtering and local list antispam filtering are enabled, checks are done in the following order:

1. The local whitelist is checked. If there is a match, no further checking is done. If there is no match...
2. The local blacklist is checked. If there is a match, no further checking is done. If there is no match...
3. The SBL server list is checked.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Local blacklist and whitelist matching continues after the antispam license key is expired.

Related Documentation

- Antispam Filtering Overview on page 5599
- Local List Antispam Filtering Configuration Overview on page 5610
- Example: Configuring Local List Antispam Filtering on page 5610
- Server-Based Antispam Filtering Configuration Overview on page 5602
- Handling Spam Messages on page 5599

Local List Antispam Filtering Configuration Overview

For each UTM feature, configure feature parameters in the following order:

1. Configure UTM custom objects for the feature:

   user@host# set security utm custom-objects url-pattern *url-pattern-name*

2. Configure the main feature parameters, using feature profiles.

   user@host# set security utm feature-profile anti-spam *as-profile-name*

3. Configure a UTM policy for each protocol, and attach this policy to a profile.

   user@host# set security utm utm-policy utmp1 anti-spam smtp-profile smtp1

4. Attach the UTM policy to a security policy.

   user@host# set security policies from-zone trust to-zone untrust policy p1 then permit application-services utm-policy utmp1

Related Documentation

- Antispam Filtering Overview on page 5599
- Understanding Local List Antispam Filtering on page 5609
- Example: Configuring Local List Antispam Filtering on page 5610
- Understanding Server-Based Antispam Filtering on page 5601
- Handling Spam Messages on page 5599

Example: Configuring Local List Antispam Filtering

This example shows how to configure local list antispam filtering.

- Requirements on page 5611
- Overview on page 5611
- Configuration on page 5611
- Verification on page 5616

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, review how to configure the feature parameters for each UTM feature. See "Local List Antispam Filtering Configuration Overview" on page 5610.

Overview

Antispam filtering uses local lists for matching. When creating your own local whitelist and blacklist for antispam filtering, you can filter against domain names, e-mail addresses, and/or IP addresses.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects url-pattern as-black value [150.61.8.134]
set security utm custom-objects url-pattern as-white value [150.1.2.3]
set security utm feature-profile anti-spam address-whitelist as-white
set security utm feature-profile anti-spam sbl profile localprofile1
set security utm feature-profile anti-spam sbl profile localprofile1 spam-action block
set security utm feature-profile anti-spam sbl profile localprofile1 custom-tag-string
  ***spam***
set security utm utm-policy spampolicy2 anti-spam smtp-profile localprofile1
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy2 match
  source-address any
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy2 match
  destination-address any
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy2 match
  application junos-smtp
set security policies from-zone trust to-zone untrust policy utmsecuritypolicy2 then permit
  application-services utm-policy spampolicy2
```

GUI Step-by-Step Procedure

To configure local list antispam filtering:

1. Create local whitelist and blacklist custom objects by configuring a URL pattern list.
   a. Select **Configure>Security>UTM>Custom Objects**.
   b. In the UTM custom objects configuration window, select the **URL Pattern List** tab.
   c. Click **Add** to create URL pattern lists.
   d. Next to URL Pattern Name, type a unique name.

Copyright © 2016, Juniper Networks, Inc.

NOTE: If you are creating a whitelist, it is helpful to indicate this in the list name. The same applies to a blacklist. The name you enter here becomes available in the Address Whitelist and Address Blacklist fields when you are configuring your antispam profiles.

e. Next to URL Pattern Value, type the URL pattern for whitelist or blacklist antispam filtering.

2. Configure antispam filtering to use the whitelist and blacklist custom objects.

   a. Select **Configure>Security>UTM>Global options**.

   b. In the right pane, select the **Anti-Spam** tab.

   c. Under Anti-Spam, select an Address Whitelist and/or an Address Blacklist from the list for local lists for spam filtering. (These lists are configured as custom objects.)

   d. Click **OK**.

   e. If the configuration item is saved successfully, you receive a confirmation, and you must click **OK** again. If it is not saved successfully, click **Details** in the pop-up window to discover why.

   f. In the left pane under Security, select the **Anti-Spam** tab.

   g. Click **Add** to configure an anti-spam profile. The profile configuration pop-up window appears.

   h. In the Profile name box, enter a unique name.

   i. If you are using the default server, select **Yes** beside Default SBL server. If you are not using the default server, select **No**.

NOTE: If you select No, you are disabling server-based spam filtering. You disable it only if you are using local lists or if you do not have a license for server-based spam filtering.

Copyright © 2016, Juniper Networks, Inc.

j. In the Custom tag string box, type a custom string for identifying a message as spam. By default, the device uses ***SPAM***.

k. In the Actions list, select the action that the device should take when it detects spam. Options include Tag subject, Block email, and Tag header.

3. Configure a UTM policy for SMTP to which you attach the antispam profile.

a. Select **Configure>Security>Policy>UTM Policies**.

b. In the UTM policy configuration window, click **Add** to configure a UTM policy. The policy configuration pop-up window appears.

c. Select the **Main** tab.

d. In the Policy name box, type a unique name.

e. In the Session per client limit box, type a session per client limit. Valid values range from 0 through 2000.

f. From the Session per client over limit list, select the action that the device should take when the session per client limit for this UTM policy is exceeded. Options include Log and permit and Block.

g. Select the **Anti-Spam profiles** tab.

h. From the SMTP profile list, select the antispam profile that you are attaching to this UTM policy.

4. Attach the UTM policy to a security policy.

a. Select **Configure>Security>Policy>FW Policies**.

b. In the Security Policy window, click **Add** to configure a security policy with UTM. The policy configuration pop-up window appears.

c. In the Policy tab, type a name in the Policy Name box.

d. Next to From Zone, select a zone from the list.

e. Next to To Zone, select a zone from the list.

f. Choose a source address.

g. Choose a destination address.

h. Choose an application by selecting **junos-smtp** (for antispam) in the Application Sets box and move it to the Matched box.

i. Next to Policy Action, select one of the following: **Permit**, **Deny**, or **Reject**.

NOTE: When you select Permit for policy action, several additional fields become available in the Applications Services tab, including UTM Policy.

Copyright © 2016, Juniper Networks, Inc.

j. Select the **Application Services** tab.

k. Next to UTM Policy, select the appropriate policy from the list. This attaches your UTM policy to the security policy.

l. Click **OK** to check your configuration and save it as a candidate configuration.

m. If the policy is saved successfully, you receive a confirmation, and you must click **OK** again. If the profile is not saved successfully, click **Details** in the pop-up window to discover why.

NOTE: You must activate your new policy to apply it.

n. If you are done configuring the device, click **Commit Options>Commit**.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure local list antispam filtering:

1. Configure the local list spam blocking by first creating your global local spam lists.

```
[edit security]
user@host# set utm custom-objects url-pattern as-black value [150.61.8.134]
user@host# set utm custom-objects url-pattern as-white value [150.1.2.3]
```

2. Configure the local list antispam feature profile by first attaching your custom-object blacklist or whitelist or both.

```
[edit security]
user@host# set utm feature-profile anti-spam address-whitelist as-white
```

NOTE: When both the whitelist and the blacklist are in use, the whitelist is checked first. If there is no match, then the blacklist is checked.

3. Configure a profile for your local list spam blocking.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile localprofile1
```

NOTE: Although you are not using the SBL for local list spam blocking, you configure your profile from within that command similar to the server-based spam blocking procedure.

4. Configure the action to be taken by the device when spam is detected (block, tag-header, tag-subject).

Copyright © 2016, Juniper Networks, Inc.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile localprofile1 spam-action
  block
```

5. Configure a custom string for identifying a message as spam.

```
[edit security]
user@host# set utm feature-profile anti-spam sbl profile localprofile1
  custom-tag-string ***spam***
```

6. Attach the spam feature profile to the UTM policy.

```
[edit security]
user@host# set utm utm-policy spampolicy2 anti-spam smtp-profile localprofile1
```

7. Configure a security policy for UTM, and attach the UTM policy to the security policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match source-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match destination-address any
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 match application junos-smtp
user@host# set security policies from-zone trust to-zone untrust policy
  utmsecuritypolicy2 then permit application-services utm-policy spampolicy2
```

Results From configuration mode, confirm your configuration by entering the **show security utm** and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm
custom-objects {
  anti-spam {
    url-pattern patternwhite;
    address-whitelist as-white;
    sbl {
      profile localprofile1 {
        spam-action block;
        custom-tag-string ***spam***;
      }
    }
  }
utm-policy spampolicy2 {
  anti-spam {
    smtp-profile localprofile1;
  }
}
```

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy utmsecuritypolicy2 {
    match {
      source-address any;
      destination-address any;
      application junos-smtp;
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        then {
          permit {
            application-services {
              utm-policy spampolicy2;
            }
          }
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Antispam Statistics

Purpose Verify the antispam statistics.

Action From operational mode, enter the **show security utm anti-spam status** and **show security utm anti-spam statistics** commands.

The following information appears:

```
SBL Whitelist Server:
SBL Blacklist Server:
msgsecurity.example.net
DNS Server:
Primary : 1.2.3.4, Src Interface: ge-0/0/0
Secondary: 2.3.4.5, Src Interface: ge-0/0/1
Ternary : 0.0.0.0, Src Interface: fe-0/0/2

Total connections: #
Denied connections: #
Total greetings: #
Denied greetings: #
Total e-mail scanned: #
White list hit: #
Black list hit: #
Spam total: #
Spam tagged: #
Spam dropped: #
DNS errors: #
Timeout errors: #
Return errors: #
Invalid parameter errors: #
Statistics start time:
Statistics for the last 10 days.
```

Related Documentation

- Antispam Filtering Overview on page 5599
- Understanding Local List Antispam Filtering on page 5609
- Local List Antispam Filtering Configuration Overview on page 5610
- Handling Spam Messages on page 5599

Copyright © 2016, Juniper Networks, Inc.

- spam-action on page 5960

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 73

Configuring Express Antivirus Protection and Pattern Updates

- Configuring Express Antivirus Protection on page 5621
- Configuring Express Antivirus Pattern Updates on page 5641

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 237

Configuring Express Antivirus Protection

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Example: Configuring Express Antivirus Custom Objects on page 5624
- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Example: Configuring Express Antivirus Feature Profiles on page 5629
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634
- Example: Configuring Express Antivirus UTM Policies on page 5636
- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637
- Example: Attaching Express Antivirus UTM Policies to Security Policies on page 5638
- Attaching Express Antivirus UTM Policies to Security Policies (J-Web Procedure) on page 5639

Express Antivirus Protection Overview

Express antivirus scanning is offered as a less CPU intensive alternative to the full file-based antivirus feature. Express antivirus supports the same protocols as full antivirus and functions in much the same manner, however, it has a smaller memory footprint, compatible with the smaller system memory present on lower end devices.

This topic includes the following sections:

- Express Antivirus Packet-Based Scanning Versus File-Based Scanning on page 5621
- Express Antivirus Expanded MIME Decoding Support on page 5622
- Express Antivirus Scan Result Handling on page 5622
- Express Antivirus Intelligent Prescreening on page 5622
- Express Antivirus Limitations on page 5622

Express Antivirus Packet-Based Scanning Versus File-Based Scanning

Express antivirus uses a different antivirus scan engine than the full file-based antivirus feature and a different back-end hardware engine to accelerate pattern matching for higher data throughput.

Copyright © 2016, Juniper Networks, Inc.

The packet-based scanning done by express antivirus provides virus scanning data buffers without waiting for entire file to be received by the firewall, whereas the file-based scanning done by full antivirus can only start virus scanning when entire file is received.

Express Antivirus Expanded MIME Decoding Support

Express antivirus offers MIME decoding support for HTTP, POP3, SMTP, and IMAP. MIME decoding support includes the following for each supported protocol:

- Multi-part and nested header decoding
- Base64 decoding, printed quote decoding, and encoded word decoding (in the subject field)

Express Antivirus Scan Result Handling

With express antivirus, the TCP traffic is closed gracefully when a virus is found and the data content is dropped.

NOTE: **Express antivirus supports the following fail mode options: default, engine-not-ready, out-of-resource, and too-many-requests. Fail mode handling of supported options with express antivirus is much the same as with full antivirus.**

Express Antivirus Intelligent Prescreening

Intelligent prescreening functionality is identical in both express antivirus and full antivirus.

Express Antivirus Limitations

Express antivirus has the following limitations when compared to full antivirus functionality:

- Express antivirus provides limited support for the scanning of file archives and compressed file formats. Express antivirus can only support gzip, deflate and compressed compressing formats.
- Express antivirus provides limited support for decompression. Decompression is only supported with HTTP (supports only gzip, deflate, and compress for HTTP and only supports one layer of compression) and POP3 (supports only gzip for POP3 and only supports one layer of compression).
- Express antivirus does not support scanning by extension.
- Express antivirus scanning is interrupted when the scanning database is loading.
- Express antivirus may truncate a warning message if a virus has been detected and the replacement warning message that is sent is longer than the original content it is replacing.
- If you switch from express antivirus protection to full file-based antivirus protection, you must reboot the device in order for full file-based antivirus to begin working.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Because express antivirus does only packet-based string matching, if you use the standard EICAR file to test express antivirus, you will see false positives. To avoid these false positives, Juniper Networks has disabled scanning on the standard EICAR file to create a modified EICAR file for testing express antivirus. You can download this modified EICAR file from the following links:

http://www.juniper.net/security/avtest/ss-eicar.txt

http://www.juniper.net/security/avtest/ss-eicar.com

http://www.juniper.net/security/avtest/ss-eicar.zip

NOTE: The modified EICAR file must be tested with express antivirus only. The Kaspersky antivirus and Sophos antivirus do not detect this file.

NOTE: The express antivirus feature provides better performance but lower security. Note that if you switch from full file-based antivirus protection to express antivirus protection, you must reboot the device in order for express antivirus to begin working.

Related Documentation

- Understanding Express Antivirus Scanner Pattern Updates on page 5641
- Express Antivirus Configuration Overview on page 5623
- Example: Automatically Updating Express Antivirus Patterns on page 5642

Express Antivirus Configuration Overview

For each UTM feature, you should configure feature parameters in the following order:

1. Configure UTM custom objects for the UTM features. The following example enables the mime-pattern, url-pattern, and custom-url-category custom objects:

   user@host# set security utm custom-objects mime-pattern
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category

2. Configure main feature parameters using feature profiles. The following examples enables the anti-virus feature profile:

   user@host# set security utm feature-profile anti-virus juniper-exress-engine

3. Configure a UTM policy for each protocol and attach this policy to a profile. The following example creates the utmp3 UTM policy for the HTTP protocol:

   user@host# set security utm utm-policy utmp3 anti-virus http-profile http1

Copyright © 2016, Juniper Networks, Inc.

4. Attach the UTM policy to a security policy. The following example attaches the utmp3 UTM policy to the p3 security policy:

   user@host# set security policies from-zone trust to-zone untrust policy p3 then permit application-services utm-policy utmp3

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634
- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637

Example: Configuring Express Antivirus Custom Objects

This example shows how to configure express antivirus custom objects.

- Requirements on page 5624
- Overview on page 5624
- Configuration on page 5625
- Verification on page 5626

Requirements

Before you begin:

- Decide the type of express antivirus protection you require. See "Express Antivirus Protection Overview" on page 5621.
- Understand the order in which express antivirus parameters are configured. See "Express Antivirus Configuration Overview" on page 5623.

Overview

In this example, you define custom objects that are used to create express antivirus feature profiles. You perform the following tasks to define custom objects:

- Create two MIME lists called avmime2 and ex-avmime2, and add patterns to the list.
- Configure a URL pattern list called urllist2.

  When entering the URL pattern, note the following wildcard character support:

  - The *\.[]\?* wildcard characters are supported.
  - You must precede all wildcard URLs with http://.
  - You can use the asterisk * wildcard character only if it is at the beginning of the URL and is followed by a period.
  - You can use the question mark ? wildcard character only at the end of the URL.

Copyright © 2016, Juniper Networks, Inc.

  - The following wildcard syntax is supported: http://*example.net, http://www.example.ne?, http://www.example.n??.
  - The following wildcard syntax is not supported: *example.net , www.example.ne?, http://*example.net, http://* .
- Configure a custom URL category list called custurl2, using the urllist2 URL pattern list.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects mime-pattern avmime2 value [video/quicktime
  image/x-portable-anymap x-world/x-vrml]
set security utm custom-objects mime-pattern ex-avmime2 value
  [video/quicktime-inappropriate]
set security utm custom-objects url-pattern urllist2 value [http://www.example.net
  1.2.3.4]
set security utm custom-objects custom-url-category custurl2 value urllist2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure express antivirus filtering custom objects:

1. Create MIME lists, and add MIME patterns to the lists.

   ```
   [edit security utm]
   user@host# set custom-objects mime-pattern avmime2 value [video/quicktime
     image/x-portable-anymap x-world/x-vrml]
   user@host# set custom-objects mime-pattern ex-avmime2 value
     [video/quicktime-inappropriate]
   ```

   NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category list.

2. Configure a URL pattern list custom object.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllist2 value [http://www.example.net
     1.2.3.4]
   ```

3. Configure a custom URL category list.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custurl2 value urllist2
   ```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm
custom-objects {
  mime-pattern {
    avmime2 {
      value [ video/quicktime image/x-portable-anymap x-world/x-vrml ];
    }
    ex-avmime2 {
      value video/quicktime-inappropriate;
    }
  }
  url-pattern {
    urllist2 {
      value [ http://www.example.net 1.2.3.4 ];
    }
  }
  custom-url-category {
    custurl2 {
      value urllist2;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Express Antivirus Custom Objects

Purpose Verify the express antivirus custom objects.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Attaching Express Antivirus UTM Policies to Security Policies (J-Web Procedure) on page 5639
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634

Configuring Express Antivirus Custom Objects (J-Web Procedure)

To configure express antivirus protection using the J-Web configuration editor, you must first create your custom objects (MIME pattern list, URL pattern list, and custom URL category list).

Copyright © 2016, Juniper Networks, Inc.

Configure a MIME pattern list custom object as follows:

1. Select **Configure>Security>UTM Custom Objects**.
2. From the MIME Pattern List tab, click **Add** to create MIME pattern lists.
3. In the Add MIME Pattern pop-up window, next to **MIME Pattern Name**, enter a unique name.

   NOTE: Keep in mind that you are creating a MIME whitelist and a MIME exception list (if necessary). Both MIME lists appear in the MIME Whitelist and Exception MIME Whitelist fields when you configure antivirus. Therefore, the MIME list names you create should be as descriptive as possible.

4. Next to MIME Pattern Value, enter the MIME pattern.
5. Click **Add** to add your MIME pattern to the Values list box. Within this box, you can also select an entry and use the Delete button to delete it from the list. Continue to add MIME patterns in this manner.
6. Optionally, create a new MIME list to act as an exception list. The exception list is generally a subset of the main MIME list.
7. Click **OK** to check your configuration and save the selected values as part of the MIME list, then click **Commit Options**>**Commit**.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a URL pattern list custom object as follows:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure>Security>UTM>Custom Objects**.
2. From the URL Pattern List tab, click **Add** to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to URL Pattern Value, enter the URL or IP address you want added to list for bypassing scanning.

   When entering the URL pattern, note the following wildcard character support:

   - The ***\.[]\?*** wildcard characters are supported.
   - You must precede all wildcard URLs with **http://**.

Copyright © 2016, Juniper Networks, Inc.

- You can only use the asterisk * wildcard character if it is at the beginning of the URL and is followed by a period.
- You can only use the question mark ? wildcard character at the end of the URL.
- The following wildcard syntax IS supported: **http://*.example.net**, **http://www.example.ne?**, **http://www.example.n??**.
- The following wildcard syntax is NOT supported: ***.example.net** , **www.example.ne?**, **http://*example.net, http://***.

5. Click **Add** to add your URL pattern to the Values list box. The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.
6. Click **OK** to check your configuration and save the selected values as part of the URL pattern list, then click **Commit Options**>**Commit**.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a custom URL category list custom object using the URL pattern list that you created:

1. Select **Configure>Security>UTM>Custom Objects**.
2. From the URL Category List tab, click **Add** to create URL category lists.
3. Next to URL Category Name, enter a unique name. This name appears in the URL Whitelist list when you configure antivirus global options.
4. In the Available Values box, select a **URL Pattern List** name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.
5. Click **OK** to check your configuration and save the selected values as part of the URL list, then click **Commit Options**>**Commit**.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634
- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Express Antivirus Feature Profiles

This example shows how to configure an express antivirus feature profile.

- Requirements on page 5629
- Overview on page 5629
- Configuration on page 5630
- Verification on page 5633

Requirements

Before you begin:

- Decide the type of express antivirus protection you require. See "Express Antivirus Protection Overview" on page 5621.
- Understand the order in which express antivirus parameters are configured. See "Express Antivirus Configuration Overview" on page 5623.
- MIME patterns must be defined for lists and exception lists. See "Example: Configuring MIME Whitelists to Bypass Antivirus Scanning" on page 5706.
- Custom objects must be defined. See "Example: Configuring Express Antivirus Custom Objects" on page 5624
- SMTP must be configured on the device. See "Understanding SMTP Antivirus Scanning" on page 5698

Overview

In this example, you configure a feature profile called junexprof1 and specify custom objects to be used for filtering content.

- Select and configure the Juniper Express Engine as the engine type.
- Select 120 as the time interval for updating the pattern database. The default antivirus pattern-update interval is once a day.

NOTE: The command for changing the URL for the pattern database is:

```
[edit]
user@host# set security utm feature-profile anti-virus juniper-express-engine
   pattern-update url http://...
```

Under most circumstances, you should not need to change the default URL.

- Enable an e-mail notification with a custom message as pattern file was updated and a custom subject line as AV pattern file updated.
- Configure a list of fallback options as block.

Copyright © 2016, Juniper Networks, Inc.

- Configure the notification options for fallback blocking for virus detection. Configure a custom message for the fallback blocking action, and send a notification.
- Configure a notification for protocol-only virus detection, and send a notification as Antivirus Alert.
- Configure content size parameters as 20000.
- Enable intelligent prescreening and set its timeout setting to 1800 seconds and trickling setting (applicable only to HTTP) to 600 seconds. This means that if the device receives a packet within a 600-second period during a file transfer or while performing an antivirus scan, it should not time out.

NOTE: Intelligent prescreening is intended only for use with non-encoded traffic. It is not applicable to mail protocols (SMTP, POP3, IMAP) or HTTP POST.

- Configure the antivirus scanner to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list, called junos-default-bypass-mime, which ships with the device. The following example enables the avmime2 and ex-avmime2 lists.
- Configure the antivirus module to use URL bypass lists. If you are using a URL whitelist (valid only for HTTP traffic), this is a custom URL category that you previously configured as a custom object. For this example, you enable the custurl1 bypass list.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile anti-virus juniper-express-engine pattern-update interval
    120
set security utm feature-profile anti-virus juniper-express-engine pattern-update
    email-notify admin-email administrator@example.net custom-message “pattern file
    was updated” custom-message-subject “AV pattern file updated”
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options content-size block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options default block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options engine-not-ready block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options out-of-resources block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options timeout block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    fallback-options too-many-requests block
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
    notification-options fallback-block custom-message “Dropped due to fallback condition”
```


Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
  notification-options virus-detection type protocol-only
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
  notification-options virus-detection custom-message ***virus-found***
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
  scan-options content-size-limit 20000
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
  scan-options intelligent-prescreening
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1
  scan-options timeout 1800
set security utm feature-profile anti-virus juniper-express-engine profile junexprof1 trickling
  timeout 600
set security utm feature-profile anti-virus mime-whitelist list avmime2
set security utm feature-profile anti-virus mime-whitelist list avmime2 exception
  ex-avmime2
set security utm feature-profile anti-virus url-whitelist custurl2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure express antivirus feature profiles:

1. Select and configure the engine type.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus type juniper-express-engine
   ```

2. Select a time interval for updating the pattern database.

   ```
   [edit security utm feature-profile anti-virus juniper-express-engine]
   user@host# set pattern-update interval 120
   ```

3. Configure the device to notify a specified administrator when patterns are updated.

   ```
   [edit security utm feature-profile anti-virus juniper-express-engine]
   user@host# set pattern-update email-notify admin-email
     administrator@example.net custom-message “pattern file was updated”
     custom-message-subject “AV pattern file updated”
   ```

4. Create a profile for the Juniper Express Engine, and configure fallback options as block.

   ```
   [edit security utm feature-profile anti-virus juniper-express-engine]
   user@host# set profile junexprof1 fallback-options content-size block
   user@host# set profile junexprof1 fallback-options default block
   user@host# set profile junexprof1 fallback-options engine-not-ready block
   user@host# set profile junexprof1 fallback-options out-of-resources block
   user@host# set profile junexprof1 fallback-options timeout block
   user@host# set profile junexprof1 fallback-options too-many-requests block
   ```

5. Configure a custom notification for the fallback blocking action, and send a notification.

   ```
   [edit security utm feature-profile anti-virus juniper-express-engine]
   user@host# set profile junexprof1 notification-options fallback-block
     custom-message “Dropped due to fallback condition”
   ```

6. Configure a notification for protocol-only virus detection, and send a notification.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
user@host# set profile junexprof1 notification-options virus-detection type
  protocol-only
```

7. Configure a custom notification for virus detection.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
set profile junexprof1 notification-options virus-detection custom-message
  ***virus-found***
```

8. Configure content size parameter.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
user@host# set profile junexprof1 scan-options content-size-limit 20000
```

9. Configure intelligent prescreening.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
user@host# set profile junexprof1 scan-options intelligent-prescreening
```

10. Configure the timeout setting.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
user@host# set profile junexprof1 scan-options timeout 1800
```

11. Configure trickling setting.

```
[edit security utm feature-profile anti-virus juniper-express-engine]
user@host# set profile junexprof1 trickling timeout 600
```

12. Configure the antivirus scanner to use MIME bypass lists and exception lists.

```
[edit security utm feature-profile anti-virus]
user@host# set mime-whitelist list avmime2
user@host# set mime-whitelist list avmime2 exception ex-avmime2
```

13. Configure the antivirus module to use URL bypass lists.

```
[edit security utm feature-profile anti-virus]
user@host# set url-whitelist custurl2
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile anti-virus** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm feature-profile anti-virus
mime-whitelist {
   list avmime2;
   exception ex-avmime2;
}
url-whitelist custurl2;
juniper-express-engine {
   pattern-update {
      email-notify {
         admin-email "administrator@example.net";
         custom-message "pattern file was updated";
         custom-message-subject "AV pattern file updated";
      }
      interval 120;
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
    profile junexprof1 {
      fallback-options {
        default block;
        content-size block;
        engine-not-ready block;
        timeout block;
        out-of-resources block;
        too-many-requests block;
      }
      scan-options {
        intelligent-prescreening;
        content-size-limit 20000;
        timeout 1800;
      }
      trickling timeout 600;
      notification-options {
        virus-detection {
          type protocol-only;
          custom-message ***virus-found***;
        }
        fallback-block {
          custom-message “Dropped due to fallback condition”;
        }
      }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration of Express Antivirus Feature Profile

Purpose Verify the express antivirus feature profile.

Action From operational mode, enter any of the following commands:

- **show configuration security utm**
- **show security utm anti-virus status**
- **show security utm anti-virus statistics**

Related Documentation

- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637
- Example: Configuring Express Antivirus UTM Policies on page 5636
- Example: Attaching Express Antivirus UTM Policies to Security Policies on page 5638

Copyright © 2016, Juniper Networks, Inc.

Configuring Express Antivirus Feature Profiles (J-Web Procedure)

After you create your custom objects, configure the antivirus feature profile:

1. Select Configure>Security>UTM>Global options.
2. In the Anti-Virus tab, next to MIME whitelist, select the custom object you created from the list.
3. Next to Exception MIME whitelist, select the custom object you created from the list.
4. Next to URL Whitelist, select the custom object you created from the list.
5. In the Engine Type section, select the type of engine you are using. For express antivirus protection, you should select Juniper Express.
6. Next to Pattern update URL, enter the URL for the pattern database in the box. Note that the URL is http://update.juniper-updates.net/EAV/<device version> and you should not change it.
7. Next to Pattern update interval, enter the time interval for automatically updating the pattern database in the box. The default for express antivirus checking is once per day.
8. Select whether you want the pattern file to update automatically (Auto update) or not (No Auto update).
9. Click OK to save the selected values.
10. If the configuration item is saved successfully, you receive a confirmation and you must click OK again. If it is not saved successfully, you can click Details in the pop-up window that appears to discover why.
11. Under Security, in the left pane, select Anti-Virus.
12. Click Add in the right window to create a profile for the antivirus Juniper Express Engine. To edit an existing item, select it and click **Edit**.
13. In the Main tab, next to Profile name, enter a unique name for this antivirus profile.
14. Select the Profile Type. In this case, select Juniper Express.
15. Next to Trickling timeout, enter timeout parameters.

NOTE: Trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing out during a file transfer or during antivirus scanning.

16. Next to Intelligent prescreening, select Yes or No.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

Copyright © 2016, Juniper Networks, Inc.

17. Next to Content Size Limit, enter content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
18. Next to Scan engine timeout, enter scanning timeout parameters.
19. Select the Fallback settings tab.
20. Next to Default (fallback option), select Log and permit or Block from the list. In most cases, Block is the default fallback option.
21. Next to Decompress Layer (fallback option), select Log and permit or Block from the list.
22. Next to Content Size (fallback option), select Log and permit or Block from the list.
23. Next to Engine Not Ready (fallback option), select Log and permit or Block from the list.
24. Next to Timeout (fallback option), select Log and permit or Block from the list.
25. Next to Out of Resource (fallback option), select Log and permit or Block from the list.
26. Next to Too Many Requests (fallback option), select Log and permit or Block from the list.
27. Select the Notification options tab.
28. In the Fallback block section, next to Notification type, select Protocol Only or Message to select the type of notification that is sent when a fallback option of block is triggered.
29. Next to Notify mail sender, select Yes or No.
30. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
31. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
32. In the Fallback non block section, next to Notify mail recipient, select Yes or No.
33. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).
34. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
35. Select the Notification options cont tab.
36. In the Virus detection section, next to Notification type, select Protocol Only or Message to select the type of notification that is sent when a fallback option of block is triggered.
37. Next to Notify mail sender, select Yes or No.
38. If you selected Yes, next to Custom Message, enter text for the message body of your custom message for this notification (if you are using a custom message).

Copyright © 2016, Juniper Networks, Inc.

39. Next to Custom message subject, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message). The limit is 255 characters.
40. Click OK to check your configuration and save it as a candidate configuration, then click Commit Options>Commit.
41. If the configuration item is saved successfully, you receive a confirmation and you must click OK again. If it is not saved successfully, you can click Details in the pop-up that appears window to discover why.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for antivirus, the UTM policy configuration page provides separate antivirus profile selection fields for each supported protocol.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637

Example: Configuring Express Antivirus UTM Policies

This example shows how to create an express antivirus UTM policy to attach to your feature profile.

- Requirements on page 5636
- Overview on page 5636
- Configuration on page 5636
- Verification on page 5637

Requirements

Before you begin, create an antivirus feature profile. See “Example: Configuring Express Antivirus Feature Profiles” on page 5629.

Overview

In this example, you configure an express antivirus UTM policy called utmp3 and attach the policy to the antivirus profile called junexprof1.

Configuration

Step-by-Step Procedure

To configure an express antivirus UTM policy:

1. Create a UTM policy for HTTP antivirus scanning and attach the policy to the profile.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security utm utm-policy utmp3 anti-virus http-profile junexprof1
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Example: Attaching Express Antivirus UTM Policies to Security Policies on page 5638

Configuring Express Antivirus UTM Policies (J-Web Procedure)

After you have created an antivirus feature profile, configure a UTM policy to which you can attach the feature profile:

1. Select **Configure>Security>Policy>UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy. The policy configuration pop-up window appears.
3. Select the **Main** tab.
4. In the **Policy name** box, enter a unique name.
5. In the **Session per client limit** box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. In the Session per client over limit list, select the action that the device should take when the session per client limit for this UTM policy is exceeded. Options include **Log and permit** and **Block**.
7. Select the **Anti-Virus profiles** tab.
8. Select the appropriate profile you have configured from the list for the corresponding protocol listed.
9. Click **OK**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634

Copyright © 2016, Juniper Networks, Inc.

Example: Attaching Express Antivirus UTM Policies to Security Policies

This example shows how to attach an express antivirus UTM policy to a security policy.

- Requirements on page 5638
- Overview on page 5638
- Configuration on page 5638
- Verification on page 5638

Requirements

Before you begin, create a UTM policy. See "Example: Configuring Express Antivirus UTM Policies" on page 5636.

Overview

In this example, you attach the express antivirus UTM policy called utmp3 to the security policy called p3.

Configuration

Step-by-Step Procedure

To attach an express antivirus UTM policy to a security policy:

1. Enable and configure the security policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy p3 match
  source-address any
user@host# set security policies from-zone trust to-zone untrust policy p3 match
  destination-address any
user@host# set security policies from-zone trust to-zone untrust policy p3 match
  application junos-http
```

2. Attach the UTM policy to the security policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy p3 then
  permit application-services utm-policy utmp3
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter **show security policies detail** from operational mode.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623
- Example: Configuring Express Antivirus Feature Profiles on page 5629

Copyright © 2016, Juniper Networks, Inc.

- Attaching Express Antivirus UTM Policies to Security Policies (J-Web Procedure) on page 5639

Attaching Express Antivirus UTM Policies to Security Policies (J-Web Procedure)

After you create a UTM policy, create a security policy and attach the UTM policy to the security policy:

1. Select **Configure>Security>Policy>FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM. The policy configuration pop-up window appears.
3. In the Policy tab, enter a name in the **Policy Name** box.
4. Next to Default Policy Action, select one of the following: **Deny-All** or **Permit-All**.
5. Next to **From Zone**, select a zone from the list.
6. Next to **To Zone**, select a zone from the list.
7. Under Zone Direction, click **Add a Policy**.
8. Choose a **Source Address**.
9. Choose a **Destination Address**.
10. Choose an application by selecting **junos-*protocol*** (for all protocols that support antivirus scanning) in the Application Sets box and clicking the —> button to move it to the Matched box.
11. Next to Policy Action, select **Permit**.

NOTE: When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.

12. Select the **Application Services** tab.
13. Next to **UTM Policy**, select the appropriate policy from the list. This action attaches your UTM policy to the security policy.
14. Click **OK**.
15. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.
16. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your new policy to apply it.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Express Antivirus Configuration Overview on page 5623

Copyright © 2016, Juniper Networks, Inc.

- Configuring Express Antivirus Custom Objects (J-Web Procedure) on page 5626
- Configuring Express Antivirus Feature Profiles (J-Web Procedure) on page 5634
- Configuring Express Antivirus UTM Policies (J-Web Procedure) on page 5637

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 238

Configuring Express Antivirus Pattern Updates

- Understanding Express Antivirus Scanner Pattern Updates on page 5641
- Example: Automatically Updating Express Antivirus Patterns on page 5642
- Example: Automatically Updating Express Antivirus Patterns (J-Web) on page 5643
- Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure) on page 5643

Understanding Express Antivirus Scanner Pattern Updates

Express antivirus uses a different signature database than the full antivirus signature database. The express antivirus signature database is called Juniper Express antivirus database and it is compatible with the hardware engine. The express signature database targets only critical viruses and malware, including worms, Trojans, and spyware. This is a smaller sized database, providing less coverage than the full antivirus signature database.

The express antivirus pattern database is updated over HTTP or HTTPS and can occur automatically or manually. This is similar functionality to that found in full antivirus with some minor differences:

- With express antivirus, the signature database auto-update interval, is once a day.
- With express antivirus, there is no support for the downloading of multiple database types.
- With express antivirus, during database loading, all scan operations are interrupted. Scan operations for existing traffic flows are stopped and no new scan operations are initiated for newly established traffic flows. You can specify the desired action for this interruption period using the fall-back parameter for engine-busy-loading-database. The available actions are block or log-and-permit.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Once your subscription expires, you have a 30 day grace period during which you can continue to update the antivirus pattern file. Once that grace period expires, the update server no longer permits antivirus pattern file updates.

The express Antivirus scanning feature is a separately licensed subscription service. When your antivirus license key expires, you can continue to use locally stored antivirus signatures. But in that case, if the local database is deleted, antivirus scanning is disabled.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Example: Automatically Updating Express Antivirus Patterns (J-Web) on page 5643
- Example: Automatically Updating Express Antivirus Patterns on page 5642
- Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure) on page 5643

Example: Automatically Updating Express Antivirus Patterns

This example shows how to update the pattern file automatically on a security device.

- Requirements on page 5642
- Overview on page 5642
- Configuration on page 5642
- Verification on page 5643

Requirements

Before you begin:

- Obtain a valid antivirus scanner license. See "Full Antivirus Protection Overview" on page 5647.
- Get network connectivity and access to the pattern database server. See "Understanding Full Antivirus Pattern Updates" on page 5669.
- Configure your DNS settings and port settings (port 80) correctly. See *DNS Overview*.

Overview

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is once a day.)

Configuration

Step-by-Step Procedure

To configure the security device to update the pattern file automatically:

1. Set the interval.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security utm feature-profile anti-virus juniper-express-engine
  pattern-update interval 120
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Understanding Express Antivirus Scanner Pattern Updates on page 5641
- Example: Automatically Updating Express Antivirus Patterns (J-Web) on page 5643
- Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure) on page 5643

Example: Automatically Updating Express Antivirus Patterns (J-Web)

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is once a day.)

To automatically update antivirus patterns:

1. Select **Configure>Security>UTM>Anti-Virus**.
2. Next to Interval, in the Juniper Express Engine section, enter **120** in the box.
3. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options>Commit**.

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Understanding Express Antivirus Scanner Pattern Updates on page 5641
- Example: Automatically Updating Express Antivirus Patterns on page 5642
- Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure) on page 5643

Manually Updating, Reloading, and Deleting Express Antivirus Patterns (CLI Procedure)

To manually update antivirus patterns, enter the following CLI statement:

```
user@host> request security utm anti-virus juniper-express-engine pattern-update
```

To manually reload antivirus patterns, enter the following CLI statement:

```
user@host> request security utm anti-virus juniper-express-engine pattern-reload
```

Copyright © 2016, Juniper Networks, Inc.

To manually delete antivirus patterns, enter the following CLI statement:

```
user@host> request security utm anti-virus juniper-express-engine pattern-delete
```

Related Documentation

- Express Antivirus Protection Overview on page 5621
- Understanding Express Antivirus Scanner Pattern Updates on page 5641
- Example: Automatically Updating Express Antivirus Patterns (J-Web) on page 5643
- Example: Automatically Updating Express Antivirus Patterns on page 5642

Copyright © 2016, Juniper Networks, Inc.

PART 74

Configuring Full Antivirus Protection and Pattern Updates

- Configuring Full Antivirus Protection on page 5647
- Configuring Full Antivirus Pattern Updates on page 5669
- Configuring File Scanning on page 5675
- Configuring Scan Results and Fallback Options on page 5687
- Configuring Application Protocol Scanning on page 5695
- Configuring Whitelists on page 5705
- Configuring HTTP Trickling on page 5709
- Configuring Notifications on page 5711

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 239

Configuring Full Antivirus Protection

- Full Antivirus Protection Overview on page 5647
- Full Antivirus Configuration Overview on page 5648
- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Configuring Full Antivirus Custom Objects on page 5650
- Configuring Full Antivirus Custom Objects (J-Web Procedure) on page 5653
- Example: Configuring Full Antivirus Feature Profiles on page 5655
- Configuring Full Antivirus Feature Profiles (J-Web Procedure) on page 5661
- Example: Configuring Full Antivirus UTM Policies on page 5664
- Configuring Full Antivirus UTM Policies (J-Web Procedure) on page 5665
- Example: Attaching Full Antivirus UTM Policies to Security Policies on page 5665
- Attaching Full Antivirus UTM Policies to Security Policies (J-Web Procedure) on page 5666

Full Antivirus Protection Overview

A virus is executable code that infects or attaches itself to other executable code in order to reproduce itself. Some malicious viruses erase files or lock up systems, while other viruses merely infect files and can overwhelm the target host or network with bogus data. The full file-based antivirus feature provides file-based scanning on specific Application Layer traffic checking for viruses against a virus signature database. It collects the received data packets until it has reconstructed the original application content, such as an e-mail file attachment, and then scans this content.

The full file-based antivirus scanning feature is a separately licensed subscription service. Kaspersky Lab provides the scan engine for full file-based antivirus. When your antivirus license key expires, you can continue to use locally stored antivirus signatures without any updates. But in that case, if the local database is deleted, antivirus scanning is disabled.

NOTE: The express antivirus feature provides better performance but lower security. Note that if you switch from full file-based antivirus protection to express antivirus protection, you must reboot the device in order for express antivirus to begin working.

Copyright © 2016, Juniper Networks, Inc.

The Kaspersky scan engine is provided as a downloadable UTM module. To download the Kaspersky scan engine, your SRX Series device must have an active UTM license. When you install the KAV license, the system automatically downloads the Kaspersky module from the Juniper Networks server and runs it.

When you set the antivirus type to KAV, and if the SRX Series device had a preinstalled Kaspersky engine, then the downloaded module replaces the original module on the device. Regardless of the UTM license status, when the KAV license is deleted from the device, the Kaspersky engine and all files associated with KAV are removed from the system immediately.

Use the **set security utm feature-profile anti-virus type kaspersky-lab-engine** command to set the antivirus type to KAV. If Kaspersky engine is not available on the device, and if the Kaspersky engine cannot be downloaded from the predefined URL, then use the **set security utm feature-profile anti-virus kaspersky-lab-engine pattern-update url *url*** command to configure the downloading application URL.

Related Documentation

- Understanding Full Antivirus Pattern Updates on page 5669
- Full Antivirus Pattern Update Configuration Overview on page 5649
- Understanding Full Antivirus Scan Level Settings on page 5678
- Understanding the Full Antivirus Internal Scan Engine on page 5675
- Full Antivirus Configuration Overview on page 5648

Full Antivirus Configuration Overview

When configuring antivirus protection, you must first create the antivirus custom objects you are using. Those custom objects may include the MIME pattern list, MIME exception list, and the filename extension list. Once you have created your custom objects, you can configure full antivirus protection, including intelligent prescreening, and content size limits.

To configure full file-based antivirus protection:

1. Configure UTM custom objects for the UTM feature. The following example enables the mime-pattern, filename-extension, url-pattern, and custom-url-category custom-objects:

   ```
   user@host# set security utm custom-objects mime-pattern
   user@host# set security utm custom-objects filename-extension
   user@host# set security utm custom-objects url-pattern
   user@host# set security utm custom-objects custom-url-category
   ```

2. Configure the main feature parameters using feature profiles. The following example enables options using the anti virus feature profile:

   ```
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
     pattern-update
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
     fallback-options
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
  notification-options
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
  scan-options
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
  trickling
user@host# set security utm feature-profile anti-virus mime-whitelist
user@host# set security utm feature-profile anti-virus url-whitelist
```

3. Configure a UTM policy for each protocol and attach this policy to a profile. The following example configure the utmp2 UTM policy for the HTTP protocol:

```
user@host# set security utm utm-policy utmp2 anti-virus http-profile http1
```

4. Attach the UTM policy to a security policy. The following example attaches the utmp2 UTM policy to the p2 security policy:

```
user@host# set security policies from-zone trust to-zone untrust policy p2 then permit
  application-services utm-policy utmp2
```

Related Documentation

- Full Antivirus Protection Overview on page 5647

Full Antivirus Pattern Update Configuration Overview

Before you begin, there are several prerequisites that must be met in order to perform a successful pattern database update:

- You must have a valid antivirus scanner license.
- You must have network connectivity and access to the pattern database server.
- Your DNS settings and port settings (port 80) must be correct.

To update the patterns for the antivirus signature database:

1. On the security device, specify the URL address of the pattern-update server.
2. (Optional) Specify how often the device should automatically check for pattern-server updates.

After the security device downloads the server-initialization file, the device checks that the pattern file is valid. The device then parses the file to obtain information about it, including the file version, size, and location of the pattern file server.

If the pattern file on the security device is out-of-date (or nonexistent because this is the first time you are loading it), and, if the antivirus pattern-update service subscription is still valid, the device automatically retrieves an updated pattern file from the pattern file server.

The following is an example of the CLI for configuring the database update feature:

```
utm {
  feature-profile {
    anti-virus {
      type
```

Copyright © 2016, Juniper Networks, Inc.

```
            kaspersky-lab-engine {
              pattern-update
                url url
              interval minutes
            }
          }
        }
      }
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Pattern Updates on page 5669
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670
- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Example: Automatically Updating Full Antivirus Patterns on page 5672
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Example: Configuring Full Antivirus Custom Objects

This example shows how to configure full antivirus custom objects.

- Requirements on page 5650
- Overview on page 5650
- Configuration on page 5651
- Verification on page 5653

Requirements

Before you begin:

- Decide the type of full antivirus protection you require. See "Full Antivirus Protection Overview" on page 5647.
- Understand the order in which full antivirus parameters are configured. See "Full Antivirus Pattern Update Configuration Overview" on page 5649.

Overview

In this example, you define custom objects that are used to create full antivirus feature profiles. You perform the following tasks to define custom objects:

1. Configure a filename extension list called extlist1 and add extensions such as .zip, .js, and .vbs to the list.
2. Create two MIME lists called avmime1 and ex-avmime1 and add patterns to the list.
3. Configure a URL pattern list called urllist1.
4. Configure a custom URL category list called custurl1 using the urllist1 URL pattern list.

Copyright © 2016, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects filename-extension extlist1 value [zip js vbs]
set security utm custom-objects mime-pattern avmime1 value [video/quicktime
  image/x-portable-anymap x-world/x-vrml]
set security utm custom-objects mime-pattern ex-avmime1 value
  [video/quicktime-inappropriate]
set security utm custom-objects url-pattern urllist1 value [http://www.url.com 5.6.7.8]
set security utm custom-objects custom-url-category custurl1 value urllist1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure full antivirus filtering custom objects:

1. Configure the filename extension list and add extensions to it.

   ```
   [edit security utm]
   user@host# set custom-objects filename-extension extlist1 value [zip js vbs]
   ```

   NOTE: The Kaspersky scan engine ships with a read-only default extension list that you can use.

2. Create MIME lists and add MIME patterns to the lists.

   ```
   [edit security utm]
   user@host# set custom-objects mime-pattern avmime1 value [video/quicktime
     image/x-portable-anymap x-world/x-vrml]
   user@host# set custom-objects mime-pattern ex-avmime1 value
     [video/quicktime-inappropriate]
   ```

3. Configure a URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllist1 value [http://www.url.com
     5.6.7.8]
   ```

   When entering the URL pattern, note the following wildcard character support:

   - The *\.[]\?* wildcard characters are supported.
   - You must precede all wildcard URLs with http://.
   - You can only use the asterisk * wildcard character if it is at the beginning of the URL and is followed by a period.
   - You can only use the question mark ? wildcard character at the end of the URL.

Copyright © 2016, Juniper Networks, Inc.

- The following wildcard syntax is supported: http://*.example.net, http://www.example.ne?, http://www.example.n??.
- The following wildcard syntax is not supported: *.example.net , www.example.ne?, http://*example.net, http://*.

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

4. Configure a custom URL category list.

```
[edit security utm]
user@host# set custom-objects custom-url-category custurl1 value urllist1
```

Results From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost# show security utm
custom-objects {
  mime-pattern {
    avmime1 {
      value [ video/quicktime image/x-portable-anymap x-world/x-vrml ];
    }
    ex-avmime1 {
      value video/quicktime-inappropriate;
    }
  }
  filename-extension {
    extlist1 {
      value [ zip js vbs ];
    }
  }
  url-pattern {
    urllist1 {
      value [ http://www.url.com 5.6.7.8 ];
    }
  }
  custom-url-category {
    custurl1 {
      value urllist1;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Verifying Full Antivirus Custom Objects

Purpose Verify the full antivirus custom objects.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Example: Configuring Full Antivirus Feature Profiles on page 5655
- Configuring Full Antivirus Feature Profiles (J-Web Procedure) on page 5661

Configuring Full Antivirus Custom Objects (J-Web Procedure)

To configure antivirus protection, you must first create your custom objects (MIME Pattern List, Filename Extension List, URL Pattern List, and Custom URL Category List).

Configure a MIME pattern list custom object:

1. Select **Configure>Security>UTM>Custom Objects**.
2. From the MIME Pattern List tab, click the **Add** button to create MIME pattern lists.
3. In the Add MIME Pattern pop-up window, next to **MIME Pattern Name**, enter a unique name.

   NOTE: Keep in mind that you are creating a MIME whitelist and a MIME exception list (if necessary). Both MIME lists appear in the MIME Whitelist and Exception MIME Whitelist fields when you configure antivirus. Therefore, the MIME list names you create should be as descriptive as possible.

4. Next to **MIME Pattern Value**, enter the MIME pattern.
5. Click **Add** to add your MIME pattern to the Values list box. Within this box, you can also select an entry and use the Delete button to delete it from the list. Continue to add MIME patterns in this manner.
6. Optionally, create a new MIME list to act as an exception list. The exception list is generally a subset of the main MIME list.
7. Click **OK** to check your configuration and save the selected values as part of the MIME list, then click **Commit Options>Commit**.
8. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Copyright © 2016, Juniper Networks, Inc.

Configure a filename extension list custom object:

1. Select **Configure>Security>UTM>Custom Objects**.
2. From the Filename Extension List tab, click the **Add** button to create filename extension lists.
3. Next to **File Extension Name**, enter a unique name. This name appears in the Scan Option By Extension list when you configure an antivirus profile.
4. In the **Available Values** box, select one or more default values (press Shift to select multiple concurrent items or press Ctrl to select multiple separate items) and click the right arrow button to move the value or values to the Selected Values box.
5. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a URL pattern list custom object:

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.

1. Select **Configure>Security>UTM>Custom Objects**.
2. From the URL Pattern List tab, click the **Add** button to create URL pattern lists.
3. Next to URL Pattern Name, enter a unique name. This name appears in the Custom URL Category List Custom Object page for selection.
4. Next to **URL Pattern Value**, enter the URL or IP address you want added to the list for bypassing scanning.

   When entering the URL pattern, note the following wildcard character support:

   - The ***\.[]\?*** wildcard characters are supported.
   - You must precede all wildcard URLs with **http://**.
   - You can only use the asterisk ***** wildcard character if it is at the beginning of the URL and is followed by a period.
   - You can only use the question mark **?** wildcard character at the end of the URL.
   - The following wildcard syntax IS supported: **http://*.example.net**, **http://www.example.ne?**, **http://www.example.n??**.
   - The following wildcard syntax is NOT supported: ***.example.net** , **www.example.ne?**, **http://*example.net, http://***.

Copyright © 2016, Juniper Networks, Inc.

5. Click **Add** to add your URL pattern to the Values list box. The list can contain up to 8192 items. You can also select an entry and use the Delete button to delete it from the list. Continue to add URLs or IP addresses in this manner.
6. Click **OK** to check your configuration and save the selected values as part of the URL pattern list you have created, then click **Commit Options**>**Commit**.
7. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Configure a custom URL category list custom object:

NOTE: **Because you use URL Pattern Lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure a custom URL category list.**

1. Select **Configure>Security>UTM>Custom Objects**.
2. In the URL Category List tab, click **Add** to create URL category lists.
3. Next to **URL Category Name**, enter a unique name. This name appears in the URL Whitelist list when you configure antivirus global options.
4. In the **Available Values** box, select a URL Pattern List name from the list for bypassing scanning and click the right arrow button to move it to the Selected Values box.
5. Click **OK** to check your configuration and save the selected values as part of the URL list that you have created, then click **Commit Options**>**Commit**.

   Click **OK** to save the selected values as part of the custom URL list you have created.
6. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Feature Profiles (J-Web Procedure) on page 5661

Example: Configuring Full Antivirus Feature Profiles

This example shows how to configure a full antivirus feature profile.

- Requirements on page 5656
- Overview on page 5656
- Configuration on page 5657
- Verification on page 5661

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

- Decide the type of full antivirus protection you require. See "Full Antivirus Protection Overview" on page 5647.
- Understand the order in which full antivirus parameters are configured. See "Full Antivirus Configuration Overview" on page 5648.
- MIME patterns must be defined for lists and exception lists. See "Example: Configuring MIME Whitelists to Bypass Antivirus Scanning" on page 5706.

Overview

In this example, you configure a feature profile called kasprof1 and specify custom objects to be used for filtering content:

- Select and configure the engine type as Kaspersky Lab Engine.
- Select 120 as the time interval for updating the pattern database. The default full file-based antivirus pattern-update interval is 60 minutes.

NOTE: The command for changing the URL for the pattern database is:

```
[edit]
user@host# edit security utm feature-profile anti-virus kaspersky-lab-engine
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set pattern-update url http://..
```

The default URL is http://update.juniper-updates.net/AV/<device-version>. You should not change this URL unless you are experiencing problems with it and have called for support.

- Enable an e-mail notification with a custom message as pattern file was updated and a custom subject line as AV pattern file updated.
- Configure a list of fallback options as block.
- Configure the notification options for fallback blocking for virus detection. Configure a custom message for the fallback blocking action.
- Configure a notification for protocol-only virus detection.
- Configure scan options. For this example, configure the device to perform a TCP payload content size check before the scan request is sent.
- Configure the decompression layer limit. For this example configure the device to decompress three layers of nested compressed files before it executes the virus scan.
- Configure content size parameters as 20000.
- Configure scan extension settings. The default list is junos-default-extension. For this example, you select extlist1, which you created as a custom object.

Copyright © 2016, Juniper Networks, Inc.

- Configure the scan mode setting to configure the device to use a custom extension list. Although you can choose to scan all files, for this example you select only files with the extensions that you specify.
- Enable intelligent prescreening and set its timeout setting to 1800 seconds and trickling setting (applicable only to HTTP) to 600 seconds. This means that if the device receives a packet within a 600-second period during a file transfer or while performing an antivirus scan, it should not time out.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP) and HTTP POST.

The following example disables intelligent prescreening for the kasprof1 profile:

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options no-intelligent-prescreening
```

- Configure the antivirus scanner to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list that ships with the device called junos-default-bypass-mime. For this example, you use the avmime1 and ex-avmime1 lists.
- Configure the antivirus module to use URL bypass lists. If you are using a URL whitelist (valid only for HTTP traffic), this is a custom URL category that you have previously configured as a custom object. For this example, you enable the custurl1 bypass list.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile anti-virus type kaspersky-lab-engine
set security utm feature-profile anti-virus kaspersky-lab-engine pattern-update interval
  120
set security utm feature-profile anti-virus kaspersky-lab-engine pattern-update
  email-notify admin-email administrator@example.net custom-message
  patternfilewasupdated custom-message-subject AVpatternfileupdated
set security utm feature-profile anti-virus type kaspersky-lab-engine
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options content-size block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options corrupt-file block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options decompress-layer block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options default block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options engine-not-ready block
```

Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options out-of-resources block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options password-file block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options timeout block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options too-many-requests block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  notification-options fallback-block custom-message "Dropped due to fallback settings"
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  notification-options virus-detection type protocol-only
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options content-size-limit 20000
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options decompress-layer-limit 3
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options intelligent-prescreening
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options scan-extension extlist1
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options scan-mode by-extension
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  scan-options timeout 1800
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1 trickling
  timeout 600
set security utm feature-profile anti-virus mime-whitelist list avmime1
set security utm feature-profile anti-virus mime-whitelist list avmime1 exception
  ex-avmime1
set security utm feature-profile anti-virus url-whitelist custurl1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure full antivirus feature profiles:

1. Select and configure the engine type.

```
[edit]
user@host# set security utm feature-profile anti-virus type kaspersky-lab-engine
user@host#set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update interval 120
```

2. Configure the device to notify a specified administrator when patterns are updated.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set pattern-update email-notify admin-email
  administrator@example.net custom-message patternfilewasupdated
  custom-message-subject AVpatternfileupdated
```

3. Create a profile for the Kaspersky Lab engine and configure fallback options as block.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 fallback-options content-size block
user@host# set profile kasprof1 fallback-options corrupt-file block
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set profile kasprof1 fallback-options decompress-layer block
user@host# set profile kasprof1 fallback-options default block
user@host# set profile kasprof1 fallback-options engine-not-ready block
user@host# set profile kasprof1 fallback-options out-of-resources block
user@host# set profile kasprof1 fallback-options password-file block
user@host# set profile kasprof1 fallback-options timeout block
user@host# set profile kasprof1 fallback-options too-many-requests block
```

4. Configure a custom notification for the fallback blocking action and send a notification.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 notification-options fallback-block custom-message
  "Dropped due to fallback settings"
```

5. Configure a notification for protocol-only virus detection.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 notification-options virus-detection type
  protocol-only
```

6. Configure content size parameter.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options content-size-limit 20000
```

7. Configure the decompression layer limit.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options decompress-layer-limit 3
```

8. Configure intelligent prescreening.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options intelligent-prescreening
```

9. Configure scan extension setting.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options scan-extension extlist1
```

10. Configure the scan mode setting.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options scan-mode by-extension
```

11. Configure the timeout setting.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 scan-options timeout 1800
```

12. Configure trickling setting.

```
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set profile kasprof1 trickling timeout 600
```

13. Configure the antivirus scanner to use MIME bypass lists and exception lists.

```
[edit security utm feature-profile anti-virus]
user@host# set mime-whitelist list avmime1
user@host# set mime-whitelist list avmime1 exception ex-avmime1
```

14. Configure the antivirus module to use URL bypass lists.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security utm feature-profile anti-virus]
user@host# set url-whitelist custurl1
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile anti-virus** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm feature-profile anti-virus
mime-whitelist {
  list avmime1;
  exception ex-avmime1;
}
url-whitelist custurl1;
kaspersky-lab-engine {
  pattern-update {
    email-notify {
      admin-email "administrator@example.net";
      custom-message patternfilewasupdated;
      custom-message-subject AVpatternfileupdated;
    }
    interval 120;
  }
  profile kasprof1 {
    fallback-options {
      default block;
      corrupt-file block;
      password-file block;
      decompress-layer block;
      content-size block;
      engine-not-ready block;
      timeout block;
      out-of-resources block;
      too-many-requests block;
    }
    scan-options {
      intelligent-prescreening;
      scan-mode by-extension;
      scan-extension extlist1;
      content-size-limit 20000;
      timeout 1800;
      decompress-layer-limit 3;
    }
    trickling timeout 600;
    notification-options {
      virus-detection {
        type protocol-only;
        custom-message ***virus-found***;
      }
      fallback-block {
        custom-message “Dropped due to fallback settings”;
      }
    }
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration of Full Antivirus Feature Profile

Purpose Verify the full antivirus feature profile.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Full Antivirus Configuration Overview on page 5648
- Example: Configuring Full Antivirus UTM Policies on page 5664
- Example: Attaching Full Antivirus UTM Policies to Security Policies on page 5665

Configuring Full Antivirus Feature Profiles (J-Web Procedure)

After you have created your custom object, configure an antivirus feature profile:

1. Select **Configure>Security>UTM>Global options**.
2. In the Anti-Virus tab, next to **MIME whitelist**, select the custom object you created from the list.
3. Next to **Exception MIME whitelist**, select the custom object you created from the list.
4. Next to **URL Whitelist**, select the custom object you created from the list.
5. In the **Engine Type** section, select the type of engine you are using. For full antivirus protection, you should select **Kaspersky Lab**.
6. In the Kaspersky Lab Engine Option section, in the **Pattern update URL** box, enter the URL for the pattern database.

   NOTE: The URL is http://update.juniper-updates.net/AV/<device version> and you should not change it.

7. Next to **Pattern update interval**, enter the time interval, in seconds, for automatically updating the pattern database in the box. The default interval is 60.
8. Select whether you want the pattern file to update automatically (**Auto update**) or not (**No Auto update**).
9. Click **OK** to save the selected values.
10. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in a pop-up window that appears to discover why.
11. Under Security, in the left pane, select **Anti-Virus**.

Copyright © 2016, Juniper Networks, Inc.

12. In the right window, click **Add** to create a profile for the antivirus Kaspersky Lab Engine. (To edit an existing item, select it and click the **Edit** button.)
13. Next to **Profile name**, enter a unique name for this antivirus profile.
14. Select the **Profile Type**. In this case, select **Kaspersky**.
15. Next to **Trickling timeout**, enter timeout parameters.

NOTE: Trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing out during a file transfer or during antivirus scanning.

16. Next to Intelligent prescreening, select **Yes** or **No**.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

17. In the Scan Options section, next to Intelligent prescreening, select **Yes** if you are using it.

NOTE: Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for mail protocols (SMTP, POP3, IMAP, and HTTP POST).

18. Next to **Content Size Limit**, enter content size parameters. The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
19. Next to **Scan engine timeout**, enter scanning timeout parameters.
20. Next to **Decompress Layer Limit**, enter decompression layer limit parameters.
21. In the Scan mode section, select either **Scan all files**, if you are scanning all content, or **Scan files with specified extension**, if you are scanning by file extensions.

NOTE: If you select Scan files with specified extension, you must select a filename extension list custom object from the Scan engine filename extention list that appears.

22. Select the **Fallback settings** tab.
23. Next to Default (fallback option), select **Log and permit** or **Block** from the list. In most cases, Block is the default fallback option.
24. Next to Corrupt File (fallback option), select **Log and permit** or **Block** from the list.
25. Next to Password File (fallback option), select **Log and permit** or **Block** from the list.

Copyright © 2016, Juniper Networks, Inc.

26. Next to Decompress Layer (fallback option), select **Log and permit** or **Block** from the list.
27. Next to Content Size (fallback option), select **Log and permit** or **Block** from the list.
28. Next to Engine Not Ready (fallback option), select **Log and permit** or **Block** from the list.
29. Next to Timeout (fallback option), select **Log and permit** or **Block** from the list.
30. Next to Out Of Resources (fallback option), select **Log and permit** or **Block** from the list.
31. Next to Too Many Request (fallback option), select **Log and permit** or **Block** from the list.
32. Select the **Notification options** tab.
33. In the Fallback block section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.
34. Next to Notify mail sender, select **Yes** or **No**.
35. If you selected Yes, next to **Custom Message**, enter text for the message body of your custom message for this notification (if you are using a custom message).
36. Next to **Custom message subject**, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
37. In the Fallback non block section, next to Notify mail recipient, select **Yes** or **No**.
38. If you selected Yes, next to **Custom Message**, enter text for the message body of your custom message for this notification (if you are using a custom message).
39. Next to **Custom message subject**, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message).
40. Select the **Notification options cont** tab.
41. In the Virus detection section, next to Notification type, select **Protocol Only** or **Message** to select the type of notification that is sent when a fallback option of block is triggered.
42. Next to Notify mail sender, select **Yes** or **No**.
43. If you selected Yes, next to **Custom Message**, enter text for the message body of your custom message for this notification (if you are using a custom message).
44. Next to **Custom message subject**, enter text to appear in the subject line of your custom message for this notification (if you are using a custom message). The limit is 255 characters.
45. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.
46. If the configuration item is saved successfully, you receive a confirmation and you must click **OK** again. If it is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Copyright © 2016, Juniper Networks, Inc.

NOTE: You create a separate antivirus profile for each antivirus protocol. These profiles may basically contain the same configuration information, but when you are creating your UTM policy for an antivirus profile, the UTM policy configuration page provides separate antivirus profile selection fields for each supported protocol.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Custom Objects (J-Web Procedure) on page 5653

Example: Configuring Full Antivirus UTM Policies

This example shows how to create a UTM policy to attach to a feature profile.

- Requirements on page 5664
- Overview on page 5664
- Configuration on page 5664
- Verification on page 5664

Requirements

Before you begin, create an antivirus feature profile. See "Example: Configuring Full Antivirus Feature Profiles" on page 5655.

Overview

In this example, you configure a full antivirus UTM policy called utmp2 and attach the policy to an HTTP profile called kasprofile1 HTTP.

Configuration

Step-by-Step Procedure

To configure a full antivirus UTM policy:

1. Create a UTM policy for HTTP antivirus scanning and attach the policy to the profile.

   ```
   [edit]
   user@host# set security utm utm-policy utmp2 anti-virus http-profile kasprofile1
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Configuration Overview on page 5648
- Example: Configuring Full Antivirus Feature Profiles on page 5655

Copyright © 2016, Juniper Networks, Inc.

- Example: Attaching Full Antivirus UTM Policies to Security Policies on page 5665

Configuring Full Antivirus UTM Policies (J-Web Procedure)

After you have created an antivirus feature profile, configure a UTM policy to which you can attach the feature profile:

1. Select **Configure>Security>Policy>UTM Policies**.
2. From the UTM policy configuration window, click **Add** to configure a UTM policy. This action takes you to the policy configuration pop-up window.
3. Select the **Main** tab in pop-up window.
4. In the **Policy name** box, enter a unique name for the UTM policy.
5. In the **Session per client limit** box, enter a session per client limit from 0 to 20000 for this UTM policy.
6. In the **Session per client over limit** list, select the action that the device should take when the session per client limit for this UTM policy is exceeded. Options include **Log and permit** and **Block**.
7. Select the **Anti-Virus profiles** tab in the pop-up window.
8. Select the appropriate profile you have configured from the list for the corresponding protocol listed.
9. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.
10. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Custom Objects (J-Web Procedure) on page 5653
- Configuring Full Antivirus Feature Profiles (J-Web Procedure) on page 5661

Example: Attaching Full Antivirus UTM Policies to Security Policies

This example shows how to attach a UTM policy to a security policy.

- Requirements on page 5666
- Overview on page 5666
- Configuration on page 5666
- Verification on page 5666

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, create a UTM policy. See "Example: Configuring Full Antivirus UTM Policies" on page 5664.

Overview

In this example, you attach the UTM policy called utmp2 to the security policy called p2.

Configuration

Step-by-Step Procedure

To attach a full antivirus UTM policy to a security policy:

1. Enable and configure the security policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy p2 match
  source-address any
user@host# set security policies from-zone trust to-zone untrust policy p2 match
  destination-address any
user@host# set security policies from-zone trust to-zone untrust policy p2 match
  application junos-http
```

2. Attach the UTM policy to the security policy.

```
[edit]
user@host# set security policies from-zone trust to-zone untrust policy p2 then
  permit application-services utm-policy utmp2
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security policies** command.

Related Documentation

- Full Antivirus Configuration Overview on page 5648
- Example: Configuring Full Antivirus UTM Policies on page 5664

Attaching Full Antivirus UTM Policies to Security Policies (J-Web Procedure)

After you create a UTM policy, create a security policy and attach the UTM policy to the security policy:

1. Select **Configure>Security>Policy>FW Policies**.
2. From the Security Policy window, click **Add** to configure a security policy with UTM. This action takes you to the policy configuration pop-up window.
3. In the Policy tab, enter a name in the **Policy Name** box.
4. Next to **From Zone**, select a zone from the list.

Copyright © 2016, Juniper Networks, Inc.

5. Next to **To Zone**, select a zone from the list.
6. Choose a **Source Address**.
7. Choose a **Destination Address**.
8. Choose an application by selecting **junos-*protocol*** (for all protocols that support antivirus scanning) in the Application Sets box and clicking the —> button to move it to the Matched box.
9. Next to Policy Action, select **Permit**.

NOTE: **When you select Permit for Policy Action, several additional fields become available in the Applications Services tab, including UTM Policy.**

10. Select the **Application Services** tab in the pop-up window.
11. Next to **UTM Policy**, select the appropriate policy from the list. This action attaches your UTM policy to the security policy.
12. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options**>**Commit**.
13. If the policy is saved successfully, you receive a confirmation and you must click **OK** again. If the profile is not saved successfully, you can click **Details** in the pop-up window that appears to discover why.

    You must activate your new policy to apply it.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Custom Objects (J-Web Procedure) on page 5653
- Configuring Full Antivirus Feature Profiles (J-Web Procedure) on page 5661
- Configuring Full Antivirus UTM Policies (J-Web Procedure) on page 5665

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 240

Configuring Full Antivirus Pattern Updates

- Understanding Full Antivirus Pattern Updates on page 5669
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670
- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Example: Automatically Updating Full Antivirus Patterns on page 5672
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Understanding Full Antivirus Pattern Updates

The full file-based antivirus protection signature database is called the Juniper Full antivirus database (downloaded by the pattern-update command). This database is different from the database used by express antivirus. It detects all destructive malicious code, including viruses (polymorphic and other advanced virus types), worms, Trojans, and malware.

Updates to the pattern file are added as new viruses are discovered. When Kaspersky Lab updates the signatures in its pattern database, the security device downloads these updates so that the antivirus scanner is using the latest, most up-to-date signatures when scanning traffic. The security device can perform these updates automatically (the default), or you can perform pattern update downloads manually.

The database pattern server is accessible through HTTP or HTTPS. By default, the antivirus module checks for database updates automatically every 60 minutes. You can change this interval and you can trigger updates manually, as well. The number of files that are downloaded during an update and the duration of the download process can vary.

A local copy of the pattern database is saved in persistent data storage (that is, the flash disk). If the device is rebooted, the local copy remains available for the antivirus scan engine to use during the antivirus scan engine initialization time, without the need for network access to the pattern database server.

NOTE: If the auto-update fails, the updater automatically retries to update three more times. If the database download continues to fail, the updater stops trying and waits for the next periodic update before trying again.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Once your subscription expires, you have a 30 day grace period during which you can continue to update the antivirus pattern file. Once that grace period expires, the update server no longer permits antivirus pattern file updates.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670
- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Example: Automatically Updating Full Antivirus Patterns on page 5672
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Example: Configuring the Full Antivirus Pattern Update Server

This example shows how to configure the pattern-update server on the security device.

- Requirements on page 5670
- Overview on page 5670
- Configuration on page 5671
- Verification on page 5671

Requirements

Before you begin:

- Obtain a valid antivirus scanner license. See "Full Antivirus Protection Overview" on page 5647.
- Get network connectivity and access to the pattern database server. See "Understanding Full Antivirus Pattern Updates" on page 5669.
- Configure your DNS settings and port settings (port 80) correctly. See *DNS Overview*.

Overview

To configure the pattern-update server on the security device, enter the URL address of the pattern-update server.

By default, the Juniper-Kaspersky URL for full antivirus protection is http://update.juniper-updates.net/AV/*device-name*, where *device-name* is the name of your device.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To configure the pattern-update server on a security device:

1. Specify the URL of the pattern-update server.

```
[edit]
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update url http://update.juniper-updates.net/AV/device-name
```

NOTE: Other than the platform name, you should not change this URL unless you are experiencing problems with it and have called for support.

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Example: Automatically Updating Full Antivirus Patterns on page 5672
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Example: Automatically Updating Full Antivirus Patterns (J-Web)

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is 60 minutes.)

To automatically update antivirus patterns:

1. Select **Configure>UTM>Anti-Virus**.
2. Next to Interval, in the Kaspersky Lab Engine section, enter **120** in the box.
3. Click **OK** to check your configuration and save it as a candidate configuration, then click **Commit Options>Commit**.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Pattern Updates on page 5669
- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670

Copyright © 2016, Juniper Networks, Inc.

- Example: Automatically Updating Full Antivirus Patterns on page 5672
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Example: Automatically Updating Full Antivirus Patterns

This example shows how to update the pattern file automatically on a security device.

- Requirements on page 5672
- Overview on page 5672
- Configuration on page 5672
- Verification on page 5672

Requirements

Before you begin:

- Obtain a valid antivirus scanner license. See "Full Antivirus Protection Overview" on page 5647.
- Get network connectivity and access to the pattern database server. See "Understanding Full Antivirus Pattern Updates" on page 5669.
- Configure your DNS settings and port settings (port 80) correctly. See *DNS Overview*.

Overview

In this example, you configure the security device to update the pattern file automatically every 120 minutes. (The default antivirus pattern-update interval is 60 minutes.)

Configuration

Step-by-Step Procedure

To configure the security device to update the pattern file automatically:

1. Set the interval.

```
[edit]
user@host# set security utm feature-profile anti-virus kaspersky-lab-engine
  pattern-update interval 120
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670

Copyright © 2016, Juniper Networks, Inc.

- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure) on page 5673

Manually Updating, Reloading, and Deleting Full Antivirus Patterns (CLI Procedure)

To manually update antivirus patterns, enter the following CLI command:

```
user@host> request security utm anti-virus kaspersky-lab-engine pattern-update
```

To manually reload antivirus patterns, enter the following CLI command:

```
user@host> request security utm anti-virus kaspersky-lab-engine pattern-reload
```

To manually delete antivirus patterns, enter the following CLI command:

```
user@host> request security utm anti-virus kaspersky-lab-engine pattern-delete
```

You can update the Kaspersky antivirus signature database offline without using a direct Internet connection. This is required in some security installations and for sites that access the Internet through a proxy server.

To update the Kaspersky antivirus signature database offline, you must configure a local webserver.

Copyright © 2016, Juniper Networks, Inc.

To configure a webserver, use the following CLI statement.

user@host# set security utm feature-profile anti-virus kaspersky-lab-engine pattern-update url <*http_server*>

user@host# commit

To update the Kaspersky antivirus signature database, perform the following tasks:

1. Based on your hardware platform, enter the following URLs in your computer browser.
2. Copy all the files to a directory on your local webserver. You might want to use a download manager for your browser to get all the files more quickly.
3. Download the Kaspersky Lab engine from http://update.juniper-updates.net/KAV_engine/.
   - For JSR, the URL is http://update.juniper-updates.net/KAV_engine/i386/.
   - For branch SRX Series devices, the URL is http://update.juniper-updates.net/KAV_engine/octeon32/.
4. Copy all the files to the same directory on your local server.

   NOTE: **The Kaspersky Lab engine is automatically loadable. For updating the Kaspersky antivirus signature database offline, both pattern update files and Kaspersky Lab engine files must be placed in the same folder on the local webserver.**

5. Set the directory as a sharepoint that can be accessed through HTTP from the SRX Series device.
6. Run the update command in the CLI.

   user@host>request security utm anti-virus kaspersky-lab-engine pattern-update

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Pattern Updates on page 5669
- Full Antivirus Pattern Update Configuration Overview on page 5649
- Example: Configuring the Full Antivirus Pattern Update Server on page 5670
- Example: Automatically Updating Full Antivirus Patterns (J-Web) on page 5671
- Example: Automatically Updating Full Antivirus Patterns on page 5672

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 241

Configuring File Scanning

- Understanding the Full Antivirus Internal Scan Engine on page 5675
- Understanding Full Antivirus Scan Mode Support on page 5676
- Configuring Full Antivirus File Extension Scanning (CLI Procedure) on page 5677
- Example: Configuring Full Antivirus File Extension Scanning on page 5677
- Understanding Full Antivirus Scan Level Settings on page 5678
- Example: Configuring Full Antivirus Scan Settings at Different Levels on page 5679
- Understanding Full Antivirus Intelligent Prescreening on page 5681
- Example: Configuring Full Antivirus Intelligent Prescreening on page 5681
- Understanding Full Antivirus Content Size Limits on page 5682
- Configuring Full Antivirus Content Size Limits (CLI Procedure) on page 5683
- Understanding Full Antivirus Decompression Layer Limits on page 5683
- Configuring Full Antivirus Decompression Layer Limits (CLI Procedure) on page 5684
- Understanding Full Antivirus Scanning Timeouts on page 5684
- Configuring Full Antivirus Scanning Timeouts (CLI Procedure) on page 5684
- Understanding Full Antivirus Scan Session Throttling on page 5685
- Configuring Full Antivirus Scan Session Throttling (CLI Procedure) on page 5685

Understanding the Full Antivirus Internal Scan Engine

The full file-based antivirus module is the software subsystem on the gateway device that scans specific Application Layer traffic to protect users from virus attacks and to prevent viruses from spreading. The antivirus software subsystem consists of a virus signature database, an application proxy, the scan manager, and the scan engine.

Kaspersky Lab provides the scan engine and it works in the following manner:

1. A client establishes a TCP connection with a server and then starts a transaction.
2. If the application protocol in question is marked for antivirus scanning, the traffic is forwarded to an application proxy for parsing.
3. When the scan request is sent, the scan engine scans the data by querying a virus pattern database.

Copyright © 2016, Juniper Networks, Inc.

4. The scan manager monitors antivirus scanning sessions, checking the properties of the data content against the existing antivirus settings.
5. After scanning has occurred, the result is then handled by the scan manager.

The Kaspersky Lab scan engine supports regular file scanning and script file scanning. With regular file scanning, the input object is a regular file. The engine matches the input content with all possible signatures. With script file scanning, the input object is a script file. It can be JavaScript, VBScript, mIRC script, bat scripts (DOS bat files), and other text scripts. The engine matches the input content only with signatures for script files. Script scanning is only applicable for HTML content over the HTTP protocol. There are two criteria for this scan type. First, the content-type field of this HTML document must be text or HTML. Second, there is no content encoding in the HTTP header. If those two criteria are met, an HTML parser is used to parse the HTML document for scripts.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scan Level Settings on page 5678
- Example: Configuring Full Antivirus Scan Settings at Different Levels on page 5679
- Understanding Full Antivirus Scan Mode Support on page 5676

Understanding Full Antivirus Scan Mode Support

The Kaspersky Lab scan engine supports two modes of scanning:

- scan-all—This option tells the scan engine to scan all the data it receives.
- scan-by-extension—This option bases all scanning decisions on the file extensions found in the traffic in question.

When scanning content, you can use a file extension list to define a set of file extensions that are used in file extension scan mode (scan-by-extension). The antivirus module can then scan files with extensions on the scan-extension list. If an extension is not defined in an extension list, the file with that extension is not scanned in scan-by-extension mode. If there is no extension present, the file in question is scanned.

When using a file extension list to scan content, please note the following requirements:

- File extension entries are case-insensitive.
- The maximum length of the file extension list name is 29 bytes.
- The maximum length of each file extension entry is 15 bytes.
- The maximum entry number in a file extension list is 255.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding the Full Antivirus Internal Scan Engine on page 5675
- Understanding Full Antivirus Scan Level Settings on page 5678
- Example: Configuring Full Antivirus Scan Settings at Different Levels on page 5679

Copyright © 2016, Juniper Networks, Inc.

Configuring Full Antivirus File Extension Scanning (CLI Procedure)

To configure file-extension scanning, use the following CLI configuration statements:

```
security utm {
  custom-objects {
    filename-extension { ; set of list
      name extension-list-name; #mandatory
        value windows-extension-string;
      }
    }
  }

security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  scan-options {
    scan-extension ext-list
    }
  }
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Example: Configuring Full Antivirus File Extension Scanning on page 5677
- Understanding Full Antivirus Scan Mode Support on page 5676

Example: Configuring Full Antivirus File Extension Scanning

This example shows how to configure full antivirus file extension scanning.

- Requirements on page 5677
- Overview on page 5677
- Configuration on page 5678
- Verification on page 5678

Requirements

Before you begin, decide the mode of scanning you require. See "Understanding Full Antivirus Scan Mode Support" on page 5676.

Overview

In this example, you perform the following tasks:

1. Create a file called extlist1 for the kasprof1 profile, and add extensions such as .zip, .js, and .vbs to the extlist1.
2. Configure the scan mode setting. You can choose to scan all files or to scan only the files that have the extensions that you specify. This example uses the scan by-extension option to configure the device to use the extlist1 file.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To configure full antivirus file extension scanning:

1. Create a extension for the list and add extensions to the filename extension list.

   ```
   [edit]
   user@host# set security utm custom-objects filename-extension extlist1 value [zip
     js vbs]
   ```

2. Configure scan extension settings.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
     kasprof1 scan-options scan-extension extlist1
   ```

3. Configure the scan mode setting.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
     kasprof1 scan-options scan-mode by-extension
   ```

4. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus File Extension Scanning (CLI Procedure) on page 5677

Understanding Full Antivirus Scan Level Settings

The antivirus module allows you to configure scanning options on a global level, on a UTM profile level, or on a firewall policy level. Each configuration level has the following implications:

- Global antivirus settings—Settings are applied to all antivirus sessions. Global settings are general overall configurations for the antivirus module or settings that are not specific for profiles.
- Profile-based settings—Antivirus settings are different for different protocols within the same policy.
- Policy-based settings—Antivirus settings are different for different policies. Policy-based antivirus settings are applied to all scan-specified traffic defined in a firewall policy.

The majority of antivirus settings are configured within an antivirus profile, bound to specified protocols, and used by designated policies. These UTM policies are then applied to the traffic according to firewall policies. If a firewall policy with an antivirus setting matches the properties of a traffic flow, the antivirus setting is applied to the traffic

Copyright © 2016, Juniper Networks, Inc.

session. Therefore, you can apply different antivirus settings for different protocols and for different traffic sessions.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding the Full Antivirus Internal Scan Engine on page 5675
- Example: Configuring Full Antivirus Scan Settings at Different Levels on page 5679
- Understanding Full Antivirus Scan Mode Support on page 5676

Example: Configuring Full Antivirus Scan Settings at Different Levels

This example shows how to configure full antivirus scan settings at different levels.

- Requirements on page 5679
- Overview on page 5679
- Configuration on page 5679
- Verification on page 5680

Requirements

Before you begin, decide the type of scanning option you require. See "Understanding Full Antivirus Scan Level Settings" on page 5678.

Overview

In this example, you define antivirus scanning options on any of the following levels:

- Global level
- UTM profile level using the kasprof1 UTM profile
- Firewall policy level using the p1 UTM policy

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile anti-virus kaspersky-lab-engine pattern-update interval
    20
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options default block
set utm-policy p1 anti-virus http-profile av-profile
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure antivirus scanning options at different levels:

1. Configure scanning options at the global level.

```
[edit security utm]
user@host# set feature-profile anti-virus kaspersky-lab-engine pattern-update
  interval 20
```

2. Configure scanning options at the UTM profile level.

```
[edit security utm]
user@host# set feature-profile anti-virus kaspersky-lab-engine profile kasprof1
  fallback-options default block
```

3. Configure scanning options at the UTM policy level.

```
[edit security utm]
user@host# set utm-policy p1 anti-virus http-profile av-profile
```

Results

From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
user@host# show security utm
...
  utm-policy p1 {
  anti-virus {
    http-profile av-profile
      ftp {
        upload-profile av-profile
        download-profile av-profile
      }
    }
  }
...
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Scan Settings at Different Levels

Purpose

Verify the scan settings at different levels.

Action

From operational mode, enter the **show configuration security utm** command.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding the Full Antivirus Internal Scan Engine on page 5675
- Understanding Full Antivirus Scan Mode Support on page 5676

Understanding Full Antivirus Intelligent Prescreening

By default, intelligent prescreening is enabled to improve antivirus scanning performance. The antivirus module generally begins to scan data after the gateway device has received all the packets of a file. Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if it finds that it is unlikely that the file is infected, it then decides that it is safe to bypass the normal scanning procedure.

NOTE: **Intelligent prescreening is only intended for use with non-encoded traffic. It is not applicable for MIME encoded traffic, mail protocols (SMTP, POP3, IMAP) and HTTP POST.**

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Example: Configuring Full Antivirus Intelligent Prescreening on page 5681
- Understanding Full Antivirus Scan Mode Support on page 5676

Example: Configuring Full Antivirus Intelligent Prescreening

This example shows how to configure full antivirus intelligent prescreening. By default, intelligent prescreening is enabled to improve antivirus scanning performance.

- Requirements on page 5681
- Overview on page 5681
- Configuration on page 5682
- Verification on page 5682

Requirements

Before you begin, understand how intelligent prescreening enables the improvement of antivirus scanning performance. See "Understanding Full Antivirus Intelligent Prescreening" on page 5681.

Overview

In this example, you perform the following tasks:

- Enable intelligent prescreening for the kasprof1 profile.
- Disable intelligent prescreening for the kasprof1 profile.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To enable or disable full antivirus intelligent prescreening:

1. Enable intelligent prescreening for the kasprof1 profile.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
     kasprof1 scan-options intelligent-prescreening
   ```

2. Disable intelligent prescreening for the kasprof1 profile.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus kaspersky-lab-engine profile
     kasprof1 scan-options no-intelligent-prescreening
   ```

   NOTE: Intelligent prescreening is intended only for use with non-encoded traffic. It is not applicable to mail protocols (SMTP, POP3, IMAP) or HTTP POST.

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scan Level Settings on page 5678

Understanding Full Antivirus Content Size Limits

Due to resource constraints, there is a default, device-dependent limit on maximum content size for the database. The content size value is configurable. There is also a lower and upper limit for maximum content size. (This range is device dependent and is not configurable.)

The content size check occurs before the scan request is sent. The exact timing of this is protocol dependent. If the protocol header contains an accurate content length field, the content size check takes place when the content length field is extracted during header parsing. The content size usually refers to file size. If there is no content length field, the size is checked while the antivirus module is receiving packets. The content size, in this case, refers to accumulated TCP payload size.

NOTE: This setting can be used in all protocols.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Content Size Limits (CLI Procedure) on page 5683

Configuring Full Antivirus Content Size Limits (CLI Procedure)

To configure content size limits, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
    scan-options {
        content-size-limit KB;
    }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Content Size Limits on page 5682

Understanding Full Antivirus Decompression Layer Limits

The decompression layer limit specifies how many layers of nested compressed files and files with internal extractable objects, such as archive files (tar), MS Word and PowerPoint files, the internal antivirus scanner can decompress before it executes the virus scan. For example, if a message contains a compressed .zip file that contains another compressed .zip file, there are two compression layers. Decompressing both files requires a decompress layer setting of 2.

It is worth noting that during the transfer of data, some protocols use content encoding. The antivirus scan engine must decode this layer, which is considered a decompression level, before it scans for viruses.

There are three kinds of compressed data:

- compressed file (zip, rar, gzip)
- encoded data (MIME)
- packaged data (OLE, .CAP, .MSI, .TAR, .EML)

A decompression layer could be a layer of a zipped file or an embedded object in packaged data. The antivirus engine scans each layer before unpacking the next layer, until it either reaches the user-configured decompress limit, reaches the device decompress layer limit, finds a virus or other malware, or decompresses the data completely, whichever comes first.

As the virus signature database becomes larger and the scan algorithms become more sophisticated, the scan engine has the ability to look deeper into the data for embedded malware. As a result, it can uncover more layers of compressed data. The Juniper Networks device's level of security is limited by decompress limit, which is based on the memory allocated to the security service. If a virus is not found within the decompress limit, the user has an option to either pass or drop the data.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This setting can be used in all protocols.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Decompression Layer Limits (CLI Procedure) on page 5684

Configuring Full Antivirus Decompression Layer Limits (CLI Procedure)

To configure decompression layer limits, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  scan-options {
    decompress-layer-limit number
  }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Decompression Layer Limits on page 5683

Understanding Full Antivirus Scanning Timeouts

The scanning timeout value includes the time frame from when the scan request is generated to when the scan result is returned by the scan engine. The time range can be 1 to 1800 seconds. By default, it is 180 seconds.

NOTE: This timeout parameter is used by all supported protocols. Each protocol can have a different timeout value.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Scanning Timeouts (CLI Procedure) on page 5684

Configuring Full Antivirus Scanning Timeouts (CLI Procedure)

To configure scanning timeouts, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  scan-options {
    timeout-value seconds {
    }
  }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scanning Timeouts on page 5684

Copyright © 2016, Juniper Networks, Inc.

Understanding Full Antivirus Scan Session Throttling

In an attempt to consume all available resources and hinder the ability of the scan engine to scan other traffic, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, a session throttle is imposed for antivirus resources, thereby restricting the amount of traffic a single source can consume at one time. The limit is an integer with 100 as the default setting. This integer refers to the maximum allowed sessions from a single source. You may change this default limit, but understand that if this limit is set high, that is comparable to no limit.

Over-limit is a fallback setting for the connection-per-client limit. The default behavior of over-limit is to block sessions. This is a per-policy setting. You can specify different settings for different UTM policies.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Full Antivirus Scan Session Throttling (CLI Procedure) on page 5685

Configuring Full Antivirus Scan Session Throttling (CLI Procedure)

To configure scan session throttling, use the following CLI configuration statements:

```
security utm utm-policy name
  traffic-options {
    sessions-per-client {
    limit number;
      over-limit { log-and-permit | block}
    }
  }
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scan Session Throttling on page 5685

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 242

Configuring Scan Results and Fallback Options

- Understanding Full Antivirus Scan Result Handling on page 5687
- Monitoring Antivirus Scan Engine Status on page 5687
- Monitoring Antivirus Session Status on page 5688
- Monitoring Antivirus Scan Results on page 5689
- Understanding Antivirus Scanning Fallback Options on page 5691
- Example: Configuring Antivirus Scanning Fallback Options on page 5692

Understanding Full Antivirus Scan Result Handling

Different antivirus scan results are handled in different manners. For example, if a scan result is clean, the traffic is forwarded to the receiver. If the scan result is infected, the traffic is dropped. If the scan results in an error, the result handling depends on the cause of the failure and the configuration (fallback settings).

The following is a list of actions based on scan results:

- Scan Result = Pass

  The scan result handling action is to pass the message. In this case, no virus is detected and no error code is returned. Or, an error code is returned, but the fallback option for this error code is set to log-and-permit.

- Scan Result = Block

  The scan result handling action is to block the message. In this case, either a virus is detected or an error code is returned and the fallback option for this error code is BLOCK.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scan Level Settings on page 5678

Monitoring Antivirus Scan Engine Status

Purpose Using the CLI, you can view the following scan engine status items:

Copyright © 2016, Juniper Networks, Inc.

Antivirus license key status

- View license expiration dates.

Scan engine status and settings

- View last action result.
- View default file extension list.

Antivirus pattern update server settings

- View update URL (HTTP or HTTPS-based).
- View update interval.

Antivirus pattern database status

- View auto update status.
- View last result of database loading.
- If the download completes, view database version timestamp virus record number.
- If the download fails, view failure reason.

Action In the CLI, enter the **user@host> show security utm anti-virus status** command.

Example status result:

```
AV Key Expire Date: 03/01/2010 00:00:00
Update Server: http://update.juniper-updates.net/AV/device-name
interval: 60 minutes
auto update status: next update in 12 minutes
last result: new database loaded
AV signature version: 12/21/2008 00:35 GMT, virus records: 154018
Scan Engine Info: last action result: No error(0x00000000)
```

Related Documentation

- Full Antivirus Configuration Overview on page 5648
- Monitoring Antivirus Session Status on page 5688
- Monitoring Antivirus Scan Results on page 5689

Monitoring Antivirus Session Status

Purpose Using the CLI, you can view the following session status items:

Antivirus session status displays a snapshot of current antivirus sessions. It includes

- Maximum supported antivirus session numbers.
- Total allocated antivirus session numbers.
- Total freed antivirus session numbers.
- Current active antivirus session numbers.

Copyright © 2016, Juniper Networks, Inc.

Action In the CLI, enter the **user@host> show security utm session status** command.

Related Documentation

- Full Antivirus Configuration Overview on page 5648
- Monitoring Antivirus Scan Engine Status on page 5687
- Monitoring Antivirus Scan Results on page 5689

Monitoring Antivirus Scan Results

Purpose View statistics for antivirus requests, scan results, and fallback counters.

Scan requests provide

- The total number of scan request forwarded to the engine.
- The number of scan request being pre-windowed.
- The number of scan requests using scan-all mode.
- The number of scan requests using scan-by-extension mode.

Scan code counters provide

- Number of clean files.
- Number of infected files.
- Number of password protected files.
- Number of decompress layers.
- Number of corrupt files.
- When the engine is out of resources.
- When there is an internal error.

Fallback applied status provides either a log-and-permit or block result when the following has occurred

- Scan engine not ready.
- Maximum content size reached.
- Too many requests.
- Password protected file found.
- Decompress layer too large.
- Corrupt file found.
- Timeout occurred.
- Out of resources.
- Other.

Copyright © 2016, Juniper Networks, Inc.

Action To view antivirus scan results using the CLI editor, enter the **user@host> show security utm anti-virus statistics status** command.

To view antivirus scan results using J-Web:

1. Select **Monitor>Security>UTM>Anti-Virus**.

   The following information becomes viewable in the right pane.

   Antivirus license key status

   - View license expiration dates.

   Antivirus pattern update server settings

   - View update URL (HTTP or HTTPS-based).
   - View update interval.

   Antivirus pattern database status

   - View auto update status.
   - View last result of database loading.
   - If the download completes, view database version timestamp virus record number.
   - If the download fails, view failure reason.

   Antivirus statistics provide

   - The number of scan request being pre-windowed.
   - The total number of scan request forwarded to the engine.
   - The number of scan requests using scan-all mode.
   - The number of scan requests using scan-by-extension mode.

   Scan code counters provide

   - Number of clean files.
   - Number of infected files.
   - Number of password protected files.
   - Number of decompress layers.
   - Number of corrupt files.
   - When the engine is out of resources.
   - When there is an internal error.

   Fallback applied status provides either a log-and-permit or block result when the following has occurred

   - Scan engine not ready.
   - Password protected file found.

Copyright © 2016, Juniper Networks, Inc.

- Decompress layer too large.
- Corrupt file found.
- Out of resources.
- Timeout occurred.
- Maximum content size reached.
- Too many requests.
- Other.

2. You can click the **Clear Anti-Virus Statistics** button to clear all current viewable statistics and begin collecting new statistics.

Related Documentation

- Monitoring Antivirus Session Status on page 5688

Understanding Antivirus Scanning Fallback Options

Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager. The following is a list of possible errors:

- Scan engine is not ready (engine-not-ready)

  The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting.

- Corrupt file (corrupt-file)

  Corrupt file is the error returned by the scan engine when engine detects a corrupted file.

- Decompression layer (decompress-layer)

  Decompress layer error is the error returned by the scan engine when the scanned file has too many compression layers.

- Password protected file (password-file)

  Password protected file is the error returned by the scan engine when the scanned file is protected by a password.

- Max content size (content-size)

  If the content size exceeds a set limit, the content is passed or blocked depending on the max-content-size fallback option.

- Too many requests (too-many-requests)

  If the total number of messages received concurrently exceeds the device limits, the content is passed or blocked depending on the too-many-request fallback option. (The allowed request limit is not configurable.)

- Timeout

Copyright © 2016, Juniper Networks, Inc.

Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is aborted and the content is passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option.

- Out of resources (out-of-resources)

  Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. This failure could be returned by either scan engine (as a scan-code) or scan manager. When out-of-resources occurs, scanning is aborted.

- Default

  All the errors other than those in the above list fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors.

The default fallback action for all the error types is log-and-permit.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Example: Configuring Antivirus Scanning Fallback Options on page 5692

Example: Configuring Antivirus Scanning Fallback Options

This example shows how to configure antivirus scanning fallback options.

- Requirements on page 5692
- Overview on page 5692
- Configuration on page 5693
- Verification on page 5694

Requirements

Before you begin, understand the possible error types and the default fallback actions for those error types. See "Understanding Antivirus Scanning Fallback Options" on page 5691.

Overview

In this example, you configure a feature profile called kasprof, and set the fallback scanning options for default, content-size, corrupt-file, decompress-layer, engine-not-ready, out-of-resources, password-file, timeout, too-many-requests, as block.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The command for changing the URL for the pattern database is:

```
[edit]
user@host# edit security utm feature-profile anti-virus kaspersky-lab-engine
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
user@host# set pattern-update url http://.
```

The default URL is http://update.juniper-updates.net/AV/<device-version>. You should not change this URL unless you are experiencing problems with it and have called for support.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile anti-virus type kaspersky-lab-engine
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options content-size block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options corrupt-file block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options decompress-layer block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options default block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options engine-not-ready block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options out-of-resources block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options password-file block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options timeout block
set security utm feature-profile anti-virus kaspersky-lab-engine profile kasprof1
    fallback-options too-many-requests block
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure scanning fallback options:

1. Select and configure the engine type.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus type kaspersky-lab-engine
   ```

2. Create a profile for the Kaspersky Lab engine and configure a list of fallback options as block or log-and-permit.

   ```
   [edit security utm feature-profile anti-virus kaspersky-lab-engine]
   user@host# set profile kasprof1 fallback-options content-size block
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set profile kasprof1 fallback-options corrupt-file block
user@host# set profile kasprof1 fallback-options decompress-layer block
user@host# set profile kasprof1 fallback-options default block
user@host# set profile kasprof1 fallback-options engine-not-ready block
user@host# set profile kasprof1 fallback-options out-of-resources block
user@host# set profile kasprof1 fallback-options password-file block
user@host# set profile kasprof1 fallback-options timeout block
user@host# set profile kasprof1 fallback-options too-many-requests block
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile anti-virus** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host#show security utm feature-profile anti-virus
kaspersky-lab-engine {
   profile kasprof1 {
      fallback-options {
         default block;
         corrupt-file block;
         password-file block;
         decompress-layer block;
         content-size block;
         engine-not-ready block;
         timeout block;
         out-of-resources block;
         too-many-requests block;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Antivirus Scanning Fallback Options

Purpose Verify the antivirus scanning fallback options.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Scan Level Settings on page 5678

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 243

Configuring Application Protocol Scanning

- Understanding Full Antivirus Application Protocol Scanning on page 5695
- Understanding HTTP Scanning on page 5696
- Enabling HTTP Scanning (CLI Procedure) on page 5697
- Understanding FTP Antivirus Scanning on page 5697
- Enabling FTP Antivirus Scanning (CLI Procedure) on page 5698
- Understanding SMTP Antivirus Scanning on page 5698
- Enabling SMTP Antivirus Scanning (CLI Procedure) on page 5700
- Understanding POP3 Antivirus Scanning on page 5700
- Enabling POP3 Antivirus Scanning (CLI Procedure) on page 5702
- Understanding IMAP Antivirus Scanning on page 5702
- Enabling IMAP Antivirus Scanning (CLI Procedure) on page 5704

Understanding Full Antivirus Application Protocol Scanning

You can turn antivirus scanning on and off on a per protocol basis. If scanning for a protocol is disabled in an antivirus profile, there is no application intelligence for this protocol. Therefore, in most cases, traffic using this protocol is not scanned. But if the protocol in question is based on another protocol for which scanning is enabled in an antivirus profile, then the traffic is scanned as that enabled protocol.

The internal antivirus scan engine supports scanning for specific Application Layer transactions allowing you to select the content (HTTP, FTP, SMTP, POP3, or IMAP traffic) to scan. For each content type that you are scanning, you have different configuration options.

Profile-based settings, including enable/disable, scan-mode, and scan result handling settings, may not be applicable to all supported protocols. The following table lists profile-based settings and their protocol support.

Table 495: Supported Profile-based Settings By Protocol

Profile Setting	Protocol Support
Enable or disable scanning on per protocol basis	All protocols support this feature

Copyright © 2016, Juniper Networks, Inc.

Table 495: Supported Profile-based Settings By Protocol *(continued)*

Profile Setting	Protocol Support
"Understanding Full Antivirus Scan Mode Support" on page 5676, including file extension scanning	All protocols support this feature
"Understanding Full Antivirus Content Size Limits" on page 5682	All protocols support this feature
"Understanding Full Antivirus Decompression Layer Limits" on page 5683	All protocols support this feature
"Understanding Full Antivirus Scanning Timeouts" on page 5684	All protocols support this feature
"Understanding HTTP Trickling" on page 5709	HTTP only
"Understanding Antivirus Scanning Fallback Options" on page 5691	All protocols support this feature
Protocol specific messages	All protocols support this feature
"Understanding E-Mail Virus-Detected Notifications" on page 5712	SMTP, POP3, and IMAP only
"Understanding Custom Message Virus-Detected Notifications" on page 5713	All protocols support this feature

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding HTTP Scanning on page 5696
- Enabling HTTP Scanning (CLI Procedure) on page 5697
- Understanding Protocol-Only Virus-Detected Notifications on page 5711

Understanding HTTP Scanning

If antivirus scanning is enabled for Hypertext Transfer Protocol (HTTP) traffic in a content security profile, TCP traffic to defined HTTP service ports (generally port 80) is monitored. For HTTP traffic, the security device scans both HTTP responses and requests (get, post, and put commands).

NOTE: For HTTP antivirus scanning, both HTTP 1.0 and 1.1 are supported. If the protocol version is HTTP 0.x , the antivirus scanner attempts to scan the traffic. Unknown protocols are bypassed. For example, some application protocols use HTTP as the transport but do not comply with HTTP 1.0 or 1.1. These are considered unknown protocols and are not scanned.

This is a general description of how HTTP traffic is intercepted, scanned, and acted upon by the antivirus scanner:

1. An HTTP client sends an HTTP request to a webserver or a webserver responds to an HTTP request.

Copyright © 2016, Juniper Networks, Inc.

2. The security device intercepts the request and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:
   - If there is no virus, the device forwards the request to the webserver.
   - If there is a virus, the device drops the request and sends an HTTP message reporting the infection to the client.

With script-only scanning, the input object is a script file. It can be JavaScript, VBScript, mIRC script, bat scripts (DOS bat files) and other text scripts. The engine matches the input content only with signatures for script files. Script scanning is applicable only for HTML content over the HTTP protocol. There are two criteria for this scan-type. First, the content-type field of this HTML document must be text or HTML. Second, there is no content encoding in the HTTP header. If those two criteria are met, an HTML parser is used to parse the HTML document.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Application Protocol Scanning on page 5695
- Enabling HTTP Scanning (CLI Procedure) on page 5697

Enabling HTTP Scanning (CLI Procedure)

To enable antivirus scanning for HTTP traffic, enter the following CLI configuration statement:

```
user@host# set security utm utm-policy policy-name anti-virus http
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Full Antivirus Application Protocol Scanning on page 5695
- Understanding HTTP Scanning on page 5696

Understanding FTP Antivirus Scanning

If antivirus scanning is enabled for File Transfer Protocol (FTP) traffic in a content security profile, the security device monitors the control channel and, when it detects one of the FTP commands for transferring data, it scans the data sent over the data channel.

This is a general description of how FTP traffic is intercepted, scanned, and acted upon by the antivirus scanner:

1. A local FTP client opens an FTP control channel to an FTP server and requests the transfer of some data.
2. The FTP client and server negotiate a data channel over which the server sends the requested data. The security device intercepts the data and passes it to the antivirus scan engine, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:

Copyright © 2016, Juniper Networks, Inc.

- If there is no virus, the device forwards the data to the client.
- If there is a virus, the device replaces the data with a drop message in the data channel and sends a message reporting the infection in the control channel.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Enabling FTP Antivirus Scanning (CLI Procedure) on page 5698

Enabling FTP Antivirus Scanning (CLI Procedure)

To enable antivirus scanning for File Transfer Protocol (FTP) traffic, enter the following CLI configuration statement:

```
user@host# security utm utm-policy policy-name anti-virus ftp
```

NOTE: In order to scan FTP traffic, the FTP ALG must be enabled.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding FTP Antivirus Scanning on page 5697

Understanding SMTP Antivirus Scanning

If SMTP (Simple Mail Transfer Protocol) antivirus scanning is enabled in a content security profile, the security device redirects traffic from local SMTP clients to the antivirus scanner before sending it to the local mail server.

NOTE: Chunking is an alternative to the data command. It provides a mechanism to transmit a large message in small chunks. It is not supported. Messages using chunking are bypassed and are not scanned.

This is a general description of how SMTP traffic is intercepted, scanned, and acted upon by the antivirus scanner:

1. An SMTP client sends an e-mail message to a local mail server or a remote mail server forwards an e-mail message via SMTP to the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the device follows one of two courses:
   - If there is no virus, the device forwards the message to the local server.
   - If there is a virus, the device sends a replacement message to the client.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Understanding SMTP Antivirus Mail Message Replacement on page 5699
- Understanding SMTP Antivirus Sender Notification on page 5699
- Understanding SMTP Antivirus Subject Tagging on page 5700

Understanding SMTP Antivirus Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
  <filename> with virus <virusname>, so it is dropped.
```

If a scan error is returned and the fail mode is set to drop, the original message is dropped and the entire message body is truncated. The content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> is dropped for <reason>.
```

Understanding SMTP Antivirus Sender Notification

If **notify-sender-on-virus** is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender. The content of the notification may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <ENVID> contaminated file <filename>
  with virus <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If **notify-sender-on-error-drop** is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <ENVID> <reason>.
e-mail Header is:
<header of scanned e-mail>
```

NOTE: For information on the ENVID parameter, refer to RFC 3461.

Copyright © 2016, Juniper Networks, Inc.

Understanding SMTP Antivirus Subject Tagging

If a scan error is returned and the fail mode is set to **pass**, the antivirus module passes the message through to the server. If **notify-recipient-on-error-pass** is set, the following string is appended to the end of the subject field:

(No virus check: <reason>)

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Enabling SMTP Antivirus Scanning (CLI Procedure) on page 5700

Enabling SMTP Antivirus Scanning (CLI Procedure)

To enable antivirus scanning for SMTP traffic, enter the following CLI configuration statement:

user@host# **set security utm utm-policy *policy-name* anti-virus smtp-profile**

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding SMTP Antivirus Scanning on page 5698

Understanding POP3 Antivirus Scanning

If Post Office Protocol 3 (POP3) antivirus scanning is enabled in a content security profile, the security device redirects traffic from a local mail server to antivirus scanner before sending it to the local POP3 client.

This is a general description of how POP3 traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. The POP3 client downloads an e-mail message from the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the security device follows one of two courses:
   - If there is no virus, the device forwards the message to the client.
   - If there is a virus, the device sends a message reporting the infection to the client.

NOTE: See "Understanding Protocol-Only Virus-Detected Notifications" on page 5711 for information on protocol-only notifications for IMAP.

Copyright © 2016, Juniper Networks, Inc.

This topic includes the following sections:

- Understanding POP3 Antivirus Mail Message Replacement on page 5701
- Understanding POP3 Antivirus Sender Notification on page 5701
- Understanding POP3 Antivirus Subject Tagging on page 5701

Understanding POP3 Antivirus Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
  <filename> with virus <virusname>, so it is dropped.
```

Understanding POP3 Antivirus Sender Notification

If **notify-sender-on-virus** is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender.

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> contaminated file <filename> with virus
  <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If **notify-sender-on-error-drop** is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <reason>.
e-mail Header is:
<header of scanned e-mail>
```

Understanding POP3 Antivirus Subject Tagging

If a scan error is returned and the fail mode is set to **pass**, the antivirus module passes the message through to the server. If **notify-recipient-on-error-pass** is set, the following string is appended to the end of subject field:

```
(No virus check: <reason>)
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Enabling POP3 Antivirus Scanning (CLI Procedure) on page 5702

Copyright © 2016, Juniper Networks, Inc.

Enabling POP3 Antivirus Scanning (CLI Procedure)

To enable antivirus scanning for POP3 traffic, enter the following CLI configuration statement:

```
user@host# set security utm utm-policy policy-name anti-virus pop3-profile
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding POP3 Antivirus Scanning on page 5700

Understanding IMAP Antivirus Scanning

If IMAP (Internet Message Access Protocol) antivirus scanning is enabled in a content security profile, the security device redirects traffic from a local mail server to the internal antivirus scanner before sending it to the local IMAP client.

This is a general description of how IMAP traffic is intercepted, scanned, and acted upon by the antivirus scanner.

1. The IMAP client downloads an e-mail message from the local mail server.
2. The security device intercepts the e-mail message and passes the data to the antivirus scanner, which scans it for viruses.
3. After completing the scan, the security device follows one of two courses:
   - If there is no virus, the device forwards the message to the client.
   - If there is a virus, the device sends a message reporting the infection to the client.

NOTE: See "Understanding Protocol-Only Virus-Detected Notifications" on page 5711 for information on protocol-only notifications for IMAP.

This topic includes the following sections:

- Understanding IMAP Antivirus Mail Message Replacement on page 5702
- Understanding IMAP Antivirus Sender Notification on page 5703
- Understanding IMAP Antivirus Subject Tagging on page 5703
- Understanding IMAP Antivirus Scanning Limitations on page 5703

Understanding IMAP Antivirus Mail Message Replacement

If the antivirus scanner finds a virus in an e-mail message, the original message is dropped, the message body is truncated, and the content is replaced by a message that may appear as follows:

```
nContent-Type: text/plain
Your mail <src_ip> : <src_port> — <dst_port>: <dst_port> contains contaminated file
  <filename> with virus <virusname>, so it is dropped.
```


Copyright © 2016, Juniper Networks, Inc.

Understanding IMAP Antivirus Sender Notification

If **notify-sender-on-virus** is set and the message is dropped due to a detected virus, an e-mail is sent to the mail sender.

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> contaminated file <filename> with virus
  <virusname>.
e-mail Header is:
<header of scanned e-mail>
```

If **notify-sender-on-error-drop** is set and the message is dropped due to a scan error, an e-mail is sent to the mail sender of the scanned message. The content of the e-mail may appear as follows:

```
From: <admin>@<gateway_ip>
To: <sender_e-mail>
Subject: Mail Delivery Failure
This message is created automatically by mail delivery software. A message that you sent
  could not be delivered to one or more of its recipients for the reason:
<src_ip> : <src_port> — <dst_port>: <dst_port> <reason>.
e-mail Header is:
<header of scanned e-mail>
```

Understanding IMAP Antivirus Subject Tagging

If a scan error is returned and the fail mode is set to **pass**, the antivirus module passes the message through to the server. If **notify-recipient-on-error-pass** is set, the following string is appended to the end of subject field:

```
(No virus check: <reason>)
```

Understanding IMAP Antivirus Scanning Limitations

Mail Fragments — It is possible to chop one e-mail into multiple parts and to send each part through a different response. This is called mail fragmenting and most popular mail clients support it in order to send and receive large e-mails. Scanning of mail fragments is not supported by the antivirus scanner and in such cases, the message body is not scanned.

Partial Content — Some mail clients treat e-mail of different sizes differently. For example, small e-mails (less than 10 KB) are downloaded as a whole. Large e-mails (for example, less than 1 MB) are chopped into 10 KB pieces upon request from the IMAP server. Scanning of any partial content requests is not supported by the antivirus scanner.

IMAP Uploads — Only antivirus scanning of IMAP downloads is supported. IMAP upload traffic is not scanned.

Related Documentation

- Full Antivirus Protection Overview on page 5647

Copyright © 2016, Juniper Networks, Inc.

- Enabling IMAP Antivirus Scanning (CLI Procedure) on page 5704

Enabling IMAP Antivirus Scanning (CLI Procedure)

To enable antivirus scanning for IMAP traffic, enter the following CLI configuration statement:

```
user@host# security utm utm-policy policy-name anti-virus imap-profile
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding IMAP Antivirus Scanning on page 5702

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 244

Configuring Whitelists

- Understanding MIME Whitelists on page 5705
- Example: Configuring MIME Whitelists to Bypass Antivirus Scanning on page 5706
- Understanding URL Whitelists on page 5706
- Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure) on page 5707

Understanding MIME Whitelists

The gateway device uses MIME (Multipurpose Internet Mail Extension) types to decide which traffic may bypass antivirus scanning. The MIME whitelist defines a list of MIME types and can contain one or many MIME entries.

A MIME entry is case-insensitive. An empty MIME is an invalid entry and should never appear in the MIME list. If the MIME entry ends with a / character, prefix matching takes place. Otherwise, exact matching occurs.

There are two types of MIME lists used to configure MIME type antivirus scan bypassing:

- mime-whitelist list—This is the comprehensive list for those MIME types that can bypass antivirus scanning.
- exception list—The exception list is a list for excluding some MIME types from the mime-whitelist list. This list is a subset of MIME types found in the mime-whitelist.

  For example, if the mime-whitelist includes the entry,**video/** and the exception list includes the entry **video/x-shockwave-flash**, by using these two lists, you can bypass objects with “video/” MIME type but not bypass “video/x-shockwave-flash” MIME type.

  You should note that there are limits for mime-whitelist entries as follows:

  - The maximum number of MIME items in a MIME list is 50.
  - The maximum length of each MIME entry is restricted to 40 bytes.
  - The maximum length of a MIME list name string is restricted to 40 bytes.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Example: Configuring MIME Whitelists to Bypass Antivirus Scanning on page 5706
- Understanding URL Whitelists on page 5706

Copyright © 2016, Juniper Networks, Inc.

- Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure) on page 5707

Example: Configuring MIME Whitelists to Bypass Antivirus Scanning

This example shows how to configure MIME whitelists to bypass antivirus scanning.

- Requirements on page 5706
- Overview on page 5706
- Configuration on page 5706
- Verification on page 5706

Requirements

Before you begin, decide the type of MIME lists used to configure MIME type antivirus scan bypassing. See "Understanding MIME Whitelists" on page 5705.

Overview

In this example, you create MIME lists called avmime2 and ex-avmime2 and add patterns to them.

Configuration

Step-by-Step Procedure

To configure MIME whitelists to bypass antivirus scanning:

1. Create MIME lists and add patterns to the lists.

```
[edit]
user@host# set security utm custom-objects mime-pattern avmime2 value
  [video/quicktime image/x-portable-anymap x-world/x-vrml]
user@host# set security utm custom-objects mime-pattern ex-avmime2 value
  [video/quicktime-inappropriate]
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding URL Whitelists on page 5706
- Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure) on page 5707

Understanding URL Whitelists

A URL whitelist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for scanning. Because antivirus

Copyright © 2016, Juniper Networks, Inc.

scanning is CPU and memory intensive action, if there are URLs or IP addresses that you are sure do not require scanning, you might want to create this custom list and add them to it.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding MIME Whitelists on page 5705
- Example: Configuring MIME Whitelists to Bypass Antivirus Scanning on page 5706
- Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure) on page 5707

Configuring URL Whitelists to Bypass Antivirus Scanning (CLI Procedure)

To configure URL whitelists, use the following CLI configuration statements:

```
security utm custom-objects {
  custom-url-category { ; set of list
    name url-category-name; #mandatory
    value url-pattern-name;
  }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding MIME Whitelists on page 5705
- Example: Configuring MIME Whitelists to Bypass Antivirus Scanning on page 5706
- Understanding URL Whitelists on page 5706

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 245

Configuring HTTP Trickling

- Understanding HTTP Trickling on page 5709
- Configuring HTTP Trickling to Prevent Timeouts During Antivirus Scanning (CLI Procedure) on page 5709

Understanding HTTP Trickling

HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. On some slow link transferring, a large file could timeout if too much time is taken for the antivirus scanner to scan a complex file.

HTTP trickling is the forwarding of specified amounts of unscanned HTTP traffic to the requesting HTTP client to prevent the browser window from timing out while the scan manager examines downloaded HTTP files. (The security device forwards small amounts of data in advance of transferring an entire scanned file.)

HTTP Trickling is time-based and there is only one parameter, the time-out interval, to configure for this feature. By default, trickling is disabled.

NOTE: The timeout based trickling is packet driven. This means, if no packet is received within a certain time frame, HTTP trickling is discontinued. This setting is only supported for HTTP connections.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring HTTP Trickling to Prevent Timeouts During Antivirus Scanning (CLI Procedure) on page 5709

Configuring HTTP Trickling to Prevent Timeouts During Antivirus Scanning (CLI Procedure)

To configure HTTP trickling, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine {
    profile name {
        trickling timeout seconds;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding HTTP Trickling on page 5709

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 246

Configuring Notifications

- Understanding Protocol-Only Virus-Detected Notifications on page 5711
- Configuring Protocol-Only Virus-Detected Notifications (CLI Procedure) on page 5711
- Understanding E-Mail Virus-Detected Notifications on page 5712
- Configuring E-Mail Virus-Detected Notifications (CLI Procedure) on page 5712
- Understanding Custom Message Virus-Detected Notifications on page 5713
- Configuring Custom Message Virus-Detected Notifications (CLI Procedure) on page 5713

Understanding Protocol-Only Virus-Detected Notifications

When content is blocked because a virus is found or a scan error occurs, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code may be returned to the client. This way, the client determines that a virus was detected rather than interpreting that a file transfer succeeded.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Protocol-Only Virus-Detected Notifications (CLI Procedure) on page 5711

Configuring Protocol-Only Virus-Detected Notifications (CLI Procedure)

To configure protocol-only virus-detected notifications, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  notification-options {
    virus-detection {
      type { protocol-only | message }
    }
    fallback-block {
      type { protocol-only | message }
        }
    }
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

Related Documentation
- Full Antivirus Protection Overview on page 5647
- Understanding Protocol-Only Virus-Detected Notifications on page 5711

Understanding E-Mail Virus-Detected Notifications

For mail protocols (SMTP, POP3, IMAP), e-mail notification is used to notify the sender or the recipient about the detected viruses or the scanning errors. There are three settings for e-mail notifications:

- virus-detection/notify-mail-sender — This setting is used when a virus is detected. If it is enabled, an e-mail is sent to the sender upon virus detection.
- fallback-block/notify-mail-sender — This setting is used when other scan codes or scanning errors are returned and the message is dropped. If it is enabled, an e-mail is sent to the sender when an error code is returned.
- fallback-non-block/notify-mail-recipient — This setting is used when other scan codes or scanning errors are returned and the message is passed. If it is enabled, the e-mail sent to the recipient is tagged when an error code is returned.

Related Documentation
- Full Antivirus Protection Overview on page 5647
- Configuring E-Mail Virus-Detected Notifications (CLI Procedure) on page 5712

Configuring E-Mail Virus-Detected Notifications (CLI Procedure)

To configure the system to send e-mail notifications when viruses are detected, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  notification-options {
    virus-detection {
      notify-mail-sender
      }
  fallback-block {
    notify-mail-sender
    }
  fallback-non-block {
  notify-mail-recipient
  }
  }
}
}
```

Related Documentation
- Full Antivirus Protection Overview on page 5647
- Understanding E-Mail Virus-Detected Notifications on page 5712

Copyright © 2016, Juniper Networks, Inc.

Understanding Custom Message Virus-Detected Notifications

Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. When using custom messages, you can provide a customized message in the message content you can define customized subject tags.

NOTE: Custom-message in fallback-nonblock is used only by mail protocols.

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Configuring Custom Message Virus-Detected Notifications (CLI Procedure) on page 5713

Configuring Custom Message Virus-Detected Notifications (CLI Procedure)

To configure the system to send custom messages when viruses are detected, use the following CLI configuration statements:

```
security utm feature-profile anti-virus kaspersky-lab-engine profile name {
  notification-options {
    virus-detection {
      custom-message msg
      custom-message-subject subject-msg
    }
  fallback-block {
    custom-message msg
    custom-message-subject subject-msg
  }
  fallback-non-block {
    custom-message msg
    custom-message-subject subject-msg
    }
    }
  }
}
```

Related Documentation

- Full Antivirus Protection Overview on page 5647
- Understanding Custom Message Virus-Detected Notifications on page 5713

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 75

Configuring and Managing Sophos Antivirus Protection

- Configuring Sophos Antivirus Protection on page 5717

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 247

Configuring Sophos Antivirus Protection

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Features on page 5718
- Understanding Sophos Antivirus Data File Update on page 5719
- Comparison of Sophos Antivirus to Kaspersky Antivirus on page 5720
- Sophos Antivirus Configuration Overview on page 5720
- Example: Configuring Sophos Antivirus Custom Objects on page 5721
- Example: Configuring Sophos Antivirus Feature Profile on page 5724
- Example: Configuring Sophos Antivirus UTM Policies on page 5731
- Example: Configuring Sophos Antivirus Firewall Security Policies on page 5732
- Example: Configuring Sophos Antivirus Scanner with SSL Forward Proxy on page 5734
- Managing Sophos Antivirus Data Files on page 5740

Sophos Antivirus Protection Overview

Sophos antivirus scanning is offered as a less CPU-intensive alternative to the full file-based antivirus feature. Sophos supports the same protocols as full antivirus and functions in much the same manner; however, it has a smaller memory footprint and is compatible with lower end devices that have less memory.

Sophos antivirus is as an in-the-cloud antivirus solution. The virus pattern and malware database is located on external servers maintained by Sophos (Sophos Extensible List) servers, thus there is no need to download and maintain large pattern databases on the Juniper device. The Sophos antivirus scanner also uses a local internal cache to maintain query responses from the external list server to improve lookup performance.

Because a significant amount of traffic processed by Juniper Unified Threat Management (UTM) is HTTP based, Uniform Resource Identifier (URI) checking is used to effectively prevent malicious content from reaching the endpoint client or server. The following checks are performed for HTTP traffic: URI lookup, true file type detection, and file checksum lookup. The following application layer protocols are supported: HTTP, FTP, SMTP, POP3 and IMAP.

Related Documentation

- Sophos Antivirus Features on page 5718
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

Sophos Antivirus Features

Sophos antivirus has the following main features:

- **Sophos antivirus expanded MIME decoding support**—Sophos antivirus offers decoding support for HTTP, POP3, SMTP, and IMAP. MIME decoding support includes the following for each supported protocol:
  - Multipart and nested header decoding
  - Base64 decoding, printed quote decoding, and encoded word decoding in the subject field
- **Sophos antivirus supports HTTPS traffic**—Sophos antivirus over SSL forward proxy supports HTTPS traffic by intercepting HTTPS traffic passing through the SRX Series device. The security channel from the SRX Series device is divided as one SSL channel between the client and the SRX Series device and another SSL channel between the SRX Series device and the HTTPS server. SSL forward proxy acts as the terminal for both channels and forwards the cleartext traffic to UTM. UTM extracts the URL and the file checksum information from cleartext traffic. The Sophos antivirus scanner determines whether to block or permit the requests.

  SSL forward proxy does not support client authentication. If client authentication is required by the server, UTM bypasses the traffic. UTM bypasses the HTTPS traffic under the following conditions:
  - If SSL proxy does not parse the first handshake packet from the client, SSL forward proxy bypasses the traffic.
  - If the SSL proxy handshake with the client and server is incomplete because of compatibility issues, connection drops.
  - If the system resource is low, SSL forward proxy cannot handle the new connection and Sophos antivirus bypasses the traffic.
  - If HTTPS traffic hits the whitelist of SSL forward proxy, SSL forward proxy and Sophos antivirus bypass the traffic.
- **Sophos antivirus scan result handling**—With Sophos antivirus, the TCP, traffic is closed gracefully when a virus is found and the data content is dropped.

  The following fail mode options are supported: content-size, default, engine-not-ready, out-of-resource, timeout, and too-many-requests. You can set the following actions: block, log-and-permit, and permit. Fail mode handling of supported options with Sophos is much the same as with full antivirus.
- **Sophos Uniform Resource Identifier checking**—Sophos provides Uniform Resource Identifier (URI) checking, which is similar to antispam realtime blackhole list (RBL) lookups. URI checking is a way of analyzing URI content in HTTP traffic against the Sophos database to identify malware or malicious content. Because malware is predominantly static, a checksum mechanism is used to identify malware to improve performance. Files that are capable of using a checksum include .exe, .zip, .rar, .swf, .pdf, and .ole2 (doc and xls).

Copyright © 2016, Juniper Networks, Inc.

NOTE: If you have a Juniper Networks device protecting an internal network that has no HTTP traffic, or has webservers that are not accessible to the outside world, you might want to turn off URI checking. If the webservers are not accessible to the outside world, it is unlikely that they contain URI information that is in the Sophos URI database. URI checking is on by default.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720
- Example: Configuring Sophos Antivirus Feature Profile on page 5724

Understanding Sophos Antivirus Data File Update

Sophos antivirus uses a small set of data files that need to be updated periodically. These data files only contain information on guiding scanning logic and do not contain the full pattern database. The main pattern database, which includes protection against critical viruses, URI checks, malware, worms, Trojans, and spyware, is located on remote Sophos Extensible List servers maintained by Sophos.

The Sophos data files are updated over HTTP or HTTPS and can be updated manually or scheduled to update automatically. This is similar functionality to that found in the Full Antivirus solution, but with some minor differences:

- With Sophos antivirus, the signature database auto-update interval is once a day by default. This interval can be changed.
- With Sophos, there is no interruption in virus scanning capability during the data file update. If the update fails, the existing data files will continue to be used.
- By default, the URL for Sophos antivirus data file update is http://update.juniper-updates.net/SAV/.

NOTE: The Sophos antivirus scanning feature is a separately licensed subscription service. When your antivirus license key expires, functionality will no longer work because the pattern lookup database is located on remote Sophos servers. You have a 30-day grace period in which to update your license.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Managing Sophos Antivirus Data Files on page 5740
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

Comparison of Sophos Antivirus to Kaspersky Antivirus

Sophos Antivirus is much like Juniper Express Antivirus and also has similarities to the Full Antivirus feature:

- Unlike the Juniper Express and Full Antivirus solutions, the antivirus and malware database for Sophos is stored on a group of remote Sophos Extensible List servers. Queries are performed using the DNS protocol. Sophos maintains these servers, so there is no need to download and maintain large pattern databases on the Juniper device. Because the database is remote, there is no size limitation and there is a quicker response to new virus outbreaks.

  NOTE: **Sophos antivirus uses a set of data files that need to be updated on a regular basis. These are not typical virus pattern files; they are a set of small files that help guide virus scanning logic. You can manually download the data files or set up automatic download.**

- Sophos does not provide the same prescreening detection as Kaspersky Antivirus. Sophos does provide a similar solution that is part of the Sophos engine and cannot be turned on and off.
- The Sophos antivirus scanning feature is a separately licensed subscription service. Also, the pattern lookup database is located on remote servers maintained by Sophos, so when your antivirus license key expires, functionality will no longer work. You have a 30-day grace period in which to update your license.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720

Sophos Antivirus Configuration Overview

Sophos antivirus is part of the Unified Threat Management (UTM) feature set, so you first configure UTM options (custom objects), configure the Sophos Feature, then create a UTM policy and a security policy. The security policy controls all traffic that is forwarded by the device, and the UTM policy specifies which parameters to use to scan traffic. The UTM policy is also used to bind a set of protocols to one or more UTM feature profiles, including Sophos antivirus in this case.

You must complete the following tasks to configure Sophos antivirus:

1. Configure UTM custom objects and MIME lists. See "Example: Configuring Sophos Antivirus Custom Objects" on page 5721,
2. Configure the Sophos antivirus feature profile. See "Example: Configuring Sophos Antivirus Feature Profile" on page 5724.

Copyright © 2016, Juniper Networks, Inc.

3. Configure a UTM policy. See "Example: Configuring Sophos Antivirus UTM Policies" on page 5731
4. Configure a security policy. See "Example: Configuring Sophos Antivirus Firewall Security Policies" on page 5732.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717

Example: Configuring Sophos Antivirus Custom Objects

This example shows you how to create UTM global custom objects to be used with Sophos antivirus.

- Requirements on page 5721
- Overview on page 5721
- Configuration on page 5722
- Verification on page 5724

Requirements

Before you begin, read about UTM custom objects. See "Understanding UTM Custom Objects" on page 5585.

Overview

Configure MIME lists. This includes creating a MIME whitelist and a MIME exception list for antivirus scanning. In this example, you bypass scanning of QuickTime videos, unless if they contain the MIME type quicktime-inappropriate.

WARNING: When you configure the MIME whitelist feature, be aware that, because header information in HTTP traffic can be spoofed, you cannot always trust HTTP headers to be legitimate. When a Web browser is determining the appropriate action for a given file type, it detects the file type without checking the MIME header contents. However, the MIME whitelist feature does refer to the MIME encoding in the HTTP header. For these reasons, it is possible in certain cases for a malicious website to provide an invalid HTTP header. For example, a network administrator might inadvertently add a malicious website to a MIME whitelist, and, because the site is in the whitelist, it will not be blocked by Sophos even though Sophos has identified the site as malicious in its database. Internal hosts would then be able to reach this site and could become infected.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To configure a MIME list:

1. Click the **Configure** tab from the taskbar, and then select **Security>UTM>Custom Objects**.
2. Click the **MIME Pattern List** tab and then click **Add**.
3. In the MIME Pattern Name box, type **avmime2**.
4. In the MIME Pattern Value box, type **video/quicktime**, and click **Add**.
5. In the MIME Pattern Value box, type **image/x-portable-anympa**, and click **Add**.
6. In the MIME Pattern Value box, type **x-world/x-vrml**, and click **Add**.

To configure a MIME exception list:

1. Click the **Configure** tab from the taskbar, and then select **Security>UTM>Custom Objects**.
2. Click the **MIME Pattern List** tab and then select **Add**.
3. In the MIME Pattern Name box, type **exception-avmime2**.
4. In the MIME Pattern Value box, type **video/quicktime-inappropriate** and click **Add**.

Configure a URL pattern list (whitelist) of URLs or addresses that will be bypassed by antivirus scanning. After you create the URL pattern list, you will create a custom URL category list and add the pattern list to it.

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

To configure a URL pattern whitelist:

1. Click the **Configure** tab from the taskbar, and then select **Security>UTM>Custom Objects**.
2. Click the **URL Pattern List** tab, and then click **Add**.
3. In the URL Pattern Name box, enter **urlist2**.
4. In the URL Pattern Value box, enter **http://example.net**. (You can also us the IP address of the server instead of the URL.)

Save your configuration:

1. Click **OK** to check your configuration and save it as a candidate configuration.
2. If you are done configuring the device, click **Actions>Commit**.

Copyright © 2016, Juniper Networks, Inc.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can use "*" only if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax is supported: http://*.example.net, http://www.example.ne?, http://www.example.n??. The following wildcard syntax is not supported: *.example.net , www.example.ne?, http://*example.net, http://*.

Step-by-Step Procedure

To configure antivirus protection using the CLI, you must create your custom objects in the following order:

1. Create the MIME whitelist.

```
[edit security utm]
user@host# set custom-objects mime-pattern avmime2 value [video/quicktime
   image/x-portable-anymap x-world/x-vrml]
```

Create the MIME exception list.

```
[edit security utm]
user@host# set custom-objects mime-pattern exception-avmime2 value
   [video/quicktime-inappropriate]
```

2. Configure a URL pattern list (whitelist) of URLs or addresses that you want to bypass. After you create the URL pattern list, you create a custom URL category list and add the pattern list to it. Configure a URL pattern list custom object by creating the list name and adding values to it as follows.

NOTE: Because you use URL pattern lists to create custom URL category lists, you must configure URL pattern list custom objects before you configure custom URL category lists.

```
[edit security utm]
user@host# set custom-objects url-pattern urllist2 value [http://www.example.net
   192.168.1.5]
```

Copyright © 2016, Juniper Networks, Inc.

NOTE: URL pattern wildcard support—The wildcard rule is as follows: *\.[]\?* and you must precede all wildcard URLs with http://. You can only use "*" if it is at the beginning of the URL and is followed by a ".". You can only use "?" at the end of the URL.

The following wildcard syntax is supported: http://*.example.net, http://www.example.ne?, http://www.example.n??. The following wildcard syntax is not supported: *.example.net , www.example.ne?, http://*example.net, http://*.

3. Configure a custom URL category list custom object by using the URL pattern list urllist2 that you created earlier:

```
[edit security utm]
user@host# set custom-objects custom-url-category custurl2 value urllist2
```

Verification

To verify the configuration, enter the **show security utm custom-objects** command.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720
- Example: Configuring Sophos Antivirus Feature Profile on page 5724
- Understanding UTM Custom Objects on page 5585

Example: Configuring Sophos Antivirus Feature Profile

This example shows you how to configure a Sophos antivirus profile that defines the parameters that will be used for virus scanning.

- Requirements on page 5724
- Overview on page 5725
- Configuration on page 5725
- Verification on page 5730

Requirements

Before you begin:

- Install a Sophos antivirus license. See the *Installation and Upgrade Guide for Security Devices*.
- Configure custom objects for UTM. See "Example: Configuring Sophos Antivirus Custom Objects" on page 5721.

Copyright © 2016, Juniper Networks, Inc.

Overview

The following configuration defines Sophos as the antivirus engine and sets parameters, such as the data file update interval, notification options for administrators, fallback options, and file size limits.

Configuration

GUI Step-by-Step Procedure

NOTE: The following example shows you how to create a custom Sophos profile. If you want to use the Juniper Networks preconfigured profile, use the profile named junos-sophos-av-defaults in your UTM policy. See "Example: Configuring Sophos Antivirus UTM Policies" on page 5731.

1. Select and configure the engine type. Because you are configuring Sophos antivirus, you configure sophos-engine:
   a. Click the **Configure** tab from the taskbar, and then select **Security>UTM>Anti-Virus**.
   b. Click the **Global Options** tab and then click **Sophos**.
   c. Click **OK** and commit your changes.
   d. Restart the device to enable Sophos as the antivirus engine.
2. Return to the antivirus Global Options screen as you did in step 1, and set the following parameters:
   a. In the MIME whitelist list, select **exception-avmime2**.
   b. In the URL whitelist list, select **custurl2**.
   c. In the Pattern update interval (sec) box, type **2880**.
   d. In the box, type the e-mail address that will receive SophosAdmin e-mail data file update notifications. For example - admin@example.net.
   e. In the Custom message subject box, type **Sophos Data File Updated**.
   f. Click **OK** to check your configuration and save it as a candidate configuration.
3. Configure a profile for the sophos-engine and set parameters.
   a. Click the **Configure** tab from the taskbar and then select **Security>UTM>Anti-Virus**. Click **Add**.
   b. In the Add profile box, click the **Main** tab.
   c. In the Profile name box, type **sophos-prof1**.
   d. In the Trickling timeout box, type **180**.

Copyright © 2016, Juniper Networks, Inc.

WARNING: When enabling the trickling option, it is important to understand that trickling might send part of the file to the client during the antivirus scan. It is possible that some of the content could be received by the client and the client might become infected before the file is fully scanned.

e. URI checking is on by default. To turn it off, clear **yes** in the URI check box.

f. In the Content size Limit box, type **20000**.

g. In the Scan engine timeout box, type **1800**.

4. Configure fallback settings by clicking the **Fallback settings** tab. In this example, all fallback options are set to log and permit. Click **Log and permit** for the following items: Default action, Content size, Engine not ready, Timeout, Out of resource, Too many requests.

5. Configure notification options by clicking the **Notification options** tab. You can configure notifications for both fallback blocking and fallback nonblocking actions and for virus detection.

   To configure notifications for Fallback settings:

   a. For Notification type, click **Protocol**.

   b. For Notify mail sender, click **yes**.

   c. In the Custom message box, type **Fallback block action occurred**.

   d. In the Custom message subject box, type ***Antivirus fallback Alert***.

6. To configure notification options for virus detection, click the **Notification options cont...** tab.

   a. For the Notification type option button, select **Protocol**.

   b. For the Notify mail sender option button, select **yes**.

   c. In the Custom message box, type **Virus has been detected**.

   d. In the Custom message subject box, type ***Virus detected***.

7. Click **OK** to check your configuration and save it as a candidate configuration.

8. If you are done configuring the device, click **Actions>Commit**.

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To configure the Sophos antivirus feature profile using the CLI:

NOTE: The following example shows you how to create a custom Sophos profile. If you want to use the Juniper Networks preconfigured profile, use the profile named junos-sophos-av-defaults in your UTM policy. See "Example: Configuring Sophos Antivirus UTM Policies" on page 5731.

1. Select and configure the engine type. Because you are configuring Sophos antivirus, you configure sophos-engine.

   ```
   [edit]
   user@host# set security utm feature-profile anti-virus type sophos-engine
   ```

2. Commit the configuration, and restart the device. After the device restarts, enter configuration mode again.

3. Select a time interval for updating the data files. The default antivirus pattern-update interval is 1440 minutes (every 24 hours). You can choose to leave this default, or you can change it. You can also force a manual update, if needed. To change the default from every 24 hours to every 48 hours:

   ```
   [edit security utm feature-profile anti-virus]
   user@host# set sophos-engine pattern-update interval 2880
   ```

4. Configure the network device with the proxy server details, to download the pattern update from a remote server:

   ```
   [edit security utm feature-profile anti-virus]
   user@host# set sophos-engine pattern-update proxy
   ```

5. In most circumstances, you will not need to change the URL to update the pattern database. If you do need to change this option, use the following command:

   ```
   [edit security utm feature-profile anti-virus]
   user@host# set sophos-engine pattern-update url
     http://www.example.net/test-download
   ```

6. You can configure the device to notify a specified administrator when data files are updated. This is an e-mail notification with a custom message and a custom subject line.

   ```
   [edit security utm feature-profile anti-virus]
   user@host# set sophos-engine pattern-update email-notify admin-email
     admin@example.net custom-message "Sophos antivirus data file was updated"
     custom-message-subject "AV data file updated"
   ```

7. Configure a list of fallback options as block, log and permit, or permit. The default setting is log-and-permit. You can use the default settings, or you can change them.

   Configure the content size action. In this example, if the content size is exceeded, the action taken is block.

   First create the profile named sophos-prof1.

   ```
   [edit security utm feature-profile anti-virus]
   user@host# edit sophos-engine profile sophos-prof1
   ```

Copyright © 2016, Juniper Networks, Inc.

Configure the content size fallback-option to block.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options content-size block
```

Configure the default fallback option to log-and-permit.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options default log-and-permit
```

Configure log-and-permit if the antivirus engine is not ready.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options engine-not-ready log-and-permit
```

Configure log-and-permit if the device is out of resources.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options out-of-resources log-and-permit
```

Configure log-and-permit if a virus scan timeout occurs.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options timeout log-and-permit
```

Configure log-and-permit if there are too many requests for the virus engine to handle.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set fallback-options too-many-requests log-and-permit
```

8. Configure notification options. You can configure notifications for fallback blocking, fallback nonblocking actions, and virus detection.

   In this step, configure a custom message for the fallback blocking action and send a notification for protocol-only actions to the administrator and the sender.

   ```
   [edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
   user@host# set notification-options fallback-block custom-message ***Fallback
     block action occurred*** custom-message-subject Antivirus Fallback Alert
     notify-mail-sender type protocol-only allow email administrator-email
     admin@example.net
   ```

9. Configure a notification for protocol-only virus detection, and send a notification.

   ```
   [edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
   user@host#set notification-options virus-detection type protocol-only
     notify-mail-sender custom-message-subject ***Virus detected***
     custom-message Virus has been detected
   ```

10. Configure content size parameters.

Copyright © 2016, Juniper Networks, Inc.

NOTE: When you configure the content-size value, keep in mind that in certain cases, content size is available in the protocol headers, so the max-content-size fallback is applied before a scan request is sent. However, in many cases, content size is not provided in the protocol headers. In these cases, the TCP payload is sent to the antivirus scanner and accumulates until the end of the payload. If the accumulated payload exceeds the maximum content size value, then max-content-size fallback is applied. The default fallback action is log and permit, so you may want to change this option to block, in which case such a packet is dropped and a block message is sent to the client.

In this example, if the content size exceeds 20 MB, the packet is dropped.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set scan-options content-size-limit 20000
```

11. URI checking is on by default. To turn off URI checking:

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set scan-options no-uri-check
```

12. Configure the timeout setting for the scanning operation to 1800 seconds.

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set scan-options timeout 1800
```

13. The Sophos Extensible List servers contain the virus and malware database for scanning operations. Set the response timeout for these servers to 3 seconds (the default is 2 seconds).

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set scan-options sxl-timeout 3
```

14. Configure the Sophos Extensible List server retry option to 2 retries (the default is 1).

```
[edit security utm feature-profile anti-virus sophos-engine profile sophos-prof1]
user@host# set scan-options sxl-retry 2
```

15. Configure the trickling setting to 180 seconds. If you use trickling, you can also set timeout parameters. Trickling applies only to HTTP. HTTP trickling is a mechanism used to prevent the HTTP client or server from timing out during a file transfer or during antivirus scanning.

WARNING: When you enable the trickling option, keep in mind that trickling might send part of a file to the client during its antivirus scan. It is therefore possible that some of the content could be received by the client before the file has been fully scanned.

```
[edit security utm feature-profile anti-virus]
user@host# set sophos-engine profile sophos-prof1 trickling timeout 180
```

Copyright © 2016, Juniper Networks, Inc.

16. Configure the antivirus module to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list that ships with the device called junos-default-bypass-mime. In this example, you use the lists that you set up earlier.

```
[edit security utm feature-profile anti-virus]
user@host# set mime-whitelist list avmime2
[edit security utm feature-profile anti-virus]
user@host# set mime-whitelist list exception-avmime2
```

17. Configure the antivirus module to use URL bypass lists. If you are using a URL whitelist, this is a custom URL category you have previously configured as a custom object. URL whitelists are valid only for HTTP traffic. In this example you use the lists that you set up earlier.

```
[edit security utm feature-profile anti-virus]
user@host# set url-whitelist custurl2
```

Verification

Obtaining Information About the Current Antivirus Status

Action From operational mode, enter the **show security utm anti-virus status** command to view the antivirus status.

```
user@host>show security utm anti-virus status
```

Meaning

- Antivirus key expire date—The license key expiration date.
- Update server—URL for the data file update server.
  - Interval—The time period, in minutes, when the device will update the data file from the update server.
  - Pattern update status—When the data file will be updated next, displayed in minutes.
  - Last result—Result of the last update. If you already have the latest version, this will display **already have latest database**.
- Antivirus signature version—Version of the current data file.
- Scan engine type—The antivirus engine type that is currently running.
- Scan engine information—Result of the last action that occurred with the current scan engine.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Sophos Antivirus UTM Policies

This example shows how to create a UTM policy for Sophos antivirus.

- Requirements on page 5731
- Overview on page 5731
- Configuration on page 5731
- Verification on page 5732

Requirements

Before you create the UTM policy, create custom objects and the Sophos feature profile.

1. Configure UTM custom objects and MIME lists. See "Example: Configuring Sophos Antivirus Custom Objects" on page 5721.
2. Configure the Sophos antivirus feature profile. See "Example: Configuring Sophos Antivirus Feature Profile" on page 5724.

Overview

After you have created an antivirus feature profile, you configure a UTM policy for an antivirus scanning protocol and attach this policy to a feature profile. In this example, HTTP will be scanned for viruses, as indicated by the **http-profile** statement. You can scan other protocols as well by creating different profiles or adding other protocols to the profile, such as: imap-profile, pop3-profile, and smtp-profile.

Configuration

GUI Step-by-Step Procedure

To configure a UTM policy for Sophos antivirus:

1. Click the **Configure** tab from the taskbar, and then select **Security>Policy>UTM Policies**. Then click **Add**.
2. Click the **Main** tab. In the Policy name box, type **utmp3**.
3. Click the **Anti-Virus profiles** tab. In the HTTP profile list, select **sophos-prof1**.
4. Click **OK** to check your configuration and save it as a candidate configuration.
5. If you are done configuring the device, select **Actions>Commit**.

Step-by-Step Procedure

To configure a UTM policy for Sophos antivirus:

1. Go to the edit security utm hierarchy.

```
[edit]
user@host# edit security utm
```

2. Create the UTM policy utmp3 and attach it to the http-profile sophos-prof1.

```
[edit security utm]
user@host# set utm-policy utmp3 anti-virus http-profile sophos-prof1
```

Copyright © 2016, Juniper Networks, Inc.

NOTE: You can use the default Sophos feature profile settings by replacing sophos-prof1 in the above statement with junos-sophos-av-defaults.

Verification

To verify the configuration, enter the **show security utm utm-policy utmp3** command.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720
- Example: Configuring Sophos Antivirus Feature Profile on page 5724

Example: Configuring Sophos Antivirus Firewall Security Policies

This example shows how to create a security policy for Sophos antivirus.

- Requirements on page 5732
- Overview on page 5732
- Configuration on page 5733
- Verification on page 5734

Requirements

Before you create the security policy, create custom objects, the Sophos feature profile, and the UTM policy.

1. Configure UTM custom objects and MIME lists. See "Example: Configuring Sophos Antivirus Custom Objects" on page 5721.
2. Configure the Sophos antivirus feature profile. See "Example: Configuring Sophos Antivirus Feature Profile" on page 5724.
3. Configure a UTM policy. See "Example: Configuring Sophos Antivirus UTM Policies" on page 5731.

Overview

Create a firewall security policy that will cause traffic from the untrust zone to the trust zone to be scanned by Sophos antivirus using the feature profile settings defined in "Example: Configuring Sophos Antivirus Feature Profile" on page 5724. Because the match application configuration is set to any, all application types will be scanned.

Copyright © 2016, Juniper Networks, Inc.

Configuration

GUI Step-by-Step Procedure

To configure a security policy for Sophos antivirus:

1. Configure the untrust to trust policy to match any source address or destination address, and select the applications to be scanned to **any**.
   a. Click the **Configure** tab from the taskbar, and then select **Security>Policy>FW Policies**. Then select **Add**.
   b. In the Policy Name box, type **p3**.
   c. In the Policy Action box, select **permit**.
   d. In the From Zone list, select **untrust**.
   e. In the To Zone list, select **trust**.
   f. In the Source Address and Destination Address boxes, make sure that Matched is set to **any**.
   g. In the Applications boxes, select **any** from the Application/Sets list, and move it to the Matched list.
2. Attach the UTM policy named utmp3 to the firewall security policy. This will cause matched traffic to be scanned by the Sophos antivirus feature.
   a. From the Edit Policy box, click the **Application Services** tab.
   b. In the UTM Policy list, select **utmp3**.
3. Click **OK** to check your configuration and save it as a candidate configuration.
4. If you are done configuring the device, select **Actions>Commit**.

Step-by-Step Procedure

To configure a security policy for Sophos antivirus:

1. Configure the untrust to trust policy to match any source-address.

```
[edit security]
user@host# set policies from-zone untrust to-zone trust policy p3 match
  source-address any
```

2. Configure the untrust to trust policy to match any destination-address.

```
[edit security]
user@host# set policies from-zone untrust to-zone trust policy p3 match
  destination-address any
```

3. Configure the untrust to trust policy to match any application type.

```
[edit security]
user@host# set policies from-zone untrust to-zone trust policy p3 match application
  any
```

4. Attach the UTM policy named utmp3 to the firewall security policy. This will cause matched traffic to be scanned by the Sophos antivirus feature.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security]
user@host# set policies from-zone untrust to-zone trust policy p3 then permit
    application-services utm-policy utmp3
```

Verification

To verify the configuration, enter the **show security policies** command.

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Sophos Antivirus Configuration Overview on page 5720
- Example: Configuring Sophos Antivirus Feature Profile on page 5724

Example: Configuring Sophos Antivirus Scanner with SSL Forward Proxy

This example shows how to configure Sophos antivirus over SSL forward proxy to support HTTPS traffic passing through SRX Series devices.

- Requirements on page 5734
- Overview on page 5734
- Configuration on page 5734
- Verification on page 5737

Requirements

Before you begin, understand Sophos antivirus features. See "Sophos Antivirus Features" on page 5718.

Overview

In this example, you configure Sophos antivirus over SSL forward proxy to support HTTPS traffic. You load the PKI certificate, generate a self-signed CA certificate, configure a trusted CA list, configure an SSL proxy profile using the root certificate, and enable SSL forward proxy. To configure UTM over SSL forward proxy, first match the source/destination/application, set up the SSL proxy service, and perform scanning to determine whether to block or permit the requests.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **edit** hierarchy level, and then enter **commit** from configuration mode.

```
request security pki generate-key-pair certificate-id ssl-inspect-ca size 2048 type rsa
request security pki local-certificate generate-self-signed certificate-id ssl-inspect-ca
    domain-name www.example.net subject
    "CN=www.example.net,OU=IT,O=example,L=Sunnyvale,ST=CA,C=US" email
    security-admin@example.net
set security pki ca-profile trusted-ca-example ca-identity trusted-ca-example
```


Copyright © 2016, Juniper Networks, Inc.

```
request security pki ca-certificate load ca-profile trusted-ca-example filename
  trusted-ca-example.crt
set services ssl proxy profile ssl-inspect-profile root-ca ssl-inspect-ca
set services ssl proxy profile ssl-inspect-profile trusted-ca trusted-ca-example
set security policies from-zone untrust to-zone trust policy 1 then permit
  application-services ssl-proxy profile-name ssl-inspect-profile
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Sophos Antivirus over SSL forward proxy:

1. Generate a self-signed CA certificate on the device.

```
user@host> request security pki generate-key-pair certificate-id ssl-inspect-ca size
  2048 type rsa
user@host> request security pki local-certificate generate-self-signed certificate-id
  ssl-inspect-ca domain-name www.example.net subject
  "CN=www.example.net,OU=IT,O=example,L=Sunnyvale,ST=CA,C=US" email
  security-admin@example.net
```

2. Configure a trusted CA list.

```
[edit]
user@host# set security pki ca-profile trusted-ca-example ca-identity
  trusted-ca-example
```

```
user@host> request security pki ca-certificate load ca-profile trusted-ca-example
  filename trusted-ca-example.crt
```

3. Configure an SSL proxy profile using a root certificate.

```
[edit]
user@host# set services ssl proxy profile ssl-inspect-profile root-ca ssl-inspect-ca
user@host# set services ssl proxy profile ssl-inspect-profile trusted-ca
  trusted-ca-example
```

4. Enable SSL forward proxy.

```
[edit]
user@host# set security policies from-zone untrust to-zone trust policy 1 then permit
  application-services ssl-proxy profile-name ssl-inspect-profile
```

Results

From configuration mode, confirm your configuration by entering the **show security utm**, **show services**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm
  traceoptions {
    flag all;
  }
  application-proxy {
    traceoptions {
```

Copyright © 2016, Juniper Networks, Inc.

```
        flag sophos-anti-virus;
      }
    }
    feature-profile {
      anti-virus {
        type sophos-engine;
        traceoptions {
          flag all;
        }
        sophos-engine {
          sxl-timeout 4;
          profile profile1 {
            fallback-options {
              default log-and-permit;
              content-size log-and-permit;
              engine-not-ready log-and-permit;
              timeout log-and-permit;
              out-of-resources log-and-permit;
              too-many-requests log-and-permit;
            }
            scan-options {
              uri-check;
            }
            notification-options {
              virus-detection {
                type message;
              }
              fallback-block {
                type message;
              }
            }
          }
        }
      }
    }
    utm-policy policy1 {
      anti-virus {
        http-profile profile1;
      }
    }
[edit]
user@host# show services
    ssl {
      traceoptions {
        file ssl_trace size 1g;
        flag all;
      }
      proxy {
        profile ssl-p {
          root-ca haojue;
          actions {
            ignore-server-auth-failure;
          }
        }
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security policies
  from-zone trust to-zone untrust {
    policy trust_2_untrust {
      match {
        source-address any;
        destination-address any;
        application [ junos-http junos-https ];
      }
      then {
        permit {
          application-services {
            ssl-proxy {
              profile-name ssl-p;
            }
            utm-policy policy1;
          }
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Security PKI Local Certificate on page 5737
- Verifying UTM Antivirus Statistics on page 5738
- Verifying UTM Antivirus Statistics Details on page 5738
- Verifying UTM Antivirus Status on page 5739

Verifying the Security PKI Local Certificate

Purpose Verify the security PKI local certificate.

Action From configurationl mode, enter the **show security pki local-certificate** command.

```
user@host#  show security pki local-certificate
Certificate identifier: SELF-SIGNED
  Issued to: abc, Issued by: CN = abc
  Validity:
    Not before: 02-20-2015 00:49 UTC
    Not after: 02-19-2020 00:49 UTC
  Public key algorithm: rsaEncryption(2048 bits)

Certificate identifier: ssl-inspect-ca
  Issued to: www.example.net, Issued by: CN = www.example.net, OU = IT, O =
example, L = Sunnyvale, ST = CA, C = US
  Validity:
    Not before: 01-28-2016 22:28 UTC
    Not after: 01-26-2021 22:28 UTC
  Public key algorithm: rsaEncryption(2048 bits)
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The sample output confirms that the PKI local ceritificate ssl-inspect-ca is configured.

Verifying UTM Antivirus Statistics

Purpose Verify UTM antivirus statistics.

Action From operational mode, enter the **show security utm anti-virus statistics** command.

```
user@host> show security utm anti-virus statistics
MIME-whitelist passed:              0
    URL-whitelist passed:               0

    Scan Requests:
        Total           Clean           Thread found    Fall back    Abort
          0               0                 0             0           0

    Fall back:                  log-and-permit      block       permit
        Engine not ready:             0               0           0
        Out of resources:             0               0           0
        Timeout:                      0               0           0
        Maxmium content size:         0               0           0
        Too many requests:            0               0           0
        Others                        0               0           0
```

Meaning The sample output shows the list of UTM antivirus statistics.

Verifying UTM Antivirus Statistics Details

Purpose Verify UTM antivirus statistics details.

Action From operational mode, enter the **show security utm anti-virus statistics details** command.

```
user@host> show security utm anti-virus statistics details
HTTP
        MIME-whitelist passed:              0
        URL-whitelist passed:               0

        URI request:
         Total   Clean     Threat-found      Need-further-inspection  Abort
           10          1             1                      8                0

     File request:
         Total       Clean     Threat-found      Fallback      Abort
               8          6             1               1           0

        Fall back:                  log-and-permit      block      permit
            Engine not ready:             0               0          0
            Out of resources:             0               0          0
            Timeout:                      0               0          0
            Maxmium content size:         1               0          0
            Too many requests:            0               0          0
            Others                        0               0          0

    FTP
        Scan request:
         Total       Clean     Threat-found      Fallback      Abort
               10          8             1               1           0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Fall back:                  log-and-permit      block      permit
        Engine not ready:             0               0          0
        Out of resources:             0               0          0
        Timeout:                      0               0          0
        Maxmium content size:         1               0          0
        Too many requests:            0               0          0
        Others                        0               0          0

SMTP
    Scan request:
     Total       Clean     Threat-found      Fallback      Abort
           10          8             1               1          0

    Fall back:                  log-and-permit      block      permit
        Engine not ready:             0               0          0
        Out of resources:             0               0          0
        Timeout:                      0               0          0
        Maxmium content size:         1               0          0
        Too many requests:            0               0          0
        Others                        0               0          0

POP3
    Scan request:
     Total       Clean     Threat-found      Fallback      Abort
           10          8             1               1          0

    Fall back:                  log-and-permit      block      permit
        Engine not ready:             0               0          0
        Out of resources:             0               0          0
        Timeout:                      0               0          0
        Maxmium content size:         1               0          0
        Too many requests:            0               0          0
        Others                        0               0          0

IMAP
    Scan request:
     Total       Clean     Threat-found      Fallback      Abort
           10          8             1               1          0

    Fall back:                  log-and-permit      block      permit
        Engine not ready:             0               0          0
        Out of resources:             0               0          0
        Timeout:                      0               0          0
        Maxmium content size:         1               0          0
        Too many requests:            0               0          0
        Others                        0               0          0
```

Meaning The sample output shows the list of antivirus statistics details.

Verifying UTM Antivirus Status

Purpose Verify UTM antivirus status.

Action From operational mode, enter the **show security utm anti-virus status** command to view the antivirus status.

```
user@host> show security utm anti-virus status
```

Copyright © 2016, Juniper Networks, Inc.

```
Anti-virus Key Expiry Date: 07/01/2010 00:00:00
       Update server: http://update.juniper-updates.net//
          Interval: 1440 minutes
          Auto update status: next update in 1440 minutes
          Last result: No error
       Anti-virus data file info:
          Version:
       Scan engine information:
          Last action result: No error(0x00000000)
          Engine type: sophos-engine
```

Meaning

- Antivirus key expire date—The license key expiration date.
- Update server—URL for the data file update server.
  - Interval—The time period, in minutes, when the device updates the data file from the update server.
  - Auto update status—Displays the next automatic update of the data file in minutes.
  - Last result—Result of the last database update.
- Antivirus signature version—Version of the current antivirus signature data file.
- Scan engine type—The antivirus scan engine type that is currently running.
- Scan engine information—Result of the last action that occurred with the current scan engine.

Related Documentation

- Sophos Antivirus Features on page 5718

Managing Sophos Antivirus Data Files

Before you begin:

- Install a Sophos antivirus license. See the *Installation and Upgrade Guide for Security Devices*.
- Configure Sophos as the antivirus feature for the device. See "Example: Configuring Sophos Antivirus Feature Profile" on page 5724. To set the antivirus engine type, you run the **set security utm feature-profile anti-virus type sophos-engine** statement.

In this example, you configure the security device to update the data files automatically every 4320 minutes (every 3 days). The default data file update interval is 1440 minutes (every 24 hours).

To automatically update Sophos data files:

```
[edit security utm feature-profile anti-virus]
user@host# set sophos-engine pattern-update interval 4320
```

NOTE: The following commands are performed from CLI operational mode.

Copyright © 2016, Juniper Networks, Inc.

To manually update data files:

```
user@host> request security utm anti-virus sophos-engine pattern-update
```

To manually reload data files:

```
user@host> request security utm anti-virus sophos-engine pattern-reload
```

To manually delete data files:

```
user@host> request security utm anti-virus sophos-engine pattern-delete
```

To check the status of antivirus, which also shows the data files version:

```
user@host> show security utm anti-virus status
```

To check the status of the proxy server:

```
user@host> show security utm anti-virus status
```

Related Documentation

- Sophos Antivirus Protection Overview on page 5717
- Understanding Sophos Antivirus Data File Update on page 5719
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 76

Configuring and Monitoring Content Filtering

- Configuring Content Filtering on page 5745

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 248

Configuring Content Filtering

- Content Filtering Overview on page 5745
- Understanding Content Filtering Protocol Support on page 5746
- Specifying Content Filtering Protocols (CLI Procedure) on page 5747
- Content Filtering Configuration Overview on page 5748
- Example: Configuring Content Filtering Custom Objects on page 5749
- Example: Configuring Content Filtering Feature Profiles on page 5751
- Example: Configuring Content Filtering UTM Policies on page 5754
- Example: Attaching Content Filtering UTM Policies to Security Policies on page 5756
- Monitoring Content Filtering Configurations on page 5758

Content Filtering Overview

Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, and protocol command. The content filter controls file transfers across the gateway by checking traffic against configured filter lists.

The content filter module evaluates traffic before all other UTM modules, except Web Filtering. Therefore, if traffic meets criteria configured in the content-filter, the content-filter acts first upon this traffic.

You can configure the following types of content filters:

- MIME Pattern Filter — MIME patterns are used to identify the type of traffic in HTTP and MAIL protocols. There are two lists of MIME patterns that are used by the content filter to determine the action to be taken. The block MIME list contains a list of MIME type traffic that is to be blocked by the content filter. The MIME exception list contains MIME patterns that are not to be blocked by the content filter and are generally subsets of items on the block list. Note that the exception list has a higher priority than the block list. If you have MIME entries that appear on both lists, those MIME types are not blocked by the content filter because the exception list takes priority. Therefore, when adding items to the exception list, it is to your advantage to be specific.
- Block Extension List — Because the name of a file is available during file transfers, using file extensions is a highly practical way to block or allow file transfers. The content filter list contains a list of file extensions to be blocked. All protocols support the use of the block extension list.

Copyright © 2016, Juniper Networks, Inc.

- Protocol Command Block and Permit Lists — Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level.

  The block and permit command lists are intended to be used in combination, with the permit list acting as an exception list to the block list.

NOTE: **If a protocol command appears on the both the permit list and the block list, that command is permitted.**

Because not all harmful files or components can be controlled by the MIME type or by the file extension, you can also use the content filter module to block ActiveX, Java Applets, and other types of content. The following types of content blocking are supported only for HTTP:

- Block ActiveX
- Block Java applets
- Block cookies
- Block EXE files
- Block ZIP files

Related Documentation

- Unified Threat Management Overview on page 5583
- Understanding Content Filtering Protocol Support on page 5746
- Content Filtering Configuration Overview on page 5748
- Monitoring Content Filtering Configurations on page 5758

Understanding Content Filtering Protocol Support

Each supported protocol may implement available content filters differently. Not all filtering capabilities are supported for each protocol.

This topic contains the following sections:

- HTTP Support on page 5746
- FTP Support on page 5747
- E-Mail Support on page 5747

HTTP Support

The HTTP protocol supports all content filtering features. With HTTP, the content filter remains in the gateway, checking every request and response between the HTTP client and server.

Copyright © 2016, Juniper Networks, Inc.

If an HTTP request is dropped due to content filtering, the client receives a response such as:

```
<custom drop message/user-configured drop
  message>.<src_port><dst_ip>:<dst_port>Download request was dropped due to
  <reason>
```

Therefore, a message may appear as follows:

```
Juniper Networks Firewall Content Filtering blocked request. 5.5.5.1:80->4.4.4.1:55247
  Download request was dropped due to file extension block list
```

FTP Support

The FTP protocol does not support all content filtering features. It supports only the following: Block Extension List and Protocol Command Block List.

When content filtering blocks an FTP request, the following response is sent through the control channel:

```
550 <src_ip>:<src_port>-<dst_ip>:<dst_port><custom drop message/user-configured
  drop message> for Content Filtering file extension block list.>
```

Therefore, a message may appear as follows:

```
550 5.5.5.1:21->4.4.4.1:45237 Requested action not taken and the request is dropped for
  Content Filtering file extension block list
```

E-Mail Support

E-mail protocols (SMTP, IMAP, POP3) have limited content filtering support for the following features: Block Extension List, Protocol Command Block List, and MIME Pattern Filtering. Support is limited for e-mail protocols for the following reasons:

- The content filter scans only one level of an e-mail header. Therefore recursive e-mail headers and encrypted attachments are not scanned.
- If an entire e-mail is MIME encoded, the content filter can only scan for the MIME type.
- If any part of an e-mail is blocked due to content filtering, the original e-mail is dropped and replaced by a text file with an explanation for why the e-mail was blocked.

Related Documentation

- Unified Threat Management Overview on page 5583
- Specifying Content Filtering Protocols (CLI Procedure) on page 5747
- Content Filtering Configuration Overview on page 5748
- Monitoring Content Filtering Configurations on page 5758

Specifying Content Filtering Protocols (CLI Procedure)

To configure content filtering protocols, use the following CLI configuration statements:

```
content-filtering {
  profile name {
    permit-command cmd-list
```

Copyright © 2016, Juniper Networks, Inc.

```
        block-command cmd-list
      block-extension file-ext-list
        block-mime {
          list mime-list
          exception ex-mime-list
        }
    block-content-type {
      activex
        java-applet
        exe
        zip
      http-cookie
    }
    notification-options {
      type { message }
        notify-mail-sender
      custom-message msg
    }
  }
  traceoptions {
    flag {
      all
        basic
        detail
      }
    }
  }
```

Related Documentation

- Unified Threat Management Overview on page 5583
- Example: Configuring Content Filtering Custom Objects on page 5749
- Example: Configuring Content Filtering Feature Profiles on page 5751
- Example: Configuring Content Filtering UTM Policies on page 5754

Content Filtering Configuration Overview

A content security filter blocks or allows certain type of traffic base on the mime type, file extension, protocol commands and embedded object type. The content filter controls file transfers across the gateway by checking traffic against configured filter lists. The content filtering module evaluates traffic before all other UTM modules, if traffic meets the criteria configured in the content filter, the content filter acts first upon this traffic. The following procedure lists the recommended order in which you should configure content filters:

1. Configure UTM custom objects for the feature. See "Example: Configuring Content Filtering Custom Objects" on page 5749.
2. Configure the main feature parameters using feature profiles. See "Example: Configuring Content Filtering Feature Profiles" on page 5751.
3. Configure a UTM policy for each protocol and attach this policy to a profile. See "Example: Configuring Content Filtering UTM Policies" on page 5754.

Copyright © 2016, Juniper Networks, Inc.

4. Attach the UTM policy to a security policy. See "Example: Attaching Content Filtering UTM Policies to Security Policies" on page 5756.

Related Documentation

- Unified Threat Management Overview on page 5583

Example: Configuring Content Filtering Custom Objects

This example shows how to configure content filtering custom objects.

- Requirements on page 5749
- Overview on page 5749
- Configuration on page 5749
- Verification on page 5751

Requirements

Before you begin:

1. Decide on the type of content filter you require. See "Content Filtering Overview" on page 5745.
2. Understand the order in which content filtering parameters are configured. See "Content Filtering Configuration Overview" on page 5748.

Overview

In this example, you define custom objects that are used to create content filtering profiles. You perform the following tasks to define custom objects:

1. Create two protocol command lists called ftpprotocom1 and ftpprotocom2, and add user, pass, port, and type commands to it.
2. Create a filename extension list called extlist2, and add the .zip, .js, and .vbs extensions to it.
3. Define block-mime list call cfmime1 and add patterns to the list.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects protocol-command ftpprotocom1 value [user pass port
    type]
set security utm custom-objects protocol-command ftpprotocom2 value [user pass port
    type]
set security utm custom-objects filename-extension extlist2 value [zip js vbs]
set security utm custom-objects mime-pattern cfmime1 value [video/quicktime
    image/x-portable-anymap x-world/x-vrml]
```

Copyright © 2016, Juniper Networks, Inc.

```
set security utm custom-objects mime-pattern ex-cfmime1 value
  [video/quicktime-inappropriate]
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure content filtering custom objects:

1. Create two protocol command lists.

```
[edit security utm]
user@host# set custom-objects protocol-command ftpprotocom1
[edit security utm]
user@host# set custom-objects protocol-command ftpprotocom2
```

2. Add protocol commands to the list.

```
[edit security utm]
user@host# set custom-objects protocol-command ftpprotocom1 value [user pass
  port type]
[edit security utm]
user@host# set custom-objects protocol-command ftpprotocom2 value [user pass
  port type]
```

3. Create a filename extension list.

```
[edit security utm]
user@host# set custom-objects filename-extension extlist2
```

4. Add extensions to the list.

```
[edit security utm]
user@host# set custom-objects filename-extension extlist2 value [zip js vbs]
```

5. Create antivirus scanning lists.

```
[edit security utm]
user@host# set custom-objects mime-pattern cfmime1
user@host# set custom-objects mime-pattern ex-cfmime1
```

6. Add patterns to the lists.

```
[edit security utm]
user@host# set custom-objects mime-pattern cfmime1 value [video/quicktime
  image/x-portable-anymap x-world/x-vrml]
user@host# set custom-objects mime-pattern ex-cfmime1 value
  [video/quicktime-inappropriate]
```

Results

From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security utm
    custom-objects {
      mime-pattern {
        cfmime1 {
          value [ video/quicktime image/x-portable-anymap x-world/x-vrml ];
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        ex-cfmime1 {
            value video/quicktime-inappropriate;
        }
    }
    filename-extension {
        extlist2 {
            value [ zip js vbs ];
        }
    }
    protocol-command {
        ftpprotocom1 {
            value [ user pass port type ];
        }
        }
        protocol-command {
        ftpprotocom2 {
            value [ user pass port type ];
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Content Filtering Custom Objects

Purpose Verify the content filtering custom objects.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Unified Threat Management Overview on page 5583
- Content Filtering Overview on page 5745
- Content Filtering Configuration Overview on page 5748
- Example: Configuring Content Filtering Feature Profiles on page 5751
- Example: Configuring Content Filtering UTM Policies on page 5754
- Example: Attaching Content Filtering UTM Policies to Security Policies on page 5756

Example: Configuring Content Filtering Feature Profiles

This example describes how to configure the content filtering feature profiles.

- Requirements on page 5752
- Overview on page 5752
- Configuration on page 5752
- Verification on page 5754

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

1. Decide on the type of content filter you require. See "Content Filtering Overview" on page 5745.
2. Create custom objects. See "Content Filtering Configuration Overview" on page 5748.

Overview

In this example, you configure a feature profile called confilter1 and specify the following custom objects to be used for filtering content:

1. Apply the ftpprotocom1 protocol command list custom object to confilter1.
2. Apply blocks to Java applets, executable files, and HTTP cookies.
3. Apply the extension list extlist2 custom object to confilter1 for blocking extensions.
4. Apply the MIME pattern list custom objects cfmime1 and ex-cfmime1 to the confilter1 for blocking MIME types.
5. Apply the protocol permit command custom object ftpprotocom2 to confilter1. (The permit protocol command list acts as an exception list for the block protocol command list.)

NOTE: Protocol command lists, both permit and block, are created by using the same custom object.

6. Configure a custom message to send a notification.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile content-filtering profile confilter1
set security utm feature-profile content-filtering profile confilter1 block-command
    ftpprotocom1
set security utm feature-profile content-filtering profile confilter1 block-content-type
    java-applet exe http-cookie
set security utm feature-profile content-filtering profile confilter1 block-extension extlist2
set security utm feature-profile content-filtering profile confilter1 block-mime list cfmime1
    exception ex-cfmime1
set security utm feature-profile content-filtering profile confilter1 permit-command
    ftpprotocom2
set security utm feature-profile content-filtering profile confilter1 notification-options
    custom-message "the action is not taken" notify-mail-sender type message
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a content filtering feature profiles:

1. Create a content filtering profile.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1
```

2. Apply a protocol command list custom object to the profile.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1 block-command
  ftpprotocom1
```

3. Apply blocks to available content.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1 block-content-type
  java-applet exe http-cookie
```

4. Apply an extension list custom object to the profile.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1 block-extension
  extlist2
```

5. Apply pattern list custom objects to the profile.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1 block-mime list
  cfmime1 exception ex-cfmime1
```

6. Apply the protocol permit command custom object to the profile.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1 permit-command
  ftpprotocom2
```

7. Configure the notification options.

```
[edit security utm]
user@host# set feature-profile content-filtering profile confilter1m
  notification-options custom-message “the action is not taken” notify-mail-sender
  type message
```

Results

From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security utm
    feature-profile {
      content-filtering {
        profile contentfilter1;
        profile confilter1 {
          permit-command ftpprotocom2;
          block-command ftpprotocom1;
```

Copyright © 2016, Juniper Networks, Inc.

```
                block-extension extlist2;
                block-mime {
                    list cfmime1;
                    exception ex-cfmime1;
                }
                block-content-type {
                    java-applet;
                    exe;
                    http-cookie;
                }
                notification-options {
                    type message;
                    notify-mail-sender;
                    custom-message " the action is not taken";
                }
            }
        }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration of Content Filtering Feature Profile

Purpose Verify the content filtering feature profile.

Action From operational mode, enter the **show configuration security utm** command.

Related Documentation

- Unified Threat Management Overview on page 5583
- Content Filtering Overview on page 5745
- Content Filtering Configuration Overview on page 5748
- Example: Configuring Content Filtering Custom Objects on page 5749
- Example: Configuring Content Filtering UTM Policies on page 5754
- Example: Attaching Content Filtering UTM Policies to Security Policies on page 5756

Example: Configuring Content Filtering UTM Policies

This example describes how to create a content filtering UTM policy to attach to your feature profile.

- Requirements on page 5755
- Overview on page 5755
- Configuration on page 5755
- Verification on page 5755

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin:

1. Decide on the type of content filter you require. See "Content Filtering Overview" on page 5745.
2. Configure UTM custom objects for each feature and define the content-filtering profile. See "Content Filtering Configuration Overview" on page 5748.

Overview

You configure UTM policies to selectively enforce various UTM solutions on network traffic passing through a UTM-enabled device. Through feature profiles you associate custom objects to these policies and specify blocking or permitting certain types of traffic.

In this example, you configure a UTM policy called utmp4, and then assign the preconfigured feature profile confilter1 to this policy.

Configuration

Step-by-Step Procedure

To configure a content filtering UTM policy:

You can configure different protocol applications in the UTM policy. The example only shows HTTP and not other protocols. Earlier you configured custom objects for FTP (ftpprotocom1 and ftpprotocom2). Next you should add a content filter policy for FTP, for example:

set security utm utm-policy utmp4 content-filtering ftp upload-profile confilter1

set security utm utm-policy utmp4 content-filtering ftp download-profile confilter1

1. Create a UTM policy.

   ```
   [edit security utm]
   user@host# set utm-policy utmp4
   ```

2. Attach the UTM policy to the profile.

   ```
   [edit security utm]
   user@host# set utm-policy utmp4 content-filtering http-profile contentfilter1
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security utm** command.

Related Documentation

- Unified Threat Management Overview on page 5583
- Content Filtering Overview on page 5745
- Content Filtering Configuration Overview on page 5748

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Content Filtering Custom Objects on page 5749
- Example: Configuring Content Filtering Feature Profiles on page 5751
- Example: Attaching Content Filtering UTM Policies to Security Policies on page 5756

Example: Attaching Content Filtering UTM Policies to Security Policies

This example shows how to create a security policy and attach the UTM policy to the security policy.

- Requirements on page 5756
- Overview on page 5756
- Configuration on page 5756
- Verification on page 5757

Requirements

Before you begin:

1. Configure UTM custom objects, define the content filtering profile, and create a UTM policy. See "Content Filtering Configuration Overview" on page 5748.
2. Enable and configure a security policy. See "Example: Configuring a Security Policy to Permit or Deny All Traffic" on page 1094.

Overview

By attaching content filtering UTM policies to security policies, you can filter traffic transiting from one security zone to another.

In this example, you create a security policy called p4 and specify that traffic from any source address to any destination address with an HTTP application matches the criteria. You then assign a UTM policy called utmp4 to the security policy p4. This UTM policy applies to any traffic that matches the criteria specified in the security policy p4.

Configuration

CLI Quick Configuration

To quickly attach a content filtering UTM policy to a security policy, copy the following commands and paste them into the CLI.

```
[edit]
set security policies from-zone trust to-zone untrust policy p4 match source-address any
set security policies from-zone trust to-zone untrust policy p4 match destination-address
    any
set security policies from-zone trust to-zone untrust policy p4 match application
    junos-htttp
set security from-zone trust to-zone untrust policy p4 then permit application-services
    utm-policy utmp4
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To attach a UTM policy to a security policy:

1. Create a security policy.

```
[edit]
user@host# edit security policies from-zone trust to-zone untrust policy p4
```

2. Specify the match conditions for the policy.

```
[edit security policies from-zone trust to-zone untrust policy p4]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application junos-http
```

3. Attach the UTM policy to the security policy.

```
[edit security policies from-zone trust to-zone untrust policy p4]
user@host# set then permit application-services utm-policy utmp4
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
    from-zone trust to-zone untrust {
      policy p4 {
        match {
          source-address any;
          destination-address any;
          application junos-http;
        }
        then {
          permit {
            application-services {
              utm-policy utmp4;
            }
          }
        }
      }
    }
      default-policy {
        permit-all;
      }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying Attaching Content Filtering UTM Policies to Security Policies

Purpose

Verify the attachment of the content filtering UTM policy to the security policy.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security policy** command.

Related Documentation

- Unified Threat Management Overview on page 5583
- Content Filtering Overview on page 5745
- Content Filtering Configuration Overview on page 5748
- Example: Configuring Content Filtering Custom Objects on page 5749
- Example: Configuring Content Filtering Feature Profiles on page 5751
- Example: Configuring Content Filtering UTM Policies on page 5754

Monitoring Content Filtering Configurations

Purpose View content filtering statistics.

Action To view content filtering statistics in the CLI, enter the **user@host > show security utm content-filtering statistics** command.

The content filtering **show statistics** command displays the following information:

```
Base on command list: # Blocked
Base on mime list: # Blocked
Base on extension list: # Blocked
ActiveX plugin: # Blocked
Java applet: # Blocked
EXE files: # Blocked
ZIP files: # Blocked
HTTP cookie: # Blocked
```

To view content filtering statistics using J-Web:

1. Select **Clear Content filtering statisticsMonitor>Security>UTM>Content FilteringMonitor>Security>UTM>Content Filtering**.

   The following statistics become viewable in the right pane.

   ```
   Base on command list: # Passed # Blocked
   Base on mime list: # Passed # Blocked
   Base on extension list: # Passed # Blocked
   ActiveX plugin: # Passed # Blocked
   Java applet: # Passed # Blocked
   EXE files: # Passed # Blocked
   ZIP files: # Passed # Blocked
   HTTP cookie: # Passed # Blocked
   ```

2. You can click **Clear Content filtering statistics** to clear all current viewable statistics and begin collecting new statistics.

Related Documentation

- Content Filtering Overview on page 5745
- Understanding Content Filtering Protocol Support on page 5746
- Content Filtering Configuration Overview on page 5748

Copyright © 2016, Juniper Networks, Inc.

- Example: Attaching Content Filtering UTM Policies to Security Policies on page 5756

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 77

Configuring Web Filtering

- Configuring Web Filtering on page 5763

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 249

Configuring Web Filtering

- Web Filtering Overview on page 5764
- Enhanced Web Filtering Overview on page 5766
- Understanding Enhanced Web Filtering Process on page 5767
- Example: Configuring Enhanced Web Filtering on page 5772
- Understanding the Quarantine Action for Enhanced Web Filtering on page 5780
- Example: Configuring Site Reputation Action for Enhanced Web Filtering on page 5781
- Understanding Integrated Web Filtering on page 5787
- Example: Configuring Integrated Web Filtering on page 5789
- Understanding Local Web Filtering on page 5797
- Example: Configuring Local Web Filtering on page 5799
- Understanding Redirect Web Filtering on page 5805
- Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects on page 5806
- Displaying Global SurfControl URL categories on page 5815
- Monitoring Web Filtering Configurations on page 5815

Copyright © 2016, Juniper Networks, Inc.

Web Filtering Overview

Web filtering lets you manage Internet usage by preventing access to inappropriate Web content. There are four types of Web filtering solutions:

- Integrated Web filtering—The integrated Web filtering solution intercepts every HTTP request in a TCP connection. In this case, the decision making is done on the device after it identifies the category for a URL either from user-defined categories or from a category server (SurfControl Content Portal Authority provided by Websense).

  NOTE: The integrated Web filtering feature is a separately licensed subscription service. When the license key for Web filtering has expired, no URLs are sent to the category server for checking, only local user-defined categories are checked.

  NOTE: Integrated Web filtering solution is supported only on branch SRX Series devices.

- Redirect Web filtering—The redirect Web filtering solution intercepts HTTP requests and sends them to an external URL filtering server, provided by Websense, to determine whether to block the requests.

  NOTE: Redirect Web filtering does not require a license.

- Local Web filtering—The local Web filtering solution intercepts every HTTP request in a TCP connection. In this case, the decision making is done on the device after it looks up a URL to determine if it is in the whitelist or blacklist based on its user-defined category.

  NOTE: Local Web filtering does not require a license or a remote category server.

- Enhanced Web filtering—The enhanced Web filtering solution intercepts the HTTP and the HTTPS requests and sends the HTTP URL or the HTTPS source IP to the Websense ThreatSeeker Cloud (TSC). The TSC categorizes the URL into one of the 151 or more categories that are predefined and also provides site reputation information. The TSC further returns the URL category and the site reputation information to the device. The device determines if it can permit or block the request based on the information provided by the TSC.

You can bind either Web filtering profiles or antivirus profiles, or both, to a firewall policy. When both are bound to a firewall policy, Web filtering is applied first, then antivirus is applied. If a URL is blocked by Web filtering, the TCP connection is closed and no antivirus

Copyright © 2016, Juniper Networks, Inc.

scanning is necessary. If a URL is permitted, the content of the transaction is then passed to the antivirus scanning process.

NOTE: Web filtering is applied by TCP port number.

UTM Web filtering supports HTTPS protocol. UTM Web filtering solution uses the IP address of the HTTPS packet to make blacklist, whitelist, permit, or block decisions.

During a block decision, the UTM Web filtering solution does not generate a block page because the clear text is not available for a HTTPS session. However, the solution terminates the session and sends resets to the client and the server for the blocked HTTPS sessions.

UTM Web filtering configuration for HTTP is also applicable for the HTTPS sessions.

NOTE: The sessions-per-client limit CLI command, which imposes a session throttle to prevent a malicious user from generating large amounts of traffic simultaneously, does not support Web filtering.

Related Documentation

- Understanding Integrated Web Filtering on page 5787
- Understanding Redirect Web Filtering on page 5805
- Understanding Enhanced Web Filtering Process on page 5767
- Understanding Local Web Filtering on page 5797
- Monitoring Web Filtering Configurations on page 5815
- **web-filtering on page 5998**

Copyright © 2016, Juniper Networks, Inc.

Enhanced Web Filtering Overview

Enhanced Web Filtering (EWF) with Websense is an integrated URL filtering solution. When you enable the solution on the device, it intercepts the HTTP and the HTTPS requests and sends the HTTP URL or the HTTPS source IP to the Websense ThreatSeeker Cloud (TSC). The TSC categorizes the URL into one of the 95 or more categories that are predefined and also provides site reputation information. The TSC further returns the URL category and the site reputation information to the device. The device determines if it can permit or block the request based on the information provided by the TSC.

EWF supports HTTPS traffic by intercepting HTTPS traffic passing through the SRX Series device. The security channel from the SRX Series device is divided as one SSL channel between the client and the SRX Series device and another SSL channel between the SRX Series device and the HTTPS server. SSL forward proxy acts as the terminal for both channels and forwards the cleartext traffic to the UTM. UTM extracts the URL from the HTTP request message.

You can consider the EWF solution as the next-generation URL filtering solution, building upon the existing SurfControl solution.

Enhanced Web Filtering supports the following HTTP methods:

- GET
- POST
- OPTIONS
- HEAD
- PUT
- DELETE
- TRACE
- CONNECT

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- Understanding Redirect Web Filtering on page 5805
- Understanding Enhanced Web Filtering Process on page 5767
- Example: Configuring Enhanced Web Filtering on page 5772

Copyright © 2016, Juniper Networks, Inc.

Understanding Enhanced Web Filtering Process

Web filtering enables you to manage Internet access and prevent access to inappropriate Web content. This topic describes how the Enhanced Web Filtering (EWF) feature intercepts, scans, and acts upon HTTP or HTTPS traffic.

1. The device creates TCP socket connections to the Websense ThreatSeeker Cloud (TSC).
2. The device intercepts an HTTP or an HTTPS connection and extracts each URL(in the HTTP request) or IP (in the HTTPS request).
3. The device looks for the URL in the user-configured blacklist or whitelist.

   NOTE: **A blacklist or a whitelist action type is a user-defined category in which all the URLs or IP addresses are always blocked or permitted and optionally logged.**

   - If the URL is in the user-configured blacklist, the device blocks the URL.
   - If the URL is in the user-configured whitelist, the device permits the URL.
4. The device checks the user-defined categories and blocks or permits the URL based on the user-specified action for the category.
5. The device looks for the URL in the URL filtering cache.
   - If the URL is not available in the URL filtering cache, the device sends the URL in HTTP format to the TSC with a request for categorization. The device uses one of the connections made available to the TSC to send the request.
   - The TSC responds to the device with the categorization and a reputation score.
6. The device performs the following actions based on the identified category:
   - If the URL is permitted, the device forwards the HTTP request to the HTTP server.
   - If the URL is blocked, the device sends a deny page to the HTTP client and also sends a reset message to the HTTP server to close the connection
   - If the URL is quarantined, the device sends a redirect response to the HTTP client and the URL is redirected to the HTTP server.
   - If the category is not available, the device permits or blocks the URL based on the configured action for the reputation score.
   - If an action for the site reputation score is not configured, the device permits or blocks the URL based on the default action configured in the Web filtering profile.

Copyright © 2016, Juniper Networks, Inc.

Functional Requirements for Enhanced Web Filtering

- **License key**—The Enhanced Web Filtering (EWF) solution builds upon the SurfControl integrated feature on the device. Two different valid license keys are required for the SurfControl integrated solution and for EWF. You need to install a new license to upgrade to the EWF solution.

  NOTE: You can ignore the warning message "requires 'wf_key_websense_ewf' license" because it is generated by routine EWF license validation check.

  A grace period of 30 days, consistent with other UTM features, is provided for the EWF feature after the license key expires.

  NOTE: The device will continue to support the SurfControl integrated solution after the upgrade.

  When the grace period for the EWF feature has passed (or if the feature has not been installed), Web filtering is disabled, all HTTP requests bypass Web filtering, and any connections to the TSC are disabled. When you install a valid license, the connections to the server are established again.

- A **debug** command provides the following information to each TCP connection available on the device:

  - Number of processed requests
  - Number of pending requests
  - Number of errors (dropped or timed-out requests)

- **TCP connection between a Web client and a Webserver**—An App-ID module is used to identify an HTTP connection. The EWF solution identifies an HTTP connection after the device receives the first SYN packet. If an HTTP request has to be blocked, EWF sends a block message from the device to the Web client. EWF further sends a TCP FIN request to the client and a TCP reset (RST) to the server to disable the connection. The device sends all the messages through the flow session. The messages follow the entire service chain.

- **HTTP request interception**—EWF intercepts the first HTTP request on the device and performs URL filtering on all methods defined in HTTP 1.0 and HTTP 1.1. The device holds the original request while waiting for a response from the TSC. If the first packet in the HTTP URL is fragmented or if the device cannot extract the URL for some reason, then the destination IP address is used for the categorization.

  NOTE: For HTTP 1.1 persistent connections, the subsequent requests on that session are ignored by the EWF module.

Copyright © 2016, Juniper Networks, Inc.

If the device holds the original request for a long time, then the client will retransmit the request. The URL filtering code will detect the retransmitted packets. If the original HTTP request has already been forwarded, then EWF forwards the retransmitted packet to the server. However, if EWF is in the middle of first-packet processing or makes the calculation to block the session, then the solution drops the retransmitted packet. A counter tracks the number of retransmitted packets received by the device.

If the TSC does not respond in time to the categorization request from the device, then the original client request is blocked or permitted as per the timeout fallback setting.

- **HTTPS request interception**— EWF intercepts HTTPS traffic passing through the SRX Series device. The security channel from the SRX Series device is divided as one SSL channel between the client and the SRX Series device and another SSL channel between the SRX Series device and the HTTPS server. SSL forward proxy acts as the terminal for both channels and forwards the cleartext traffic to the UTM. UTM extracts the URL from the HTTP request message.
- **Blocking message**—The blocking message sent to the Web client is user-configurable and is of the following types:
  - The Juniper Networks blocking message is the default message defined in the device that can be modified by the user. The default blocking message contains the reason why the request is blocked and the category name (if it is blocked because of a category).
  - Syslog message.

  For example, if you have set the action for Enhanced_Search_Engines_and_Portals to block, and you try to access www.example.com, the blocking message is of the following form: **Juniper Web Filtering:Juniper Web Filtering has been set to block this site. CATEGORY: Enhanced_Search_Engines_and_Portals REASON: BY_PRE_DEFINED** . However, the corresponding syslog message on the DUT is: **WEBFILTER_URL_BLOCKED: WebFilter: ACTION="URL Blocked" 56.56.56.2(59418)->74.125.224.48(80) CATEGORY="Enhanced_Search_Engines_and_Portals" REASON="by predefined category" PROFILE="web-ewf" URL=www.example.com OBJ=/** .
- **Monitoring the Websense server**—The URL filtering module uses two methods to determine if the TSC is active: socket connections and heartbeat. EWF maintains persistent TCP sockets to the TSC. The server responds with a TCP ACK if it is enabled. EWF sends an application layer NOOP keepalive to the TSC. If the device does not receive responses to three consecutive NOOP keepalives in a specific period, it determines the socket to be inactive. The EWF module attempts to open a new connection to the TSC. If all sockets are inactive, the TSC is considered to be inactive. Therefore an error occurs. The error is displayed and logged. Subsequent requests and pending requests are either blocked or passed according to the server connectivity fallback setting until new connections to the TSC are opened again.
- **HTTP protocol communication with the TSC**—EWF uses the HTTP 1.1 protocol to communicate with the TSC. This ensures a persistent connection and transmission of multiple HTTP requests through the same connection. A single HTTP request or response is used for client or server communication. The TSC can handle queued requests; for optimal performance, an asynchronous request or response mechanism

Copyright © 2016, Juniper Networks, Inc.

is used. The requests are sent over TCP, so TCP retransmission is used to ensure request or response delivery. TCP also ensures valid in-order, non-retransmitted HTTP stream data is sent to the HTTP client on the device.

- **Responses**—The responses adhere to the basic HTTP conventions. Successful responses includes a 20x response code (typically 200). An error response include a 4xx or 5xx code. Error responses in the 4xx series indicate issues in the custom code. Error responses in the 5xx series indicate issues with the service.

  Error codes and meanings are as follows:

  - 400–Bad request
  - 403–Forbidden
  - 404–Not found
  - 408–Request canceled or null response
  - 500–Internal server error

  Errors in the 400 series indicate issues with the request. Errors in the 500 series indicate issues with the TSC service. Websense is notified of these errors automatically and responds accordingly.

  You can configure the default fallback setting to determine whether to pass or block the request: **set security utm feature-profile web-filtering juniper-enhanced profile juniper-enhanced fallback-settings default ?**

  The response also contains the site categorization and site reputation information.

- **Categories**—A category list is available on the device. This list consists of categories, each containing a category code, a name, and a parent ID. Categories can also be user-defined. Each category consists of a list of URLs or IP addresses. Categories are not updated dynamically and are tied to the Junos OS release because they have to be compiled into the Junos OS image. Any update in categories needs to be synchronized with the Junos OS release cycle.

- **Caching**—Successfully categorized responses are cached on the device. Uncategorized URLs are not cached. The size of the cache can be configured by the user.

- **Safe search (HTTP support only, not HTTPS)**—A safe-search solution is used to ensure that the embedded objects, such as images on the URLs received from the search engines, are safe and that no undesirable content is returned to the client.

  A URL is provided to the TSC to provide categorization information. If it is a search URL, the TSC also returns a safe-search string. For instance, the safe-search string is **safe=active**. This safe-search string is appended to the URL, and a redirect response for redirecting the client's query with safe search is turned on. This ensures that no unsafe content is returned to the client. If the TSC indicates that it needs to be safe-searched, then you can perform the safe-search redirect.

  For example, the client makes a request to the URL http://images.example.com/images?hl=en&source=imghp&biw=1183&bih=626&q=adult+movies&gbv=2&aq=f&aqi=&aql=&oq=&gs_rfai= No category action is defined for this URL . TSC returns safe-search string **safe=active**. The EWF code on the DUT

Copyright © 2016, Juniper Networks, Inc.

generates a HTTP 302 response, with the redirect URL: http://images.example.com/images?hl=en&source=imghp&biw=1183&bih=626&q=adult+movies&gbv=2&aq=f&aqi=&aql=&oq=&gs_rfai=&safe=active . This response is returned to the client. The client now sends out a safe redirect request to this URL.

NOTE: Safe-search redirect supports HTTP only. You cannot extract the URL for HTTPS. Therefore it is not possible to generate a redirect response for HTTPS search URLs. Safe-search redirects can be disabled by using the CLI option no-safe-search.

- **Site reputation**—The TSC provides site reputation information. Based on these reputations, you can choose a block or a permit action. If the URL is not handled by a whitelist or a blacklist and does not fall in a user or predefined category, then the reputation can be used to perform URL filtering decision.

  The reputation scores are as follows:

  - 100-90—Site is considered very safe.
  - 80-89—Site is considered moderately safe.
  - 70-79—Site is considered fairly safe.
  - 60-69—Site is considered suspicious.
  - 0-59—Site is considered harmful.

  The device maintains a log for URLs that are blocked or permitted based on site reputation scores.

- **Profiles**—A URL filtering profile is defined as a list of categories, with each profile having an action type (permit, log-and-permit, block, quarantine) associated with it. A predefined profile, *junos-wf-enhanced-default*, is provided to users if they choose not to define their own profile.

  You can also define an action based on site reputations in a profile to specify the action when the incoming URL does not belong to any of the categories defined in the profile. If you do not configure the site reputation handling information, then you can define a default action. All URLs that do not have a defined category or defined reputation action in their profile will be blocked, permitted, logged-and-permitted, or quarantined depending on the block or permit handling for the default action explicitly defined in the profile. If you do not specify a default action, then the URLs will be permitted. For search engine requests, if there is no explicit user-defined configuration, and the URL request is without the safe search option, then EWF generates a redirect response and sends it to the client. The client will generate a new search request with the safe-search option enabled.

Copyright © 2016, Juniper Networks, Inc.

NOTE:

A URL filtering profile can contain the following items:

- Multiple user-defined and predefined categories, each with a permit or block action
- Multiple site reputation handling categories, each with a permit or block action
- One default action with a permit or block action

The order of search is blacklist, whitelist, user-defined category, predefined category, safe-search, site reputation, and default action.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- Understanding Redirect Web Filtering on page 5805
- Enhanced Web Filtering Overview on page 5766
- Example: Configuring Enhanced Web Filtering on page 5772

Example: Configuring Enhanced Web Filtering

- Requirements on page 5772
- Overview on page 5772
- Configuration on page 5773
- Verification on page 5779

Requirements

Before you begin, you should be familiar with Web filtering and Enhanced Web Filtering (EWF). See "Web Filtering Overview" on page 5764 and "Understanding Enhanced Web Filtering Process" on page 5767.

Overview

In this example, you configure custom objects and feature profiles.

In the first example configuration, you create a custom object called urllist3 that contains the pattern http://www.example.net 1.2.3.4. The urllist3 custom object is then added to the custom URL category custurl3.

In the second example configuration, you configure the Web filtering feature profile. You set the URL blacklist filtering category to custblacklist, set the whitelist filtering category to custwhitelist, and set the type of Web filtering engine to juniper-enhanced. Then you

Copyright © 2016, Juniper Networks, Inc.

set the cache size parameters for Web filtering to 500 KB and the cache timeout parameters to 1800.

You name the EWF server as rp.cloud.example.com and enter 80 as the port number for communicating with it. (Default port is 80.) Then you create an EWF profile name called junos-wf-enhanced-default.

Next you select a category from the included whitelist and blacklist categories or select a custom URL category list you created for filtering against. Then you enter an action (permit, log and permit, block, or quarantine) to go with the filter. You do this as many times as necessary to compile your whitelists and blacklists and their accompanying actions. This example blocks URLs in the Enhanced_Hacking category. You also specify the action to be taken depending on the site reputation returned for the URL if there is no category match found.

Then you enter a custom message to be sent when HTTP requests are blocked. This example configures the device to send an ***access denied*** message. You select a default action (permit, log and permit, block, or quarantine) for this profile for requests that does not match to any explicitly configured action. This example sets the default action to block. You select fallback settings (block or log and permit) for this profile, in case errors occur in each configured category. This example sets fallback settings to block.

You can also define a redirect URL server so that instead of the device sending a block page with plain text html, the device will send an HTTP 302 redirect to this redirect server with some special variables embedded in the HTTP redirect location field. These special variables can be parsed by the redirect server and serve a special block page to the client with rich images and formatting. The CLI command hierarchy is as follows:

```
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default block-message type custom-redirect-url
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default block-message url http://10.10.121.18
```

NOTE: If you configure the **security utm feature-profile web-filtering juniper-enhanced profile junos-wf-enhanced-default block-message** , then the default block message configuration takes precedence over the **security utm feature-profile web-filtering juniper-enhanced profile junos-wf-enhanced-default custom-block-message** configuration.

Finally, you enter a timeout value in seconds. Once this limit is reached, fail mode settings are applied. The default is 15 seconds, and you can enter a value from 0 through 1800 seconds. This example sets the timeout value to 10. You also disable the safe-search functionality. By default, search requests have safe-search strings attached to them, and a redirect response is sent to ensure that all search requests are safe or strict.

Configuration

- Configuring Enhanced Web Filtering Custom Objects on page 5774
- Configuring the Enhanced Web Filtering Feature Profiles on page 5775

Copyright © 2016, Juniper Networks, Inc.

Configuring Enhanced Web Filtering Custom Objects

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects url-pattern urllist3 value http://www.example.net
set security utm custom-objects url-pattern urllist3 value 1.2.3.4
set security utm custom-objects url-pattern urllistblack value http://www.untrusted.com
set security utm custom-objects url-pattern urllistblack value 13.13.13.13
set security utm custom-objects url-pattern urllistwhite value http://www.trusted.com
set security utm custom-objects url-pattern urllistwhite value 7.7.7.7
set security utm custom-objects custom-url-category custurl3 value urllist3
set security utm custom-objects custom-url-category custblacklist value urllistblack
set security utm custom-objects custom-url-category custwhiltelist value urllistwhite
```

WARNING: A custom category does not take precedence over a predefined category when it has the same name as one of the predefined categories. Do not use the same name for a custom category that you have used for a predefined category.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the CLI User Guide.

To configure integrated Web filtering:

1. Create custom objects and create the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllist3 value [http://www.example.net
     1.2.3.4]
   ```

2. Configure the custom URL category list custom object using the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custurl3 value urllist3
   ```

3. Create a list of untrusted sites.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllistblack value
     [http://www.untrusted.com 13.13.13.13]
   ```

4. Configure the custom URL category list custom object using the URL pattern list of untrusted sites.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custblacklist value urllistblack
   ```

5. Create a list of trusted sites.

   ```
   [edit security utm]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set custom-objects url-pattern urllistwhite value
  [http://www.trusted.com 7.7.7.7]
```

6. Configure the custom URL category list custom object using the URL pattern list of trusted sites.

```
[edit security utm]
user@host# set custom-objects custom-url-category custwhitelist value urllistwhite
```

Results From configuration mode, confirm your configuration by entering the **show security utm custom-objects** command. If the output does not display the intended configuration, repeat the instructions in this example to correct.

```
[edit]
userhost#show security utm custom-objects
  url-pattern {
    urllist3 {
      value [ http://www.example.net ];
    }
    urllistblack {
      value [ http://www.untrusted.com 13.13.13.13 ];
    }
    urllistwhite {
      value [ http://www.trusted.com 7.7.7.7 ];
    }
  }
    custom-url-category {
      custurl3 {
        value urllist3;
      }
      custblacklist {
        value urllistblack;
      }
      custwhiltelist {
        value urllistwhite;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Enhanced Web Filtering Feature Profiles

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile web-filtering url-whitelist custwhitelist
set security utm feature-profile web-filtering url-blacklist custblacklist
set security utm feature-profile web-filtering juniper-enhanced cache size 500
set security utm feature-profile web-filtering juniper-enhanced cache timeout 1800
set security utm feature-profile web-filtering juniper-enhanced server host
  rp.cloud.example.com
set security utm feature-profile web-filtering juniper-enhanced server port 80
set security utm feature-profile web-filtering http-reassemble
```

Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile web-filtering http-persist
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default category Enhanced_Hacking action log-and-permit
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default site-reputation-action very-safe permit
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default site-reputation-action moderately-safe log-and-permit
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default site-reputation-action fairly-safe log-and-permit
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default site-reputation-action harmful block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default site-reputation-action suspicious block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default default block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default custom-block-message "***access denied ***"
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default default block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default fallback-settings server-connectivity block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default fallback-settings timeout block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default fallback-settings too-many-requests block
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default timeout 10
set security utm feature-profile web-filtering juniper-enhanced profile
  junos-wf-enhanced-default no-safe-search
set security utm utm-policy mypolicy web-filtering http-profile my_ewfprofile01
set security policies from-zone utm_clients to-zone mgmt policy 1 then permit
  application-services utm-policy mypolicy
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  quarantine-custom-message “**The requested webpage is blocked by your
  organization's access policy**”.
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  quarantine-message type custom-redirect-url
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  quarantine-message url besgas.spglab.example.net
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the EWF feature profiles:

1. Configure the Web filtering URL blacklist.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-blacklist custblacklist
   ```

2. Configure the Web filtering URL whitelist.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-whitelist custwhitelist
   ```

3. Specify the EWF engine, and set the cache size parameters.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced cache size 500
```

4. Set the cache timeout parameters.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced cache timeout 1800
```

5. Set the server name or IP address.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced server host rp.cloud.example.com
```

6. Enter the port number for communicating with the server.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced server port 80
```

7. Set HTTP fragment reassemble.

```
[edit security utm feature-profile web-filtering]
user@host# set security utm feature-profile web-filtering http-reassemble
```

8. Set HTTP requests in the same session.

```
[edit security utm feature-profile web-filtering]
user@host# set security utm feature-profile web-filtering http-persist
```

9. Create a profile name, and select a category from the included whitelist and blacklist categories.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile junos-wf-enhanced-default category
  Enhanced_Hacking action log-and-permit
```

10. Specify the action to be taken depending on the site reputation returned for the URL if there is no category match found.

```
[edit security utm feature-profile web-filtering]
user@host#set juniper-enhanced profile junos-wf-enhanced-default
  site-reputation-action very-safe permit
user@host#set juniper-enhanced profile junos-wf-enhanced-default
  site-reputation-action moderately-safe log-and-permit
user@host#set juniper-enhanced profile junos-wf-enhanced-default
  site-reputation-action fairly-safe log-and-permit
user@host#set juniper-enhanced profile junos-wf-enhanced-default
  site-reputation-action harmful block
user@host#set juniper-enhanced profile junos-wf-enhanced-default
  site-reputation-action suspicious block
```

11. Enter a custom message to be sent when HTTP requests are blocked.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile junos-wf-enhanced-default
  custom-block-message "***access denied ***"
```

12. Select a default action (permit, log and permit, block, or quarantine) for the profile, when no other explicitly configured action (blacklist, whitelist, custom category, predefined category actions, or site reputation actions) is matched .

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile junos-wf-enhanced-default default block
```

Copyright © 2016, Juniper Networks, Inc.

13. Select fallback settings (block or log and permit) for this profile.

```
[edit security utm feature-profile web-filtering]
set juniper-enhanced profile junos-wf-enhanced-default fallback-settings default
  block
user@host# set juniper-enhanced profile junos-wf-enhanced-default
  fallback-settings server-connectivity block
user@host# set juniper-enhanced profile junos-wf-enhanced-default
  fallback-settings timeout block
set juniper-enhanced profile junos-wf-enhanced-default fallback-settings
  too-many-requests block
```

14. Enter a timeout value in seconds.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile junos-wf-enhanced-default timeout 10
```

15. Disable the safe-search option.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile junos-wf-enhanced-default no-safe-search
```

16. Configure a UTM policy for the Web-filtering HTTP protocol and attach this policy to a security profile to implement it.

```
[edit security utm]
user@host# set utm-policy mypolicy web-filtering http-profile my_ewfprofile01
```

17. Configure a security policy.

```
[edit security]
user@host# set policies from-zone utm_clients to-zone mgmt policy 1 then permit
  application-services utm-policy mypolicy
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile** command. If the output does not display the intended configuration, repeat the instructions in this example to correct.

```
[edit]
user@host# show security utm
feature-profile{
web-filtering {
  url-whitelist custwhitelist;
  url-blacklist custblacklist;
  http-reassemble;
  http-persist;
  type juniper-enhanced;
  juniper-enhanced {
    cache {
      timeout 1800;
      size 500;
    }
    server {
      host rp.cloud.example.com;
      port 80;
    }
    profile junos-wf-enhanced-default {
      category {
```


Copyright © 2016, Juniper Networks, Inc.

```
            Enhanced_Hacking {
              action log-and-permit;
              }
              Enhanced_Government {
              action quarantine;
            }
          }
          site-reputation-action {
            very-safe permit;
            moderately-safe log-and-permit;
            fairly-safe log-and-permit;
            harmful block;
            suspicious block;
          }
          default block;
          custom-block-message "***access denied ***";
          fallback-settings {
            default block;
            server-connectivity block;
            timeout block;
            too-many-requests block;
          }
          timeout 10;
          no-safe-search;
        }
        utm-policy mypolicy {
          web-filtering {
        http-profile my_ewfprofile01;
      }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Status of the Web Filtering Server on page 5779
- Verifying the Increase in Web Filtering Statistics on page 5779

Verifying the Status of the Web Filtering Server

Purpose Verify the Web filtering server status.

Action From the top of the configuration in configuration mode, enter the **show security utm web-filtering status** command.

Verifying the Increase in Web Filtering Statistics

Purpose Verify the increase in Web filtering statistics.

Copyright © 2016, Juniper Networks, Inc.

Action From the top of the configuration in configuration mode, enter the **show security utm web-filtering statistics** command.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Redirect Web Filtering on page 5805
- Enhanced Web Filtering Overview on page 5766
- Understanding Enhanced Web Filtering Process on page 5767

Understanding the Quarantine Action for Enhanced Web Filtering

UTM Enhanced Web Filtering supports block, log-and-permit, and permit actions for HTTP/HTTPS requests. In addition to this, UTM Enhanced Web Filtering now supports the quarantine action which allows or denies access to the blocked site based on the user's response to the message.

The following sequence explains how the HTTP or HTTPs request is intercepted, redirected, and acted upon by the quarantine action:

- The HTTP client requests URL access.
- The device intercepts the HTTP request and sends the extracted URL to the Websense Thread Seeker Cloud (TSC).
- The TSC returns the URL category and the site reputation information to the device.
- If the action configured for the category is quarantine, the device logs the quarantine action and sends a redirect response to HTTP client.
- The URL is sent to the HTTP server for redirecting.
- The device shows a warning message stating that the access to the URL is blocked according to the organization's security policies and prompts the user to respond.
- If the user response is "No," the session is terminated. If the user response is "Yes," the user is allowed access to the site and such access is logged and reported to the administrator.

NOTE: **On all branch SRX Series devices, the quarantine action is supported only for UTM Enhanced Web Filtering or Juniper enhanced type of Web filtering.**

Quarantine Message

The quarantine message sent to the HTTP client is user-configurable and is of the following types:

- Default message

Copyright © 2016, Juniper Networks, Inc.

The default quarantine message is displayed when a user attempts to access a quarantined website and it contains the following information:

- URL name
- Quarantine reason
- Category (if available)
- Site-reputation (if available)

For example, if you have set the action for Enhanced_Search_Engines_and_Portals to quarantine, and you try to access www.search.example.com, the quarantine message is as follows:

*****The requested webpage is blocked by your organization's access policy***.**

- Syslog message.

  The syslog message will be logged by the system when the user access the web page that has already been quarantined and marked as block or permit.

  The corresponding syslog message on the device under test is:

  Jan 25 15:10:40 rodian utmd[3871]: WEBFILTER_URL_BLOCKED: WebFilter: ACTION="URL Blocked" 99.99.99.4(60525)->74.125.224.114(80) CATEGORY="Enhanced_Search_Engines_and_Portals" REASON="by predefined category(quarantine)" PROFILE="ewf-test-profile" URL=www.search.example.com OBJ=/

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- Understanding Redirect Web Filtering on page 5805
- Understanding Enhanced Web Filtering Process on page 5767
- Example: Configuring Enhanced Web Filtering on page 5772
- Example: Configuring Site Reputation Action for Enhanced Web Filtering on page 5781

Example: Configuring Site Reputation Action for Enhanced Web Filtering

This example shows how to configure the site reputation action for both categorized and uncategorized URLs.

- Requirements on page 5782
- Overview on page 5782
- Configuration on page 5782
- Verification on page 5785

Copyright © 2016, Juniper Networks, Inc.

Requirements

Before you begin, you should be familiar with Web Filtering and Enhanced Web Filtering. See "Web Filtering Overview" on page 5764 and "Understanding Enhanced Web Filtering Process" on page 5767.

Overview

In this example, you configure Web Filtering profiles to URLs according to defined categories using the site reputation action. You set the URL whitelist filtering category to url-cat-white and the type of Web Filtering engine to juniper-enhanced. Then you set the cache size parameters for Web Filtering and the cache timeout parameters to 1.

Then you create a juniper-enhanced profile called profile ewf-test-profile, set the URL whitelist category to cust-cat-quarantine, and set the reputation action to quarantine.

You enter a custom message to be sent when HTTP requests are quarantined. In this example, the following message is sent: ***The requested webpage is blocked by your organization's access policy*** .

You block URLs in the Enhanced_News_and_Media category and permit URLs in the Enhanced_Education category. Then you quarantine the URLs in the Enhanced_Streaming_Media category and configure the device to send the following message: ***The requested webpage is blocked by your organization's access policy***.

In this example, you set the default action to permit. You select fallback settings (block or log and permit) for this profile in case errors occur in each configured category. Finally, you set the fallback settings to block.

Configuration

Configuring Site Reputation Action

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile web-filtering set url-whitelist url-cat-white
set security utm feature-profile web-filtering juniper-enhanced cache size
set security utm feature-profile web-filtering juniper-enhanced cache timeout 1
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    category cust-cat-quarantine action quarantine
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    category Enhanced_News_and_Media action block
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    category Enhanced_Education action permit
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    category Enhanced_Education reputation-action harmful block
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    category Enhanced_Streaming_Media action quarantine
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
    default permit
```


Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  default quarantine-message “*** The requested webpage is blocked by your
  organization’s access policy***”.
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  fallback-settings server-connectivity block
set security utm feature-profile web-filtering juniper-enhanced profile ewf-test-profile
  fallback-settings timeout block
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the site reputation action:

1. Configure the Web Filtering URL whitelist.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-whitelist custwhitelist
   ```

2. Specify the Enhanced Web Filtering engine, and set the cache size parameters.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced cache size
   ```

3. Set the cache timeout parameters.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced cache timeout 1
   ```

4. Create a profile name, and select a category from the whitelist categories.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced profile ewf-test-profile category
     cust-cat-quarantine action quarantine
   ```

5. Create a profile name, and select a category from the whitelist categories.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced profile ewf-test-profile category
     Enhanced_News_and_Media action block
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced profile ewf-test-profile category
     Enhanced_Education action permit
   user@host# set juniper-enhanced profile ewf-test-profile category
     Enhanced_Education action harmful block
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced profile ewf-test-profile category
     Enhanced_Streaming_Media action quarantine
   ```

6. Enter a warning message to be sent when HTTP requests are quarantined.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-enhanced profile ewf-test-profile
     quarantine-custom-message "***The requested webpage is blocked by your
     organization's access policy ***"
   ```

7. Select a default action (permit, log and permit, block, or quarantine) for the profile, when no other explicitly configured action (blacklist, whitelist, custom category, predefined category or site reputation) is matched .

Copyright © 2016, Juniper Networks, Inc.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile ewf-test-profile default permit
```

8. Select fallback settings (block or log and permit) for this profile.

```
[edit security utm feature-profile web-filtering]
user@host# set juniper-enhanced profile ewf-test-profile fallback-settings
  server-connectivity block
user@host# set juniper-enhanced profile ewf-test-profile fallback-settings timeout
  block
```

Results From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the instructions in this example to correct.

```
[edit]
user@host# show security utm
feature-profile{
web-filtering {
  url-whitelist url-cat-white;
  type juniper-enhanced;
  traceoptions;
  flag all;
}
  juniper-enhanced {
    cache {
      timeout 1
    }
    profile ewf-test-profile {
      category {
        cust-cat-quarantine {
          action quarantine;
          }
          Enhanced_News_and_Media {
            action block;
            reputation-action;
          }
          Enhanced_Education {
          action permit;
          reputation-action;
          {
          harmful block;
          }
          }
          Enhanced_Streaming_Media {
          action quarantine;
          }
        }
      default permit;
    quarantine-custom-message "***The requested webpage is blocked by your
      organization's access policy***".
      fallback-settings {
        server-connectivity block;
        timeout block;
        }
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying the Status of UTM Service on page 5785
- Verifying the Status of UTM Session on page 5785
- Verifying the Status of UTM Web Filtering on page 5785
- Verifying the Statistics of UTM Web Filtering on page 5786

Verifying the Status of UTM Service

Purpose Verify the UTM service status.

Action From operational mode, enter the **show security utm status** command.

Sample Output

```
user@host>show security utm status
UTM service status: Running
```

Verifying the Status of UTM Session

Purpose Verify the UTM session status.

Action From operational mode, enter the **show security utm session** command.

Sample Output

```
user@host>show security utm session
UTM session info:
 Maximum sessions:                4000
 Total allocated sessions:        0
 Total freed sessions:            0
 Active sessions:                 0
```

Verifying the Status of UTM Web Filtering

Purpose Verify the UTM Web filtering status.

Action From operational mode, enter the **show security utm web-filtering status** command.

Sample Output

```
user@host>show security utm web-filtering status
UTM web-filtering status:
    Server status: Juniper Enhanced using Websense server UP
```

Copyright © 2016, Juniper Networks, Inc.

Verifying the Statistics of UTM Web Filtering

Purpose Verify the Web filtering statistics for connections including whitelist and blacklist hits and custom category hits.

Action From operational mode, enter the **show security utm web-filtering statistics** command.

Sample Output

```
user@host>show security utm web-filtering statistics
UTM web-filtering statistics:
    Total requests:                     2594
    white list hit:                     0
    Black list hit:                     0
    Queries to server:                  2407
    Server reply permit:                1829
    Server reply block:                 0
    Server reply quarantine:            517
    Server reply quarantine block:      0
    Server reply quarantine permit:     8
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
    Cache hit permit:                   41
    Cache hit block:                    0
    Cache hit quarantine:               144
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        1
    Safe-search redirect:               0
    Web-filtering sessions in total:    16000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit          block
          Default                                0              0
          Timeout                                0              0
     Connectivity                                0              1
Too-many-requests                                0              0
```

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Redirect Web Filtering on page 5805
- Enhanced Web Filtering Overview on page 5766
- Understanding Enhanced Web Filtering Process on page 5767
- Example: Configuring Enhanced Web Filtering on page 5772
- Understanding the Quarantine Action for Enhanced Web Filtering on page 5780
- **web-filtering on page 5998**

Copyright © 2016, Juniper Networks, Inc.

Understanding Integrated Web Filtering

With integrated Web filtering, the firewall intercepts every HTTP request in a TCP connection and extracts the URL from the HTTP request. Each individual HTTP request is blocked or permitted based on URL filtering profiles defined by you. The decision making is done on the device after it identifies a category for a URL.

A URL category is a list of URLs grouped by content. URL categories are predefined and maintained by SurfControl or are defined by you. SurfControl maintains about 40 predefined categories. When defining your own URL categories, you can group URLs and create categories specific to your needs.

You define your own categories using URL pattern list and custom URL category list custom objects. Once defined, you can select your categories when you configure your Web filtering profile. Each category can have a maximum of 20 URLs. When you create a category, you can add either the URL or the IP address of a site. When you add a URL to a user-defined category, the device performs DNS lookup, resolves the host name into IP addresses, and caches this information. When a user tries to access a site with the IP address of the site, the device checks the cached list of IP addresses and tries to resolve the hostname. Many sites have dynamic IP addresses, meaning that their IP addresses change periodically. A user attempting to access a site can type an IP address that is not in the cached list on the device. Therefore, if you know the IP addresses of sites you are adding to a category, enter both the URL and the IP address(es) of the site.

NOTE: If a URL appears in both a user-defined category and a predefined category, the device matches the URL to the user-defined category.

NOTE: Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1.

This topic contains the following sections:

- Integrated Web Filtering Process on page 5787
- Integrated Web Filtering Cache on page 5788
- Integrated Web Filtering Profiles on page 5788
- Profile Matching Precedence on page 5789

Integrated Web Filtering Process

This is a general description of how Web traffic is intercepted and acted upon by the Web filtering module.

1. The device intercepts a TCP connection.
2. The device intercepts each HTTP request in the TCP connection.
3. The device extracts each URL in the HTTP request and checks its URL filter cache.

Copyright © 2016, Juniper Networks, Inc.

4. Global Web filtering whitelists and blacklists are checked first for block or permit.
5. If the HTTP request URL is allowed based on cached parameters, it is forwarded to the webserver. If there is no cache match, a request for categorization is sent to the SurfControl server. (If the HTTP request URL is blocked, the request is not forwarded and a notification message is logged.)
6. In the allowed case, the SurfControl server responds with the corresponding category.
7. Based on the identified category, if the URL is permitted, the device forwards the HTTP request to the webserver. If the URL is not permitted, then a deny page is sent to the HTTP client.

Integrated Web Filtering Cache

By default, the device retrieves and caches the URL categories from the SurfControl CPA server. This process reduces the overhead of accessing the SurfControl CPA server each time the device receives a new request for previously requested URLs. You can configure the size and duration of the cache, according to the performance and memory requirements of your networking environment. The lifetime of cached items is configurable between 1 and 1800 seconds with a default value of 300 seconds.

NOTE: **Caches are not preserved across device reboots or power losses.**

Integrated Web Filtering Profiles

You configure Web filtering profiles that permit or block URLs according to defined categories. A Web filtering profile consists of a group of URL categories assigned one of the following actions:

- Permit — The device always allows access to the websites in this category.
- Block — The device blocks access to the websites in this category. When the device blocks access to this category of websites, it displays a message in your browser indicating the URL category.
- Blacklist — The device always blocks access to the websites in this list. You can create a user-defined category.
- Whitelist — The device always allows access to the websites in this list. You can create a user-defined category.

NOTE: **A predefined profile is provided and can be used if you choose not to define your own profile.**

A Web filtering profile may contain one blacklist or one whitelist, multiple user-defined and/or predefined categories each with a permit or block action, and an *Other* category with a permit or block action. You can define an action for all *Other* categories in a profile to specify what to do when the incoming URL does not belong to any of the categories defined in the profile. If the action for the *Other* category is block, the incoming URL is

Copyright © 2016, Juniper Networks, Inc.

blocked if it does not match any of the categories explicitly defined in the profile. If an action for the *Other* category is not specified, the default action of permit is applied to the incoming URL not matching any category.

Profile Matching Precedence

When a profile employs several categories for URL matching, those categories are checked for matches in the following order:

1. If present, the global blacklist is checked first. If a match is made, the URL is blocked. If no match is found...
2. The global whitelist is checked next. If a match is made, the URL is permitted. If no match is found...
3. User-defined categories are checked next. If a match is made, the URL is blocked or permitted as specified. If no match is found...
4. Predefined categories are checked next. If a match is made, the URL is blocked or permitted as specified. If no match is found...
5. The Other category is checked next. If a match is made, the URL is blocked or permitted as specified.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Redirect Web Filtering on page 5805
- Understanding Local Web Filtering on page 5797
- Example: Configuring Integrated Web Filtering on page 5789

Example: Configuring Integrated Web Filtering

This example shows how to configure integrated Web filtering.

- Requirements on page 5789
- Overview on page 5789
- Configuration on page 5790
- Verification on page 5796

Requirements

Before you begin, learn more about Web filtering. See "Web Filtering Overview" on page 5764.

Overview

In this example you configure integrated Web filtering custom objects, integrated Web filtering feature profiles, and integrated Web filtering UTM policies. You also attach integrated Web filtering UTM policies to security policies.

Copyright © 2016, Juniper Networks, Inc.

In the first example configuration you create a custom object called urllist3 that contains the pattern http://www.example.net 1.2.3.4. The urllist3 custom object is then added to the custom URL category custurl3.

In the second example configuration, you configure the Web filtering feature profile. You set the URL blacklist filtering category to custblacklist, set the whitelist filtering category to custwhitelist and the type of Web filtering engine to surf-control-integrated. Then you set the cache size parameters for Web filtering to 500 KB, which is the default, and the cache timeout parameters to 1800.

You name the Surf Control server as surfcontrolserver and enter 8080 as the port number for communicating with it. (Default ports are 80, 8080, and 8081.) Then you create a surf-control-integrated profile name called surfprofile1.

Next you select a category from the included whitelist and blacklist categories or select a custom URL category list you created for filtering against. Then you enter an action (permit, log and permit, block) to go with the filter. You do this as many times as necessary to compile your whitelists and blacklists and their accompanying actions. This example blocks URLs in the custurl3 category.

Then you enter a custom message to be sent when HTTP requests are blocked. This example configures the device to send an ***access denied*** message. You select a default action (permit, log and permit, block) for this profile for requests that experience errors. This example sets the default action to block. You select fallback settings (block or log and permit) for this profile, in case errors occur in each configured category. This example sets fallback settings to block.

Finally, you enter a timeout value in seconds. Once this limit is reached, fail mode settings are applied. The default is 10 seconds, and you can enter a value from 10 to 240 seconds. This example sets the timeout value to 10.

In the third example configuration, you create UTM policy utmp5 and attach it to profile surfprofile1.

In the final example configuration, you attach the UTM policy utmp5 to the security policy p5.

Configuration

- Configuring Integrated Web Filtering Custom Objects on page 5790
- Configuring the Integrated Web Filtering Feature Profiles on page 5792
- Configuring Integrated Web Filtering UTM Policies on page 5795
- Attaching Integrated Web Filtering UTM Policies to Security Policies on page 5795

Configuring Integrated Web Filtering Custom Objects

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects url-pattern urllist3 value http://www.example.net
```


Copyright © 2016, Juniper Networks, Inc.

```
set security utm custom-objects url-pattern urllist3 value 1.2.3.4
set security utm custom-objects url-pattern urllistblack value http://www.untrusted.com
set security utm custom-objects url-pattern urllistblack value 13.13.13.13
set security utm custom-objects url-pattern urllistwhite value http://www.trusted.com
set security utm custom-objects url-pattern urllistwhite value 7.7.7.7
set security utm custom-objects custom-url-category custurl3 value urllist3
set security utm custom-objects custom-url-category custblacklist value urllistblack
set security utm custom-objects custom-url-category custwhiltelist value urllistwhite
```

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Step-by-Step Procedure

To configure integrated Web filtering:

1. Create custom objects and create the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllist3 value [http://www.example.net
     1.2.3.4]
   ```

2. Configure the custom URL category list custom object using the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custurl3 value urllist3
   ```

3. Create a list of untrusted sites

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllistblack value
     [http://www.untrusted.com 13.13.13.13]
   ```

4. Configure the custom URL category list custom object using the URL pattern list of untrusted sites.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custblacklist value urllistblack
   ```

5. Create a list of trusted sites.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllistwhite value
     [http://www.trusted.com 7.7.7.7]
   ```

6. Configure the custom URL category list custom object using the URL pattern list of trusted sites.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custwhitelist value urllistwhite
   ```

Results

From configuration mode, confirm your configuration by entering the **show security utm custom-objects** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

For brevity, this show command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
userhost#show security utm custom-objects
  url-pattern {
    urllist3 {
      value [ http://www.example.net ];
    }
    urllistblack {
      value [ http://www.untrusted.com 13.13.13.13 ];
    }
    urllistwhite {
      value [ http://www.trusted.com 7.7.7.7 ];
    }
  }
    custom-url-category {
      custurl3 {
        value urllist3;
      }
      custblacklist {
        value urllistblack;
      }
      custwhiltelist {
        value urllistwhite;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Integrated Web Filtering Feature Profiles

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile web-filtering url-whitelist custwhitelist
set security utm feature-profile web-filtering url-blacklist custblacklist
set security utm feature-profile web-filtering surf-control-integrated cache timeout 1800
set security utm feature-profile web-filtering surf-control-integrated cache size 500
set security utm feature-profile web-filtering surf-control-integrated server host
  surfcontrolserver
set security utm feature-profile web-filtering surf-control-integrated server port 8080
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
  category custurl3 action block
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
  default block
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
  custom-block-message "***access denied ***"
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
  fallback-settings default block
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
  fallback-settings server-connectivity block
```


Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
    fallback-settings timeout block
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
    fallback-settings too-many-requests block
set security utm feature-profile web-filtering surf-control-integrated profile surfprofile1
    timeout 10
set security utm feature-profile content-filtering profile contentfilter1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure integrated Web filtering feature profiles:

1. Configure the Web filtering URL Black List.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-blacklist custblacklist
   ```

2. Configure the Web filtering URL White List.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-whitelist custwhitelist
   ```

3. Specify the surf-control-integrated Web filtering engine and set the cache size parameters.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated cache size 500
   ```

4. Set the cache timeout parameters.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated cache timeout 1800
   ```

5. Set the server name or IP address.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated server host surfcontrolserver
   ```

6. Enter the port number for communicating with the server.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated server port 8080
   ```

7. Create a profile name and select a category from the included whitelist and blacklist categories.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated profile surfprofile1 category custurl3 action
       block
   ```

8. Enter a custom message to be sent when HTTP requests are blocked.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set surf-control-integrated profile surfprofile1 custom-block-message
       "***access denied***"
   ```

9. Select a default action (permit, log and permit, block) for this profile for requests that experience errors.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security utm feature-profile web-filtering]
user@host# set surf-control-integrated profile surfprofile1 default block
```

10. Select fallback settings (block or log and permit) for this profile.

```
[edit security utm feature-profile web-filtering]
user@host# set surf-control-integrated profile surfprofile1 fallback-settings default
  block

user@host# set surf-control-integrated profile surfprofile1 fallback-settings
  server-connectivity block
user@host# set surf-control-integrated profile surfprofile1 fallback-settings timeout
  block
user@host# set surf-control-integrated profile surfprofile1 fallback-settings
  too-many-requests block
```

11. Enter a timeout value, in seconds.

```
[edit security utm feature-profile web-filtering]
user@host# set surf-control-integrated profile surfprofile1 timeout 10
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security utm feature-profile
  web-filtering {
    url-whitelist custwhitelist;
    url-blacklist custblacklist;
  type juniper-local;
    surf-control-integrated {
      cache {
        timeout 1800;
        size 500;
      }
      server {
        host surfcontrolserver;
        port 8080;
      }
      profile surfprofile1 {
        category {
          custurl3 {
            action block;
          }
        }
        default block;
        custom-block-message "***access denied ***";
        fallback-settings {
          default block;
          server-connectivity block;
          timeout block;
          too-many-requests block;
        }
        timeout 10;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
content-filtering {
  profile contentfilter1;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Integrated Web Filtering UTM Policies

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm utm-policy utmp5 web-filtering http-profile surfprofile1
```

Step-by-Step Procedure

To configure a UTM policy:

1. Create the UTM policy referencing a profile.

```
[edit]
user@host# set security utm utm-policy utmp5 web-filtering http-profile surfprofile1
```

Results

From configuration mode, confirm your configuration by entering the **show security utm utm-policy** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security utm utm-policy
...
  utm-policy utmp5 {
    content-filtering {
      http-profile contentfilter1;
    }
    web-filtering {
      http-profile surfprofile1;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Attaching Integrated Web Filtering UTM Policies to Security Policies

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone untrust policy p5 match source-address any
set security policies from-zone trust to-zone untrust policy p5 match destination-address
  any
set security policies from-zone trust to-zone untrust policy p5 match application junos-http
set security policies from-zone trust to-zone untrust policy p5 then permit
  application-services utm-policy utmp5
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To attach a UTM policy to a security policy:

1. Create and configure the security policy.

```
[edit security policies from-zone trust to-zone untrust policy p5]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application junos-http
```

2. Attach the UTM policy to the security policy.

```
[edit security policies from-zone trust to-zone untrust policy p5]
user@host# set then permit application-services utm-policy utmp5
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security policies
  from-zone trust to-zone untrust {
    policy p5 {
      match {
        source-address any;
        destination-address any;
        application junos-http;
      }
      then {
        permit {
          application-services {
            utm-policy utmp5;
          }
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Configuration of Integrated Web Filtering Custom Objects on page 5796
- Verifying the Configuration of Integrated Web Filtering Feature Profiles on page 5797
- Verifying the Configuration of Integrated Web Filtering UTM Policies on page 5797
- Verifying the Attachment of Integrated Web Filtering UTM Policies to Security Policies on page 5797

Verifying the Configuration of Integrated Web Filtering Custom Objects

Purpose

Verify the configuration of integrated Web filtering custom objects.

Copyright © 2016, Juniper Networks, Inc.

Action From the top of the configuration in configuration mode, enter the **show security utm custom-objects** command.

Verifying the Configuration of Integrated Web Filtering Feature Profiles

Purpose Verify the configuration of integrated Web filtering feature profiles.

Action From the top of the configuration in configuration mode, enter the **show security utm feature-profile** command.

Verifying the Configuration of Integrated Web Filtering UTM Policies

Purpose Verify the configuration of integrated Web filtering UTM policies.

Action From the top of the configuration in configuration mode, enter the **show security utm** command.

Verifying the Attachment of Integrated Web Filtering UTM Policies to Security Policies

Purpose Verify the attachment of integrated Web filtering UTM policies to security policies.

Action From the top of the configuration in configuration mode, enter the **show security policies** command.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Redirect Web Filtering on page 5805
- Understanding Local Web Filtering on page 5797
- Understanding Integrated Web Filtering on page 5787
- **web-filtering on page 5998**

Understanding Local Web Filtering

With local Web filtering, the firewall intercepts every HTTP request in a TCP connection and extracts the URL. The decision making is done on the device after it looks up a URL to determine if it is in the whitelist or blacklist based on its user-defined category. If the URL is in the url-blacklist, the request is blocked; if it's in the url-whitelist, the request is permitted. If the URL is not in either list, the defined default action will occur (block, log-and-permit, or permit). You can permit or block access to a requested site by binding a Web filtering profile to a firewall policy. Local Web filtering provides basic Web filtering without requiring an additional license or external category server.

Copyright © 2016, Juniper Networks, Inc.

This topic contains the following sections:

- User-Defined URL Categories on page 5798
- Local Web Filtering Process on page 5798
- Local Web Filtering Profiles on page 5798
- Profile Matching Precedence on page 5799

User-Defined URL Categories

When defining your own URL categories, you can group URLs and create categories specific to your needs. Each category can have a maximum of 20 URLs. When you create a category, you can add either the URL or the IP address of a site. When you add a URL to a user-defined category, the device performs DNS lookup, resolves the hostname into IP addresses, and caches this information. When a user tries to access a site with the IP address of the site, the device checks the cached list of IP addresses and tries to resolve the hostname. Many sites have dynamic IP addresses, meaning that their IP addresses change periodically. A user attempting to access a site can type an IP address that is not in the cached list on the device. Therefore, if you know the IP addresses of sites you are adding to a category, enter both the URL and the IP address(es) of the site.

You define your own categories using URL pattern list and custom URL category list custom objects. Once defined, you assign your categories to the global user-defined url-blacklist (block) or url-whitelist (permit) categories.

NOTE: **Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1.**

Local Web Filtering Process

This is a general description of how Web traffic is intercepted and acted upon by the Web filtering module.

1. The device intercepts a TCP connection.
2. The device intercepts each HTTP request in the TCP connection.
3. The device extracts each URL in the HTTP request and checks its URL against the user-defined whitelist and blacklist.
4. If the URL is found in the blacklist, the request is not permitted and a deny page is sent to the http client. If the URL is found in the whitelist, the request is permitted.
5. If the URL is not found in the whitelist or blacklist, the configured default fallback action is applied. If no fallback action is defined, then the request is permitted.

Local Web Filtering Profiles

You configure Web filtering profiles that permit or block URLs according to defined custom categories. A Web filtering profile consists of a group of URL categories assigned one of the following actions:

Copyright © 2016, Juniper Networks, Inc.

- Blacklist — The device always blocks access to the websites in this list. Only user-defined categories are used with local Web filtering.
- Whitelist — The device always allows access to the websites in this list. Only user-defined categories are used with local Web filtering.

A Web filtering profile can contain one blacklist or one whitelist with multiple user-defined categories each with a permit or block action. You can define a default fallback action when the incoming URL does not belong to any of the categories defined in the profile. If the action for the default category is block, the incoming URL is blocked if it does not match any of the categories explicitly defined in the profile. If an action for the default action is not specified, the default action of permit is applied to the incoming URL not matching any category.

Profile Matching Precedence

When a profile employs several categories for URL matching, those categories are checked for matches in the following order:

1. If present, the global blacklist is checked first. If a match is made, the URL is blocked. If no match is found...
2. The global whitelist is checked next. If a match is made, the URL is permitted. If no match is found...
3. User-defined categories are checked next. If a match is made, the URL is blocked or permitted as specified.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Redirect Web Filtering on page 5805
- Example: Configuring Local Web Filtering on page 5799

Example: Configuring Local Web Filtering

This example shows how to configure local Web filtering.

- Requirements on page 5799
- Overview on page 5800
- Configuration on page 5800
- Verification on page 5804

Requirements

Before you begin, learn more about Web filtering. See "Web Filtering Overview" on page 5764.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example you configure local Web filtering custom objects, local Web filtering feature profiles, and local Web filtering UTM policies. You also attach local Web filtering UTM policies to security policies.

In the first example configuration you create custom objects called urllist5 and urllist6 that contains the patterns http://www.example.net 1.2.3.4 and http://www.example.com 1.2.3.4 respectively. The urllist5 and urllist6 custom objects are then added to the custom URL category custurl5 and custurl6.

In the second example configuration, you configure the Web filtering feature profile. You set the URL blacklist filtering category to custurl4 and URL whitelist filtering category to custurl3. You set the type of Web filtering engine to juniper-local.

Then you create a juniper-local profile name called localprofile1. You select a default action (permit, log and permit, block) for this profile for requests that experience errors. This example sets the default action to permit.

Then you enter a custom message to be sent when HTTP requests are blocked. This example configures the device to send an ***Access to this site is not permitted*** message. You select fallback settings (block or log and permit) for this profile, in case errors occur in each configured category. This example sets fallback settings to block.

In the third example configuration, you create UTM policy utmp5 and attach it to profile localprofile1.

In the final example configuration, you attach the UTM policy utmp5 to the security policy p5.

Configuration

- Configuring Local Web Filtering Custom Objects on page 5800
- Configuring the Local Web Filtering Feature Profiles on page 5801
- Configuring Local Web Filtering UTM Policies on page 5803
- Attaching Local Web Filtering UTM Policies to Security Policies on page 5803

Configuring Local Web Filtering Custom Objects

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects url-pattern urllist5 value http://www.example.net
set security utm custom-objects url-pattern urllist5 value 1.2.3.4
set security utm custom-objects url-pattern urllist6 value http:// www.example.com
set security utm custom-objects url-pattern urllist6 value 1.2.3.4
set security utm custom-objects custom-url-category custurl5 value urllist5
set security utm custom-objects custom-url-category custurl6 value urllist6
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

To configure local Web filtering using the CLI:

1. Create custom objects and URL pattern lists.

```
[edit]
user@host# set security utm custom-objects url-pattern urllist5 value
  [http://www.example.net 1.2.3.4]
user@host# set security utm custom-objects url-pattern urllist6 value
  [http://www.example.com 1.2.3.4]
```

2. Configure the custom URL category list custom object using the URL pattern list.

```
[edit]
user@host# set security utm custom-objects custom-url-category custurl5 value
  urllist5
user@host# set security utm custom-objects custom-url-category custurl6 value
  urllist6
```

Results

From configuration mode, confirm your configuration by entering the **show security utm custom-objects** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security utm custom-objects
  url-pattern {
    urllist5 {
      value [ http://www.example.net 1.2.3.4 ];
    }
    urllist6 {
      value [ http://www.example.com 1.2.3.4 ];
    }
  }
    custom-url-category {
      custurl5 {
        value urllist5;
      }
      custurl6 {
        value urllist6;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Local Web Filtering Feature Profiles

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile web-filtering url-whitelist custurl3
set security utm feature-profile web-filtering url-blacklist custurl4
set security utm feature-profile web-filtering type juniper-local
set security utm feature-profile web-filtering juniper-local profile localprofile1 default
  permit
```

Copyright © 2016, Juniper Networks, Inc.

```
set security utm feature-profile web-filtering juniper-local profile localprofile1
  custom-block-message "Access to this site is not permitted."
set security utm feature-profile web-filtering juniper-local profile localprofile1
  fallback-settings default block
set security utm feature-profile web-filtering juniper-local profile localprofile1
  fallback-settings too-many-requests block
set security utm feature-profile content-filtering profile contentfilter1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure local Web filtering feature profiles:

1. Configure the Web filtering feature profiles.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-whitelist custurl3
   user@host# set url-blacklist custurl4
   ```

2. Select the Web filtering engine.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set type juniper-local
   ```

3. Select a default action (permit, log and permit, block) for this profile for requests that experience errors.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper–local profile localprofile1 default permit
   ```

4. Enter a custom message to be sent when HTTP requests are blocked.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper-local profile localprofile1 custom-block-message “Access
     to this site is not permitted”
   ```

5. Select fallback settings (block or log and permit) for this profile.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set juniper–local profile localprofile1 fallback-settings default block
   set security utm feature-profile web-filtering juniper–local profile localprofile1
     fallback-settings too-many-requests block
   ```

Results

From configuration mode, confirm your configuration by entering the **show security utm feature-profile** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security utm feature-profile
web-filtering {
  url-whitelist custurl3;
  url-blacklist custurl4;
  type juniper-local;
  juniper-local {
    profile localprofile1 {
      default permit;
      custom-block-message "Access to this site is not permitted.";
```


Copyright © 2016, Juniper Networks, Inc.

```
            fallback-settings {
              default block;
              too-many-requests block;
            }
          }
        }
      }
        content-filtering {
          profile contentfilter1;
        }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Local Web Filtering UTM Policies

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
et security utm utm-policy utmp5 web-filtering http-profile localprofile1
```

Step-by-Step Procedure

To configure a UTM policy:

1. Create the UTM policy referencing a profile.

```
[edit]
user@host# set security utm utm-policy utmp5 web-filtering http-profile localprofile1
```

Results

From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this show command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
userhost#show security utm
utm-policy utmp5 {
  web-filtering {
    http-profile localprofile1;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Attaching Local Web Filtering UTM Policies to Security Policies

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone trust to-zone untrust policy p5 match source-address any
set security policies from-zone trust to-zone untrust policy p5 match destination-address
    any
set security policies from-zone trust to-zone untrust policy p5 match application junos-http
set security policies from-zone trust to-zone untrust policy p5 then permit
    application-services utm-policy utmp5
```

Step-by-Step Procedure

To attach a UTM policy to a security policy:

1. Create and configure the security policy.

```
[edit security policies from-zone trust to-zone untrust policy p5]
user@host# set match source-address any
user@host# set match destination-address any
user@host# set match application junos-http
```

2. Attach the UTM policy to the security policy.

```
[edit security policies from-zone trust to-zone untrust policy p5]
user@host# set then permit application-services utm-policy utmp5
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost#show security policies
  from-zone trust to-zone untrust {
    policy p5 {
      match {
        source-address any;
        destination-address any;
        application junos-http;
      }
      then {
        permit {
          application-services {
            utm-policy utmp5;
          }
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Configuration of Local Web Filtering Custom Objects on page 5805
- Verifying the Configuration of Local Web Filtering Feature Profiles on page 5805

Copyright © 2016, Juniper Networks, Inc.

- Verifying the Configuration of Local Web Filtering UTM Policies on page 5805
- Verifying the Attachment of Local Web Filtering UTM Policies to Security Policies on page 5805

Verifying the Configuration of Local Web Filtering Custom Objects

Purpose Verify the configuration of local Web filtering custom objects.

Action From operational mode, enter the **show security utm custom-objects** command.

Verifying the Configuration of Local Web Filtering Feature Profiles

Purpose Verify the configuration of local Web filtering feature profiles.

Action From operational mode, enter the **show security utm feature-profile** command.

Verifying the Configuration of Local Web Filtering UTM Policies

Purpose Verify the configuration of local Web filtering UTM policies.

Action From operational mode, enter the **show security utm** command.

Verifying the Attachment of Local Web Filtering UTM Policies to Security Policies

Purpose Verify the attachment of local Web filtering UTM policies to security policies.

Action From operational mode, enter the **show security policies** command.

Related Documentation

- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- Monitoring Web Filtering Configurations on page 5815
- **web-filtering on page 5998**

Understanding Redirect Web Filtering

With redirect Web filtering, the Web filtering module intercepts an HTTP request. The URL in the request is then sent to the external Websense server, which makes a permit or a deny decision. If access is permitted to the URL in question, the original HTTP request and all the subsequent requests are sent to the intended HTTP server. But if access is denied to the URL in question, a blocking message is sent to the client.

This is a general description of how Web traffic is intercepted, redirected, and acted upon by the Web filtering module:

1. A Web client establishes a TCP connection with the webserver.
2. The Web client then sends an HTTP request.

Copyright © 2016, Juniper Networks, Inc.

3. The device intercepts the requests and extract URL. The URL is checked against Global Web filtering whitelists and blacklists. If no match is made, the Websense server configuration parameters are utilized. Otherwise go to step 6.
4. The URL is sent to the Websense server for checking,
5. The Websense server returns a response indicating whether or not the URL is to be permitted or blocked.
6. If access is allowed, then the original HTTP request is sent to the webserver. If access is denied, the device sends a blocking message to the client and tears down the TCP connection.

NOTE: Web filtering is performed on all the methods defined in HTTP 1.0 and HTTP 1.1. However, redirect Web filtering uses destination IP as URL when it is checking HTTPS traffic.

NOTE: Decision making from real-time options provides a higher level of accuracy, therefore caching for redirect Web filtering is not supported.

NOTE: Redirect Web filtering does not require a subscription license.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects on page 5806

Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects

This example shows how to manage Internet usage by configuring redirect Web filtering using custom objects and preventing access to inappropriate Web content.

- Requirements on page 5806
- Overview on page 5807
- Configuration on page 5808
- Verification on page 5813

Requirements

Before you begin, learn more about Web filtering. See "Web Filtering Overview" on page 5764.

Copyright © 2016, Juniper Networks, Inc.

Overview

The benefit of using Web filtering is that it extracts the URLs from HTTP request messages and performs filtering according to the requirements. The advantage of configuring redirect Web filtering is that it extracts the URLs from the HTTP requests and sends them to an external URL filtering server to determine whether to allow or deny access.

In this example you configure redirect Web filtering custom objects, redirect Web filtering feature profiles, and redirect Web filtering UTM policies. You also attach redirect Web filtering UTM policies to security policies.

The default websense-redirect server port number is 15868.

You select fallback settings (block or log and permit) for this profile, in case errors occur in each configured category. This example sets fallback settings to block the profile. You enter the number of sockets used for communicating between the client and the server. The default is 32 for SRX Series devices.

Finally, you enter a timeout value in seconds. Once this limit is reached, fail mode settings are applied. The default is 15 seconds, and you can enter a value from 1 to 1800 seconds. This example sets the timeout value to 10.

Copyright © 2016, Juniper Networks, Inc.

Figure 239 shows the overall architecture for the Websense redirect feature.

Figure 239: Websense Redirect Architecture

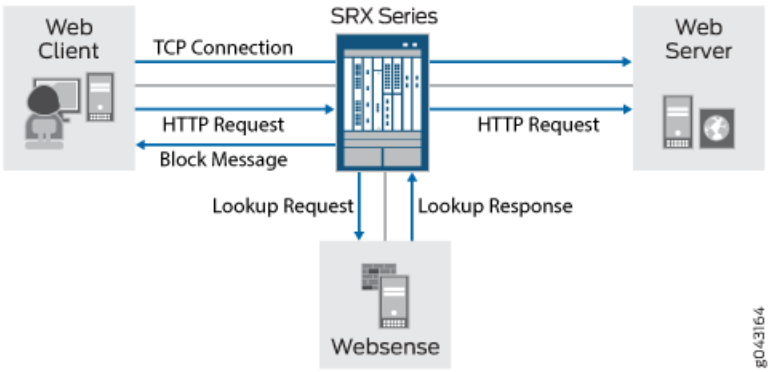

Configuration

- Configuring Redirect Web Filtering Custom Objects on page 5808
- Configuring the Redirect Web Filtering Feature Profiles on page 5810
- Configuring Redirect Web Filtering UTM Policies and Attaching the Redirect Web Filtering UTM Policies to Security Policies on page 5811

Configuring Redirect Web Filtering Custom Objects

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm custom-objects url-pattern urllist4 value 1.2.3.4
set security utm custom-objects url-pattern urllistblack value http://www.untrusted.com
set security utm custom-objects url-pattern urllistblack value 13.13.13.13
set security utm custom-objects url-pattern urllistwhite value http://www.trusted.com
set security utm custom-objects url-pattern urllistwhite value 7.7.7.7
set security utm custom-objects custom-url-category custurl4 value urllist4
set security utm custom-objects custom-url-category custblacklist value urllistblack
set security utm custom-objects custom-url-category custwhitelist value urllistwhite
```

Step-by-Step Procedure

To configure redirect Web filtering custom objects:

1. Create custom objects and create the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects url-pattern urllist4 value [http://www.example.net
     1.2.3.4]
   ```

2. Configure the custom URL category list custom object using the URL pattern list.

   ```
   [edit security utm]
   user@host# set custom-objects custom-url-category custurl4 value urllist4
   ```


Copyright © 2016, Juniper Networks, Inc.

3. Create a list of untrusted sites

```
[edit security utm]
user@host# set custom-objects url-pattern urllistblack value
  [http://www.untrusted.com 13.13.13.13]
```

4. Configure the custom URL category list custom object using the URL pattern list of untrusted sites.

```
[edit security utm]
user@host# set custom-objects custom-url-category custblacklist value urllistblack
```

5. Create a list of trusted sites.

```
[edit security utm]
user@host# set custom-objects url-pattern urllistwhite value
  [http://www.trusted.com 7.7.7.7]
```

6. Configure the custom URL category list custom object using the URL pattern list of trusted sites.

```
[edit security utm]
user@host# set custom-objects custom-url-category custwhitelist value urllistwhite
```

Results From configuration mode, confirm your configuration by entering the **show security utm custom-objects** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost# show security utm custom-objects
url-pattern {
  urllist4 {
    value [ http://www.example.net 1.2.3.4 ];
  }
  urllistblack {
    value [ http://www.untrusted.com 13.13.13.13 ];
  }
  urllistwhite {
    value [ http://www.trusted.com 7.7.7.7 ];
  }
}
  custom-url-category {
    custurl4 {
      value urllist4;
    }
    custblacklist {
      value urllistblack;
    }
    custwhitelist {
      value urllistwhite;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring the Redirect Web Filtering Feature Profiles

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security utm feature-profile web-filtering url-whitelist custwhitelist
set security utm feature-profile web-filtering url-blacklist custblacklist
set security utm feature-profile web-filtering type websense-redirect
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    server host Websenseserver
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    server port 15868
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    fallback-settings server-connectivity block
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    fallback-settings timeout block
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    fallback-settings too-many-requests block
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    timeout 10
set security utm feature-profile web-filtering websense-redirect profile websenseprofile1
    sockets 1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure redirect Web filtering feature profiles:

1. Configure the Web filtering URL Black List.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-blacklist custblacklist
   ```

2. Configure the Web filtering URL White List.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set url-whitelist custwhitelist
   ```

3. Specify the Web filtering type, create a profile name, and set the server name or IP address.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set websense-redirect profile websenseprofile1 server host
     Websenseserver
   ```

4. Enter the port number for communicating with the server.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set websense-redirect profile websenseprofile1 server port 15868
   ```

5. Select fallback settings (block or log and permit) for this profile.

   ```
   [edit security utm feature-profile web-filtering]
   user@host# set websense-redirect profile websenseprofile1 fallback-settings default
     block
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set websense-redirect profile websenseprofile1 fallback-settings
  server-connectivity block
user@host# set websense-redirect profile websenseprofile1 fallback-settings
  timeout block
user@host# set websense-redirect profile websenseprofile1 fallback-settings
  too-many-requests block
```

6. Enter the number of sockets used for communicating between the client and the server.

```
[edit security utm feature-profile web-filtering]
user@host# set websense-redirect profile websenseprofile1 sockets 1
```

7. Enter a timeout value, in seconds.

```
[edit security utm feature-profile web-filtering]
user@host# set .websense-redirect profile websenseprofile1 timeout 10
```

Results From configuration mode, confirm your configuration by entering the **show security utm feature-profile** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost# show security utm feature-profile
  web-filtering {
    url-whitelist custwhitelist;
    url-blacklist custblacklist;
    type websense-redirect {
      profile websenseprofile1 {
        server {
          host Websenseserver;
          port 15868;
        }
        fallback-settings {
          server-connectivity block;
          timeout block;
          too-many-requests block;
        }
        timeout 10;
        sockets 1;
      }
    }
  }
    content-filtering {
      profile contentfilter1;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Redirect Web Filtering UTM Policies and Attaching the Redirect Web Filtering UTM Policies to Security Policies

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set security utm utm-policy utmp6 web-filtering http-profile websenseprofile1
set security policies from-zone trust to-zone untrust policy p6 match source-address any
set security policies from-zone trust to-zone untrust policy p6 match destination-address
  any
set security policies from-zone trust to-zone untrust policy p6 match application junos-http
set security policies from-zone trust to-zone untrust policy p6 then permit
  application-services utm-policy utmp6
```

Step-by-Step Procedure

To configure a UTM policy and attach it to a security policy:

1. Create the UTM policy referencing a profile.

   ```
   [edit security utm]
   user@host# set utm-policy utmp6 web-filtering http-profile websenseprofile1
   ```

2. Create and configure the security policy.

   ```
   [edit security policies from-zone trust to-zone untrust policy p6]
   user@host# set match source-address any
   user@host#  set match destination-address any
   user@host# set match application junos-http
   ```

3. Attach the UTM policy to the security policy.

   ```
   [edit security policies from-zone trust to-zone untrust policy p6]
   user@host# set then permit application-services utm-policy utmp6
   ```

Results

From configuration mode, confirm your configuration by entering the **show security utm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost# show security utm
utm-policy utmp6 {
  web-filtering {
    http-profile websenseprofile1;
  }
}
```

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
userhost# show security policies
  from-zone trust to-zone untrust {
    policy p6 {
      match {
        source-address any;
        destination-address any;
        application junos-http;
      }
      then {
        permit {
          application-services {
            utm-policy utmp6;
          }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Configuration of Redirect Web Filtering Custom Objects on page 5813
- Verifying the Configuration of Redirect Web Filtering Feature Profiles on page 5813
- Verifying the Attachment of Redirect Web Filtering UTM Policies to Security Policies on page 5814

Verifying the Configuration of Redirect Web Filtering Custom Objects

Purpose Verify the configuration of redirect Web filtering custom objects.

Action From the top of the configuration in configuration mode, enter the **show security utm custom-objects** command.

```
[edit]
userhost# show security utm custom-objects
url-pattern {
  urllist4 {
    value [ http://www.example.net 1.2.3.4 ];
  }
  urllistblack {
    value [ http://www.untrusted.com 13.13.13.13 ];
  }
  urllistwhite {
    value [ http://www.trusted.com 7.7.7.7 ];
  }
}
  custom-url-category {
    custurl4 {
      value urllist4;
    }
    custblacklist {
      value urllistblack;
    }
    custwhitelist {
      value urllistwhite;
    }
  }
```

Meaning The sample output shows the list of custom objects created.

Verifying the Configuration of Redirect Web Filtering Feature Profiles

Purpose Verify the configuration of redirect Web filtering feature profiles.

Copyright © 2016, Juniper Networks, Inc.

Action From the top of the configuration in configuration mode, enter the **show security utm feature-profile** command.

```
[edit]
userhost# show security utm feature-profile
  web-filtering {
    url-whitelist custwhitelist;
    url-blacklist custblacklist;
    type websense-redirect {
      profile websenseprofile1 {
        server {
          host Websenseserver;
          port 15868;
        }
        fallback-settings {
          server-connectivity block;
          timeout block;
          too-many-requests block;
        }
        timeout 10;
        sockets 1;
      }
    }
  }
    content-filtering {
      profile contentfilter1;
    }
```

Meaning The sample output shows the feature profile configured for a Websense redirect server.

Verifying the Attachment of Redirect Web Filtering UTM Policies to Security Policies

Purpose Verify the attachment of the newly created redirect Web filtering UTM policies to the security policies.

Action From the top of the configuration in configuration mode, enter the **show security utm** and **show security policies** commands.

```
[edit]
userhost# show security utm
utm-policy utmp6 {
  web-filtering {
    http-profile websenseprofile1;
  }
}
```

```
[edit]
userhost# show security policies
  from-zone trust to-zone untrust {
    policy p6 {
      match {
        source-address any;
        destination-address any;
```


Copyright © 2016, Juniper Networks, Inc.

```
            application junos-http;
          }
          then {
            permit {
              application-services {
                utm-policy utmp6;
              }
            }
          }
        }
      }
```

Meaning The sample output shows the security policies to which the newly created redirect Web filtering UTM policies are attached.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Understanding Local Web Filtering on page 5797
- **web-filtering on page 5998**
- Understanding Redirect Web Filtering on page 5805

Displaying Global SurfControl URL categories

Purpose View global URL categories defined and maintained by SurfControl.

Action Enter the **user@host# show groups junos-defaults** CLI command. You can also look for **custom-url-category**.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Redirect Web Filtering on page 5805
- Understanding Local Web Filtering on page 5797
- Understanding Integrated Web Filtering on page 5787

Monitoring Web Filtering Configurations

Purpose View Web-filtering statistics.

Action To view Web-filtering statistics using the CLI, enter the following commands:

```
user@host> show security utm web-filtering status
user@host> show security utm web-filtering statistics
```

To view Web-filtering statistics using J-Web:

1. Select **Clear Web Filtering Statistics**.

   The following information is displayed in the right pane.

Copyright © 2016, Juniper Networks, Inc.

Total Requests: #
White List Hit: #
Black List Hit: #
Queries to Server: #
Server Reply Permit: #
Server Reply Block: #
Custom Category Permit: #
Custom Category Block: #
Cache Hit Permit: #
Cache Hit Block: #
Web Filtering Session Total: #
Web Ffiltering Session Inuse: #
Fall Back: Log-and-Permit Block
Default # #
Timeout # #
Server-Connectivity # #
Too-Many-Requests # #

2. You can click the **Clear Web Filtering Statistics** button to clear all current viewable statistics and begin collecting new statistics.

Related Documentation

- Web Filtering Overview on page 5764
- Understanding Integrated Web Filtering on page 5787
- Example: Configuring Local Web Filtering on page 5799

Copyright © 2016, Juniper Networks, Inc.

PART 78

Configuration Statements and Operational Commands

- Configuration Statements on page 5819
- Operational Commands on page 6001

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 250

Configuration Statements

- Security Configuration Statement Hierarchy on page 5824
- SMTP Configuration Statement Hierarchy on page 5825
- [edit security policies] Hierarchy Level on page 5826
- [edit security utm] Hierarchy Level on page 5830
- action (Security UTM Web Filtering) on page 5838
- address-blacklist on page 5838
- address-whitelist on page 5839
- admin-email on page 5839
- administrator-email (Security Fallback Block) on page 5840
- administrator-email (Security Virus Detection) on page 5840
- allow-email (Security Fallback Block) on page 5841
- allow-email (Security Virus Detection) on page 5841
- application (Security Policies) on page 5842
- application-proxy (Security UTM) on page 5843
- anti-spam (Security Feature Profile) on page 5844
- anti-spam (Security UTM Policy) on page 5844
- anti-virus (Security Feature Profile) on page 5845
- anti-virus (Security UTM Policy) on page 5849
- block-command on page 5849
- block-content-type on page 5850
- block-extension on page 5850
- block-message (Security UTM) on page 5851
- block-mime on page 5851
- cache on page 5852
- category (Security Logging) on page 5853
- category (Security Web Filtering) on page 5854
- content-filtering (Security Feature Profile) on page 5860
- content-filtering (Security UTM Policy) on page 5861

Copyright © 2016, Juniper Networks, Inc.

- content-size on page 5862
- content-size (Security Antivirus Sophos Engine) on page 5863
- content-size-limit on page 5864
- corrupt-file on page 5864
- custom-block-message on page 5865
- custom-message (Security Content Filtering) on page 5865
- custom-message (Security Email Notify) on page 5866
- custom-message (Security Fallback Block) on page 5866
- custom-message (Security Fallback Non-Block) on page 5867
- custom-message (Security Virus Detection) on page 5867
- custom-message-subject (Security Email Notify) on page 5868
- custom-message-subject (Security Fallback Block) on page 5868
- custom-message-subject (Security Fallback Non-Block) on page 5869
- custom-message-subject (Security Virus Detection) on page 5869
- custom-objects on page 5870
- custom-tag-string on page 5871
- custom-url-category on page 5872
- decompress-layer on page 5873
- decompress-layer-limit on page 5873
- default (Security Antivirus) on page 5874
- default (Security Antivirus Sophos Engine) on page 5874
- default (Security UTM) on page 5875
- default (Security Web Filtering) on page 5876
- display-host (Security Fallback Block) on page 5877
- display-host (Security Virus Detection) on page 5877
- download-profile (Security Antivirus FTP) on page 5878
- download-profile (Security Content Filtering FTP) on page 5878
- email-notify on page 5879
- engine-not-ready on page 5879
- engine-not-ready (Security Antivirus Sophos Engine) on page 5880
- exception (Security Antivirus Mime Whitelist) on page 5880
- exception (Security Content Filtering) on page 5881
- fallback-block (Security Antivirus) on page 5881
- fallback-non-block (Security Antivirus) on page 5882
- fallback-options (Security Antivirus Juniper Express Engine) on page 5883
- fallback-options (Security Antivirus Kaspersky Lab Engine) on page 5884
- fallback-options (Security Antivirus Sophos Engine) on page 5885

Copyright © 2016, Juniper Networks, Inc.

- fallback-settings (Security Web Filtering) on page 5886
- fallback-settings (Security Web Filtering Juniper Local) on page 5886
- fallback-settings (Security Web Filtering Websense Redirect) on page 5887
- feature-profile on page 5888
- filename-extension on page 5893
- flag (SMTP) on page 5894
- format (Security Log Stream) on page 5895
- from-zone (Security Policies) on page 5896
- ftp (UTM Policy Anti-Virus) on page 5898
- ftp (UTM Policy Content Filtering) on page 5899
- host (Security Web Filtering) on page 5899
- http-profile (Security Antivirus) on page 5900
- http-profile (Security Content Filtering) on page 5900
- http-profile (Security Web Filtering) on page 5901
- imap-profile (Security UTM Policy Antivirus) on page 5901
- imap-profile (Security UTM Policy Content Filtering) on page 5902
- intelligent-prescreening on page 5902
- interval (Security Antivirus) on page 5903
- ipc on page 5904
- juniper-enhanced on page 5905
- juniper-express-engine on page 5906
- juniper-local on page 5908
- limit (UTM Policy) on page 5908
- list (Security Antivirus Mime Whitelist) on page 5909
- list (Security Content Filtering Block Mime) on page 5909
- log (Security) on page 5910
- mime-pattern on page 5912
- mime-whitelist on page 5913
- no-autoupdate on page 5914
- no-intelligent-prescreening on page 5915
- no-notify-mail-recipient on page 5916
- no-notify-mail-sender (Security Content Filtering Notification Options) on page 5916
- no-notify-mail-sender (Security Fallback Block) on page 5917
- no-notify-mail-sender (Security Virus Detection) on page 5917
- no-sbl-default-server on page 5918
- notification-options (Security Antivirus) on page 5919
- notification-options (Security Content Filtering) on page 5920

Copyright © 2016, Juniper Networks, Inc.

- notify-mail-recipient on page 5920
- notify-mail-sender (Security Content Filtering Notification Options) on page 5921
- notify-mail-sender (Security Fallback Block) on page 5921
- notify-mail-sender (Security Virus Detection) on page 5922
- no-uri-check on page 5922
- out-of-resources on page 5923
- out-of-resources (Security Antivirus Sophos Engine) on page 5924
- over-limit on page 5924
- packet-filter on page 5925
- password (Security Antivirus) on page 5926
- password-file on page 5926
- pattern-update (Security Antivirus) on page 5927
- permit-command on page 5928
- policies on page 5929
- pop3-profile (Security UTM Policy Antivirus) on page 5933
- pop3-profile (Security UTM Policy Content Filtering) on page 5934
- port (Security Antivirus) on page 5934
- port (Security Web Filtering Server) on page 5935
- primary-server on page 5935
- profile (Security Antispam SBL) on page 5936
- profile (Security Antivirus Juniper Express Engine) on page 5937
- profile (Security Antivirus Kaspersky Lab Engine) on page 5939
- profile (Security Content Filtering) on page 5940
- profile (Security Sophos Engine Antivirus) on page 5941
- profile (Security Web Filtering Juniper Enhanced) on page 5942
- profile (Security Web Filtering Juniper Local) on page 5943
- profile (Security Web Filtering Surf Control Integrated) on page 5944
- profile (Security Web Filtering Websense Redirect) on page 5945
- protocol-command on page 5946
- proxy (Security Antivirus) on page 5946
- quarantine-message (Security UTM) on page 5947
- sbl on page 5948
- sbl-default-server on page 5948
- scan-extension on page 5949
- scan-mode on page 5949
- scan-options (Security Antivirus Juniper Express Engine) on page 5950
- scan-options (Security Antivirus Kaspersky Lab Engine) on page 5950

Copyright © 2016, Juniper Networks, Inc.

- scan-options (Security Antivirus Sophos Engine) on page 5951
- secondary-server on page 5951
- server (Security Antivirus) on page 5952
- server (Security Web Filtering) on page 5952
- server-connectivity on page 5953
- sessions-per-client on page 5954
- site-reputation-action on page 5955
- size (Security Web Filtering Cache) on page 5956
- smtp-profile (Security UTM Policy Antispam) on page 5956
- smtp-profile (Security UTM Policy Antivirus) on page 5957
- smtp-profile (Security UTM Policy Content Filtering) on page 5957
- sockets on page 5958
- sophos-engine on page 5959
- spam-action on page 5960
- surf-control-integrated on page 5961
- sxl-retry on page 5962
- sxl-timeout on page 5962
- timeout (Security Antivirus Fallback Options) on page 5963
- timeout (Security Antivirus Fallback Options Sophos Engine) on page 5964
- timeout (Security Antivirus Scan Options) on page 5964
- timeout (Security Web Filtering) on page 5965
- timeout (Security Web Filtering Cache) on page 5965
- timeout (Security Web Filtering Fallback Settings) on page 5966
- too-many-requests (Security Antivirus Fallback Options) on page 5967
- too-many-requests (Security Antivirus Fallback Options Sophos Engine) on page 5967
- too-many-requests (Security Web Filtering Fallback Settings) on page 5968
- to-zone (Security Policies) on page 5969
- traceoptions (Security Antispam) on page 5971
- traceoptions (Security Antivirus) on page 5972
- traceoptions (Security Application Proxy) on page 5973
- traceoptions (Security Content Filtering) on page 5974
- traceoptions (Security UTM) on page 5975
- traceoptions (Security Web Filtering) on page 5976
- traceoptions (SMTP) on page 5977
- traffic-options on page 5977
- trickling on page 5978
- type (Security Antivirus Feature Profile) on page 5979

Copyright © 2016, Juniper Networks, Inc.

- type (Security Content Filtering Notification Options) on page 5979
- type (Security Fallback Block) on page 5980
- type (Security Virus Detection) on page 5981
- type (Security Web Filtering) on page 5982
- upload-profile (Security Antivirus FTP) on page 5982
- upload-profile (Security Content Filtering FTP) on page 5983
- uri-check on page 5983
- url (Security Antivirus) on page 5984
- url-blacklist on page 5984
- url-pattern on page 5985
- url-whitelist (Security Antivirus) on page 5985
- url-whitelist (Security Web Filtering) on page 5986
- username (Security Antivirus) on page 5986
- utm on page 5987
- utm-policy on page 5995
- utm-policy (Application Services) on page 5996
- virus-detection (Security Antivirus) on page 5997
- web-filtering on page 5998
- websense-redirect on page 5999

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*

Copyright © 2016, Juniper Networks, Inc.

- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*
- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

SMTP Configuration Statement Hierarchy

```
smtp {
  primary-server {
    address ipv4-address;
    login sender-email-address {
      password password;
    }
  }
  secondary-server {
    address ipv4-address;
```

Copyright © 2016, Juniper Networks, Inc.

```
    login sender-email-address {
      password password;
    }
  }
  traceoptions {
    flag {
      all;
      configuration;
      IPC;
      protocol-exchange;
      send-request;
    }
  }
}
```

Related Documentation

- *CLI User Guide*

[edit security policies] Hierarchy Level

```
security {
  policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          destination-address-excluded;
          source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
            }
          source-address-excluded;
          source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
          }
        }
        scheduler-name scheduler-name;
        then {
```


Copyright © 2016, Juniper Networks, Inc.

```
count {
  alarm {
    per-minute-threshold number;
    per-second-threshold number;
  }
}
deny;
log {
  session-close;
  session-init;
}
permit {
  application-services {
    application-firewall {
      rule-set rule-set-name;
    }
    application-traffic-control {
      rule-set rule-set-name;
    }
    gprs-gtp-profile profile-name;
    gprs-sctp-profile profile-name;
    idp;
    redirect-wx | reverse-redirect-wx;
    ssl-proxy {
      profile-name profile-name;
    }
    uac-policy {
      captive-portal captive-portal;
    }
    utm-policy policy-name;
  }
  destination-address {
    drop-translated;
    drop-untranslated;
  }
  firewall-authentication {
    pass-through {
      access-profile profile-name;
      client-match user-or-group-name;
      ssl-termination-profile profile-name;
      web-redirect;
      web-redirect-to-https;
    }
    user-firewall {
      access-profile profile-name;
      domain domain-name
      ssl-termination-profile profile-name;
    }
    web-authentication {
      client-match user-or-group-name;
    }
  }
  services-offload;
  tcp-options {
    sequence-check-required;
    syn-check-required;
```

Copyright © 2016, Juniper Networks, Inc.

```
            }
            tunnel {
              ipsec-group-vpn group-vpn;
              ipsec-vpn vpn-name;
              pair-policy pair-policy;
            }
          }
          reject;
        }
      }
    }
    global {
      policy policy-name {
        description description;
        match {
          application {
            [application];
            any;
          }
          destination-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
          }
          from-zone {
            [zone-name];
            any;
          }
          source-address {
            [address];
            any;
            any-ipv4;
            any-ipv6;
          }
          source-identity {
            [role-name];
            any;
            authenticated-user;
            unauthenticated-user;
            unknown-user;
          }
          to-zone {
            [zone-name];
            any;
          }
        }
        scheduler-name scheduler-name;
        then {
          count {
            alarm {
              per-minute-threshold number;
              per-second-threshold number;
            }
          }
          deny;
```


Copyright © 2016, Juniper Networks, Inc.

```
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
          captive-portal captive-portal;
        }
        utm-policy policy-name;
      }
      destination-address {
        drop-translated;
        drop-untranslated;
      }
      firewall-authentication {
        pass-through {
          access-profile profile-name;
          client-match user-or-group-name;
          ssl-termination-profile profile-name;
          web-redirect;
          web-redirect-to-https;
        }
        user-firewall {
          access-profile profile-name;
          domain domain-name
          ssl-termination-profile profile-name;
        }
        web-authentication {
          client-match user-or-group-name;
        }
      }
      services-offload;
      tcp-options {
        initial-tcp-mss mss-value;
        reverse-tcp-mss mss-value;
        sequence-check-required;
        syn-check-required;
      }
    }
    reject;
  }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable);
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        (no-world-readable | world-readable);
        size maximum-file-size;
      }
      flag flag;
      no-remote-trace;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045
- Unified Threat Management Overview on page 5583

[edit security utm] Hierarchy Level

```
security {
  utm {
    application-proxy {
      traceoptions {
        flag flag;
      }
    }
    custom-objects {
      custom-url-category object-name {
        value [value];
      }
      filename-extension object-name {
        value [value];
      }
      mime-pattern object-name {
        value [value];
      }
      protocol-command object-name {
        value [value];
      }
      url-pattern object-name {
        value [value];
      }
    }
    feature-profile {
      anti-spam {
        address-blacklist list-name;
        address-whitelist list-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
    sbl {
      profile profile-name {
        custom-tag-string [string];
        (sbl-default-server | no-sbl-default-server);
        spam-action (block | tag-header | tag-subject);
      }
    }
    traceoptions {
      flag flag;
    }
  }
  anti-virus {
    juniper-express-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile profile-name {
        fallback-options {
          content-size (block | log-and-permit);
          default (block | log-and-permit);
          engine-not-ready (block | log-and-permit);
          out-of-resources (block | (log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        notification-options {
          fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
          fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
          }
          virus-detection {
            custom-message message;
            custom-message-subject message-subject;
```

```
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
        }
        scan-options {
          content-size-limit value;
          (intelligent-prescreening | no-intelligent-prescreening);
          timeout value;
        }
        trickling {
          timeout value;
        }
      }
    }
    kaspersky-lab-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile profile-name {
        fallback-options {
          content-size (block | log-and-permit);
          corrupt-file (block | log-and-permit);
          decompress-layer (block | log-and-permit);
          default (block | log-and-permit);
          engine-not-ready (block | log-and-permit);
          out-of-resources (block | (log-and-permit);
          password-file (block | (log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        notification-options {
          fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
          fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
```


Copyright © 2016, Juniper Networks, Inc.

```
            (notify-mail-recipient | no-notify-mail-recipient);
          }
          virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
        }
        scan-options {
          content-size-limit value;
          decompress-layer-limit value;
          (intelligent-prescreening | no-intelligent-prescreening);
          scan-extension filename;
          scan-mode (all | by-extension);
          timeout value;
        }
        trickling {
          timeout value;
        }
      }
    }
    mime-whitelist {
      exception listname;
      list listname {
        exception listname;
      }
    }
    sophos-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile <name> {
        fallback-options {
          content-size (block | log-and-permit | permit);
          default (block | log-and-permit | permit);
          engine-not-ready (block | log-and-permit | permit);
          out-of-resources (block | log-and-permit | permit);
          timeout (block | log-and-permit | permit);
          too-many-requests (block | log-and-permit | permit);
        }
        notification-options {
          fallback-block {
```

Copyright © 2016, Juniper Networks, Inc.

```
                administrator-email email-address;
                allow-email;
                custom-message message;
                custom-message-subject message-subject;
                display-host;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
              }
              fallback-non-block {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
              }
              virus-detection {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
              }
            }
            scan-options {
              content-size-limit value;
              (no-uri-check | uri-check);
              timeout value;
            }
            trickling {
              timeout value;
            }
          }
          sxl-retry value;
          sxl-timeout seconds;
        }
        traceoptions {
          flag flag;
        }
        type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
        url-whitelist listname;
      }
      content-filtering {
        profile profile-name {
          block-command protocol-command-list;
          block-content-type (activex | exe | http-cookie | java-applet | zip);
          block-extension extension-list;
          block-mime {
            exception list-name;
            list list-name;
          }
          notification-options {
            custom-message message;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
          permit-command protocol-command-list;
        }
        traceoptions {
          flag flag;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  web-filtering {
    juniper-enhanced {
      cache {
        size value;
        timeout value;
      }
      profile profile-name {
        block-message {
          type {
            custom-redirect-url;
          }
          url url;
        }
        quarantine-message {
          type {
            custom-redirect-url;
          }
          url url;
        }
        category customurl-list name {
          action (block | log-and-permit | permit | quarantine);
        }
        custom-block-message value;
        custom-quarantine-message value;
        default (block | log-and-permit | permit | quarantine);
        fallback-settings {
          default (block | log-and-permit);
          server-connectivity (block | log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        no-safe-search;
        site-reputation-action {
          fairly-safe (block | log-and-permit | permit | quarantine);
          harmful (block | log-and-permit | permit | quarantine);
          moderately-safe (block | log-and-permit | permit | quarantine);
          suspicious (block | log-and-permit | permit | quarantine);
          very-safe (block | log-and-permit | permit |quarantine);
        }
        timeout value;
      }
      server {
        host host-name;
        port number;
      }
    }
    juniper-local {
      profile profile-name {
        custom-block-message value;
        default (block | log-and-permit | permit);
        fallback-settings {
          default (block | log-and-permit);
          server-connectivity (block | log-and-permit);
          timeout (block | log-and-permit);
```

```
            too-many-requests (block | log-and-permit);
          }
          timeout value;
        }
      }
      surf-control-integrated {
        cache {
          size value;
          timeout value;
        }
        profile profile-name {
          category customurl-list name {
            action (block | log-and-permit | permit);
          }
          custom-block-message value;
          default (block | log-and-permit | permit);
          fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
          }
          timeout value;
        }
        server {
          host host-name;
          port number;
        }
      }
      traceoptions {
        flag flag;
      }
      type (juniper-enhanced | juniper-local | surf-control-integrated |
        websense-redirect);
      url-blacklist listname;
      url-whitelist listname;
      websense-redirect {
        profile profile-name {
          account value;
          custom-block-message value;
          fallback-settings {
            default (block | log-and-permit);
            server-connectivity (block | log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
          }
          server {
            host host-name;
            port number;
          }
          sockets value;
          timeout value;
        }
      }
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    ipc {
      traceoptions flag flag;
    }
    traceoptions {
      flag flag;
    }
    utm-policy policy-name {
      anti-spam {
        smtp-profile profile-name;
      }
      anti-virus {
        ftp {
          download-profile profile-name;
          upload-profile profile-name;
        }
        http-profile profile-name;
        imap-profile profile-name;
        pop3-profile profile-name;
        smtp-profile profile-name;
      }
      content-filtering {
        ftp {
          download-profile profile-name;
          upload-profile profile-name;
        }
        http-profile profile-name;
        imap-profile profile-name;
        pop3-profile profile-name;
        smtp-profile profile-name;
      }
      traffic-options {
        sessions-per-client {
          limit value;
          over-limit (block | log-and-permit);
        }
      }
      web-filtering {
        http-profile profile-name;
      }
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Unified Threat Management Overview on page 5583

Copyright © 2016, Juniper Networks, Inc.

action (Security UTM Web Filtering)

Syntax action (block | log-and-permit | permit | quarantine);

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name* category *customurl-last-name*]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* category *customurl-last-name*]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for UTM Enhanced Web Filtering.

Description Enter an action to go with the customurl-list filter.

Options
- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit—Permit the traffic.
- quarantine—Show the warning message and permit/block the traffic based on user input.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

address-blacklist

Syntax address-blacklist *list-name*;

Hierarchy Level [edit security utm feature-profile anti-spam]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enter an address blacklist (or whitelist) custom object for local list spam filtering.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

address-whitelist

Syntax address-whitelist *list-name*;

Hierarchy Level [edit security utm feature-profile anti-spam]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Enter an address-whitelist (or blacklist) custom-object for local list spam filtering.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

admin-email

Syntax admin-email *email-address*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]

Release Information Statement introduced in Junos OS Release 9.5.

Description You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

administrator-email (Security Fallback Block)

Syntax administrator-email *email-address*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Configure the administrator e-mail address that will be notified when a fallback-block occurs. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

administrator-email (Security Virus Detection)

Syntax administrator-email *email address*;

Hierarchy Level
[edit security utm feature-profile anti-virus sophos-engine profile *profile name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure the administrator e-mail address that will be notified when a virus is detected by Sophos antivirus. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

allow-email (Security Fallback Block)

Syntax allow-email;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Enable e-mail notification to notify a specified administrator when a fallback-block occurs.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

allow-email (Security Virus Detection)

Syntax allow–email;

Hierarchy Level [edit security utm feature-profile anti-virus profile notification-options virus-detect]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Enable e-mail notification to notify a specified administrator when a virus is detected.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

application (Security Policies)

Syntax

```
application {
  [application];
  any;
}
```

Hierarchy Level

[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* match]

[edit security policies global policy *policy-name* match]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Specify the IP or remote procedure call (RPC) application or set of applications to be used as match criteria.

Options

application-name-or-set—Name of the predefined or custom application or application set used as match criteria.

any—Any predefined or custom applications or application sets.

NOTE: A custom application that does not use a well-known destination port for the application will not be included in the any option, and must be named explicitly.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

application-proxy (Security UTM)

Syntax

```
application-proxy {
  traceoptions {
    flag flag;
  }
}
```

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace options for the application proxy.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

anti-spam (Security Feature Profile)

Syntax

```
anti-spam {
  address-blacklist list-name;
  address-whitelist list-name;
  sbl {
    profile profile-name {
      custom-tag-string [string];
      (sbl-default-server | no-sbl-default-server);
      spam-action (block | tag-header | tag-subject);
    }
  }
  traceoptions flag flag;
}
```

Hierarchy Level [edit security utm feature-profile]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure UTM antispam features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

anti-spam (Security UTM Policy)

Syntax

```
anti-spam {
  smtp-profile profile-name;
}
```

Hierarchy Level [edit security utm utm-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antispam SMTP protocol and attach this policy to a security profile to implement it.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

anti-virus (Security Feature Profile)

Syntax

```
anti-virus {
  juniper-express-engine {
    pattern-update {
      email-notify {
        admin-email email-address;
        custom-message message;
        custom-message-subject message-subject;
      }
      interval value;
      no-autoupdate;
      proxy {
        password password-string;
        port port-number;
        server address-or-url;
        username name;
      }
      url url;
    }
    profile profile-name {
      fallback-options {
        content-size (block | log-and-permit);
        default (block | log-and-permit);
        engine-not-ready (block | log-and-permit);
        out-of-resources (block | (log-and-permit);
        timeout (block | log-and-permit);
        too-many-requests (block | log-and-permit);
      }
      notification-options {
        fallback-block {
          administrator-email email-address;
          allow-email;
          custom-message message;
          custom-message-subject message-subject;
          display-host;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
        fallback-non-block {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-recipient | no-notify-mail-recipient);
        }
        virus-detection {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
      }
      scan-options {
        content-size-limit value;
        (intelligent-prescreening | no-intelligent-prescreening);
```

Copyright © 2016, Juniper Networks, Inc.

```
            timeout value;
        }
        trickling {
            timeout value;
        }
    }
}
kaspersky-lab-engine {
    pattern-update {
        email-notify {
            admin-email email-address;
            custom-message message;
            custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
            password password-string;
            port port-number;
            server address-or-url;
            username name;
        }
        url url;
    }
    profile profile-name {
        fallback-options {
            content-size (block | log-and-permit);
            corrupt-file (block | log-and-permit);
            decompress-layer (block | log-and-permit);
            default (block | log-and-permit);
            engine-not-ready (block | log-and-permit);
            out-of-resources (block | (log-and-permit);
            password-file (block | (log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        notification-options {
            fallback-block {
                administrator-email email-address;
                allow-email;
                custom-message message;
                custom-message-subject message-subject;
                display-host;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
            }
            fallback-non-block {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      scan-options {
        content-size-limit value;
        decompress-layer-limit value;
        (intelligent-prescreening | no-intelligent-prescreening);
        scan-extension filename;
        scan-mode (all | by-extension);
        timeout value;
      }
      trickling {
        timeout value;
      }
    }
  }
  mime-whitelist {
    exception listname;
    list listname {
      exception listname;
    }
  }
  sophos-engine {
    pattern-update {
      email-notify {
        admin-email email-address;
        custom-message message;
        custom-message-subject message-subject;
      }
      interval value;
      no-autoupdate;
      proxy {
        password password-string;
        port port-number;
        server address-or-url;
        username name;
      }
      url url;
    }
    profile <name> {
      fallback-options {
        content-size (block | log-and-permit | permit);
        default (block | log-and-permit | permit);
        engine-not-ready (block | log-and-permit | permit);
        out-of-resources (block | log-and-permit | permit);
        timeout (block | log-and-permit | permit);
        too-many-requests (block | log-and-permit | permit);
      }
      notification-options {
        fallback-block {
          administrator-email email-address;
          allow-email;
          custom-message message;
          custom-message-subject message-subject;
          display-host;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        fallback-non-block {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-recipient | no-notify-mail-recipient);
        }
        virus-detection {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
      }
      scan-options {
        content-size-limit value;
        (no-uri-check | uri-check);
        timeout value;
      }
      trickling {
        timeout value;
      }
    }
    sxl-retry value;
    sxl-timeout seconds;
  }
  traceoptions flag flag;
  type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
  url-whitelist listname;
}
```

Hierarchy Level [edit security utm feature-profile]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure UTM antivirus full and express features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

anti-virus (Security UTM Policy)

Syntax

```
anti-virus {
  ftp {
    download-profile profile-name;
    upload-profile profile-name;
  }
  http-profile profile-name;
  imap-profile profile-name;
  pop3-profile profile-name;
  smtp-profile profile-name;
}
```

Hierarchy Level [edit security utm utm-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus protocols and attach this policy to a security profile to implement it.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

block-command

Syntax block-command *protocol-command-list*;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Apply protocol block command custom-objects to the content-filtering profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

block-content-type

Syntax block-content-type (activex | exe | http-cookie | java-applet | zip);

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Apply blocks to other available content such as exe, http-cookie, java-applet. This is for HTTP only.

Options

- **activex**—Block ActiveX.
- **exe**—Block EXE files.
- **http-cookie**—Block cookies.
- **java-applet**—Block Java applets.
- **zip**—Block ZIP files.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

block-extension

Syntax block-extension *extension-list*;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Apply block extensions to the content-filtering profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

block-message (Security UTM)

Syntax

```
block-message {
  type {
    custom-redirect-url;
  }
  url url;
}
```

Hierarchy Level [edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Configure Juniper enhanced block message settings.

Options

- **type**—Specify the following type of the block message:
  - **custom-redirect-url**—Specify Custom redirect URL server.
- **url *url***—Specify an URL of the block message.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

block-mime

Syntax

```
block-mime {
  exception list-name;
  list list-name;
}
```

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Apply MIME pattern list custom-objects to the content-filtering profile for blocking MIME types.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

cache

Syntax

```
cache {
  size value;
  timeout value;
}
```

Hierarchy Level

[edit security utm feature-profile web-filtering surf-control-integrated]
[edit security utm feature-profile web-filtering juniper-enhanced]

Release Information

Statement introduced in Junos OS Release 9.5 for surf-control integrated.
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description

Set the cache parameters for Surf-Control-Integrated Web filtering and Enhanced Web Filtering.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

category (Security Logging)

Syntax category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec | idp | rtlog | pst-ds-lite | appqos | secintel)

Hierarchy Level [edit security log stream *stream-name*]

Release Information Statement introduced in Junos OS Release 10.0. Statement modified in Junos OS Release 15.1X49-D40.

Description Set the category of logging to **all** or **content-security**. Note that for the WELF format, the category must be set to **content-security**.

NOTE: On SRX5600 and SRX5800 devices, if the stream configuration does not specify a destination port, the default destination port will be the system log port. If you specify a destination port in the stream configuration, then that port will be used instead.

Options

- **all**—All events are logged.
- **content-security**—Only content security events are logged.
- **fw-auth**—Firewall authentication events are logged.
- **screen**—Screen events are logged.
- **alg**—Application Layer Gateway (ALG) events are logged.
- **nat**—Network Address Translation (NAT) events are logged.
- **flow**—Flow events are logged.
- **sctp**—Stream Control Transmission Protocol (SCTP) events are logged.
- **gtp**—GTP events are logged.
- **ipsec**—IPsec events are logged.
- **idp**—Intrusion Detection and Prevention (IDP) events are logged.
- **rtlog**—RTLOG system log events are logged.
- **pst-ds-lite**—PST dual-stack lite (DS-Lite) events are logged.
- **appqos**—Application quality of service (AppQoS) events are logged.
- **secintel**—Juniper networks Security Intelligence (SecIntel) events are logged.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *AppSecure Services Feature Guide for Security Devices*
- *Logical Systems Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

category (Security Web Filtering)

Syntax

```
category customurl-list name {
  action (block | log-and-permit | permit | quarantine);
}
```

Hierarchy Level

[edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name*]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering. Support for new categories and category name updates by Websense added in Junos OS Release 12.1X47-D15 and 12.3X48-D10.

Description

Select a custom URL category list you created (custom objects) for filtering against.

Copyright © 2016, Juniper Networks, Inc.

Table 496: List of categories predefined by Websense

Category ID	Category Name	Parent ID
1	Adult Material	0
2	Business and Economy	0
3	Education	0
4	Government	0
5	News and Media	0
6	Religion	0
7	Society and Lifestyles	0
8	Special Events	0
9	Information Technology	0
10	Abortion	0
11	Advocacy Groups	0
12	Entertainment	0
13	Gambling	0
14	Games	0
15	Illegal or Questionable	0
16	Job Search	0
17	Shopping	0
18	Sports	0
19	Tasteless	0
20	Travel	0
21	Vehicles	0
22	Violence	0
23	Weapons	0
24	Drugs	0

Copyright © 2016, Juniper Networks, Inc.

Table 496: List of categories predefined by Websense *(continued)*

Category ID	Category Name	Parent ID
25	Militancy and Extremist	0
26	Intolerance	0
27	Health	0
28	Website Translation	9
29	Advertisements	110
64	User-Defined	0
65	Nudity	1
66	Adult Content	1
67	Sex	1
68	Financial Data and Services	2
69	Cultural Institutions	3
70	Media File Download	12
72	Military	4
73	Political Organizations	4
74	General Email	91
75	Proxy Avoidance	9
76	Search Engines and Portals	9
78	Web Hosting	9
79	Web Chat	91
80	Hacking	9
81	Alternative Journals	5
82	Non-Traditional Religions	6
83	Traditional Religions	6
84	Restaurants and Dining	7

Copyright © 2016, Juniper Networks, Inc.

Table 496: List of categories predefined by Websense *(continued)*

Category ID	Category Name	Parent ID
85	Gay or Lesbian or Bisexual Interest	7
86	Personals and Dating	7
87	Alcohol and Tobacco	7
88	Prescribed Medications	24
89	Nutrition	24
90	Abused Drugs	24
91	Internet Communication	0
92	Pro-Choice	10
93	Pro-Life	10
94	Sex Education	1
95	Lingerie and Swimsuit	1
96	Online Brokerage and Trading	110
97	Educational Institutions	3
98	Instant Messaging	110
99	Application and Software Download	110
100	Pay-to-Surf	110
101	Internet Auctions	17
102	Real Estate	17
103	Hobbies	7
107	Sport Hunting and Gun Clubs	18
108	Internet Telephony	116
109	Streaming Media	116
110	Productivity	0
111	Marijuana	24

Copyright © 2016, Juniper Networks, Inc.

Table 496: List of categories predefined by Websense *(continued)*

Category ID	Category Name	Parent ID
112	Message Boards and Forums	110
113	Personal Network Storage and Backup	116
114	Internet Radio and TV	116
115	Peer-to-Peer File Sharing	116
116	Bandwidth	0
117	Social Networking and Personal Sites	7
118	Educational Materials	3
121	Reference Materials	3
122	Social Organizations	0
123	Service and Philanthropic Organizations	122
124	Social and Affiliation Organizations	122
125	Professional and Worker Organizations	122
126	Security	0
128	Malicious Web Sites	126
138	Computer Security	9
146	Miscellaneous	0
147	Web Infrastructure	146
148	Web Images	146
149	Private IP Addresses	146
150	Content Delivery Networks	146
151	Dynamic Content	146
152	Network Errors	146
153	Uncategorized	146
154	Spyware	126

Copyright © 2016, Juniper Networks, Inc.

Table 496: List of categories predefined by Websense *(continued)*

Category ID	Category Name	Parent ID
156	File Download Servers	146
164	Phishing and Other Frauds	126
166	Keyloggers	126
167	Potentially Unwanted Software	126
172	Bot Networks	126
191	Extended Protection	0
192	Elevated Exposure	191
193	Emerging Exploits	191
194	Suspicious Content	191
195	Organizational Email	91
196	Text and media messaging	91
200	Web and Email Spam	9
220	Compromised Websites	0
221	Newly Registered Websites	0
222	Collaboration Office	0
223	Office Mail	222
224	Office Drive	222
225	Office Documents	222
226	Office Apps	222
227	Web Analytics	9
228	Web and Email Marketing	9
1529	Classifieds Posting	0
1530	Blog Posting	0
1531	Blog Commenting	0

Copyright © 2016, Juniper Networks, Inc.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

content-filtering (Security Feature Profile)

Syntax

```
content-filtering {
  profile profile-name {
    block-command protocol-command-list;
    block-content-type (activex | exe | http-cookie | java-applet | zip);
    block-extension extension-list;
    block-mime {
      exception list-name;
      list list-name;
    }
    notification-options {
      custom-message message;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
    permit-command protocol-command-list;
  }
  traceoptions flag flag;
}
```

Hierarchy Level [edit security utm feature-profile]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure UTM content-filtering features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

content-filtering (Security UTM Policy)

Syntax

```
content-filtering {
  ftp {
    download-profile profile-name;
    upload-profile profile-name;
  }
  http-profile profile-name;
  imap-profile profile-name;
  pop3-profile profile-name;
  smtp-profile profile-name;
}
```

Hierarchy Level [edit security utm utm-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering protocols and attach this policy to a security profile to implement it.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

content-size

Syntax content-size (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description If the content size exceeds a set limit, the content is either passed or blocked. The default action is log-and-permit.

NOTE: When you configure the content-size value, keep in mind that in certain cases, content size is available in the protocol headers, so the max-content-size fallback is applied before a scan request is sent. However, in many cases, content size is not provided in the protocol headers. In these cases, the TCP payload is sent to the antivirus scanner and accumulates until the end of the payload. If the accumulated payload exceeds the maximum content size value, then max-content-size fallback is applied. The default fallback action is log and permit, so you may want to change this option to block, in which case such a packet is dropped and a block message is sent to the client.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

content-size (Security Antivirus Sophos Engine)

Syntax content-size (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description If the content size exceeds a set limit, the content is either passed or blocked.

NOTE: When you configure the content-size value, keep in mind that in certain cases, content size is available in the protocol headers, so the max-content-size fallback is applied before a scan request is sent. However, in many cases, content size is not provided in the protocol headers. In these cases, the TCP payload is sent to the antivirus scanner and accumulates until the end of the payload. If the accumulated payload exceeds the maximum content size value, then max-content-size fallback is applied. You might want to set the fallback action to block, in which case such a packet is dropped and a block message is sent to the client.

Options

- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

content-size-limit

Syntax content-size-limit *value*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description The content size check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.

Range: 20 through 20,000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

corrupt-file

Syntax corrupt-file (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Corrupt file is the error returned by the scan engine when engine detects a corrupted file. The default action is log-and-permit.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Full Antivirus Configuration Overview on page 5648

Copyright © 2016, Juniper Networks, Inc.

custom-block-message

Syntax custom-block-message *value*;

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name*]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Enter a custom message to be sent when HTTP requests are blocked.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

custom-message (Security Content Filtering)

Syntax custom-message *message*;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* notification-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Custom message notifications are generally used when content is blocked by the content filter.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

custom-message (Security Email Notify)

Syntax custom-message *message*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]

Release Information Statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Description You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

custom-message (Security Fallback Block)

Syntax custom-message *message*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

custom-message (Security Fallback Non-Block)

Syntax custom-message *message*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-non-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

custom-message (Security Virus Detection)

Syntax custom-message *message*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

custom-message-subject (Security Email Notify)

Syntax custom-message-subject *message-subject*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update email-notify]
[edit security utm feature-profile anti-virus sophos-engine pattern-update email-notify]

Release Information Statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Description You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

custom-message-subject (Security Fallback Block)

Syntax custom-message-subject *message-subject*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

custom-message-subject (Security Fallback Non-Block)

Syntax custom-message-subject *message-subject*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-non-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

custom-message-subject (Security Virus Detection)

Syntax custom-message-subject *message-subject*;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Custom message notifications are mainly used in file replacement or in a response message when the antivirus scan result is to drop the file. As part of a custom message, you can customize the message subject line.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

custom-objects

Syntax

```
custom-objects {
  custom-url-category object-name {
    value [value];
  }
  filename-extension object-name {
    value [value];
  }
  mime-pattern object-name {
    value [value];
  }
  protocol-command object-name {
    value [value];
  }
  url-pattern object-name {
    value [value];
  }
}
```

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure custom objects before configuring UTM feature-profile features.

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

custom-tag-string

Syntax custom-tag-string [*string*];

Hierarchy Level [edit security utm feature-profile anti-spam sbl profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a custom string for identifying a message as spam.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Antispam Filtering Overview on page 5599

Copyright © 2016, Juniper Networks, Inc.

custom-url-category

Syntax

```
custom-url-category object-name {
  value [value];
}
```

Hierarchy Level [edit security utm custom-objects]

Release Information Statement introduced in Junos OS Release 9.5.

Description Use URL pattern lists to create Custom URL category lists. These are lists of patterns that bypass scanning.

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Options

- ***object-name***—Name of the URL category-list object.
- **value *value***—Value of the URL category-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding UTM Custom Objects on page 5585
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

decompress-layer

Syntax decompress-layer (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Description Decompress layer error is the error returned by the scan engine when the scanned file has too many compression layers. The default action is block.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Full Antivirus Configuration Overview on page 5648

decompress-layer-limit

Syntax decompress-layer-limit *value*;

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description The decompression layer limit specifies how many layers of nested compressed files and files with internal extractable objects, such as archive files (tar), the internal antivirus scanner can decompress before it executes the virus scan.

Range: 1 through 4

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Full Antivirus Configuration Overview on page 5648

Copyright © 2016, Juniper Networks, Inc.

default (Security Antivirus)

Syntax default (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description All errors other than those specifically listed fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

default (Security Antivirus Sophos Engine)

Syntax default (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description All errors other than those specifically listed fall into this category. This could include either unhandled system exceptions (internal errors) or other unknown errors.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

default (Security UTM)

Syntax default (block |log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4 .

Description Specify the default action to take for a URL.

Options
- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit—Permit the traffic.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

default (Security Web Filtering)

Syntax default (block | log-and-permit | permit | quarantine);

Hierarchy Level
[edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering juniper–local profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* fallback-settings]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Specify an action for the profile, for requests that experience internal errors in the Web-filtering module.

Options
- block—Log the error and deny the traffic.
- log-and-permit—Log the error and permit the traffic.
- permit —Permit the traffic.
- quarantine—Show the warning message and permit/block the traffic based on user input.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

display-host (Security Fallback Block)

Syntax display-host;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Display the computer host name in the notification e-mail sent to the administrator when a fallback-block notification occurs.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

display-host (Security Virus Detection)

Syntax display-host;

Hierarchy Level [edit security utm feature-profile anti-virus profile *profile name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Display the computer host name in the notification e-mail sent to the administrator when a virus is detected by Sophos antivirus.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

download-profile (Security Antivirus FTP)

Syntax download-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus ftp]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus FTP (download) protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

download-profile (Security Content Filtering FTP)

Syntax download-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering ftp]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering FTP (download) protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

email-notify

Syntax

```
email-notify {
  admin-email email-address;
  custom-message message;
  custom-message-subject message-subject;
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]

Release Information

Statement introduced in Junos OS Release 9.5 .

Description

You can configure the device to notify a specified administrator when patterns are updated. This is an e-mail notification with a custom message and a custom subject line.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

engine-not-ready

Syntax

```
engine-not-ready (block | log-and-permit);
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information

Statement introduced in Junos OS Release 9.5 .

Description

The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting. The default action is block.

Options

- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

engine-not-ready (Security Antivirus Sophos Engine)

Syntax default (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Release 11.1 .

Description The scan engine is initializing itself, for example, loading the signature database. During this phase, it is not ready to scan a file. A file could either pass or be blocked according to this setting.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Sophos Antivirus Configuration Overview on page 5720

exception (Security Antivirus Mime Whitelist)

Syntax exception *listname*;

Hierarchy Level [edit security utm feature-profile anti-virus mime-whitelist]
[edit security utm feature-profile anti-virus mime-whitelist list *listname*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the antivirus scanner to use an exception list to the MIME bypass list (custom objects). To use the exception list, you first create a whitelist custom-object list with the **list** statement. The system will first look at any existing whitelist mime pattern. If it matches an item, it will then continue to look for any exceptions to the whitelist and will then scan any item in the exception list.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

exception (Security Content Filtering)

Syntax exception *list-name*;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* block-mime]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the content filter to use an exception list to the MIME block list (custom objects).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

fallback-block (Security Antivirus)

Syntax
```
fallback-block {
  administrator-email email-address;
  allow-email;
  custom-message message;
  custom-message-subject message-subject;
  display-host;
  (notify-mail-sender | no-notify-mail-sender);
  type (message | protocol-only);
}
```

Hierarchy Level
```
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
  notification-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
  notification-options]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name
  notification-options]
```

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Configure notifications for fallback blocking actions. Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

fallback-non-block (Security Antivirus)

Syntax

```
fallback-non-block {
  custom-message message;
  custom-message-subject message-subject;
  (notify-mail-recipient | no-notify-mail-recipient);
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options]

Release Information

Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description

Configure notifications for fallback nonblocking actions.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

fallback-options (Security Antivirus Juniper Express Engine)

Syntax

```
fallback-options {
  content-size (block | log-and-permit);
  default (block | log-and-permit);
  engine-not-ready (block | log-and-permit);
  out-of-resources (block | (log-and-permit);
  timeout (block | log-and-permit);
  too-many-requests (block | log-and-permit);
}
```

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Fallback options tell the system how to handle the errors.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Express Antivirus Configuration Overview on page 5623

Copyright © 2016, Juniper Networks, Inc.

fallback-options (Security Antivirus Kaspersky Lab Engine)

Syntax

```
fallback-options {
  content-size (block | log-and-permit);
  corrupt-file (block | log-and-permit);
  decompress-layer (block | log-and-permit);
  default (block | log-and-permit);
  engine-not-ready (block | log-and-permit);
  out-of-resources (block | (log-and-permit);
  password-file (block | (log-and-permit);
  timeout (block | log-and-permit);
  too-many-requests (block | log-and-permit);
}
```

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Fallback options tell the system how to handle the errors returned by either the scan engine or the scan manager.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Full Antivirus Configuration Overview on page 5648

Copyright © 2016, Juniper Networks, Inc.

fallback-options (Security Antivirus Sophos Engine)

Syntax

```
fallback-options {
  content-size (block | log-and-permit | permit);
  default (block | log-and-permit | permit);
  engine-not-ready (block | log-and-permit | permit);
  out-of-resources (block | log-and-permit | permit);
  timeout (block | log-and-permit | permit);
  too-many-requests (block | log-and-permit | permit);
}
```

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure fallback options to instruct the system how to handle errors.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

fallback-settings (Security Web Filtering)

Syntax

```
fallback-settings {
  default (block | log-and-permit);
  server-connectivity (block | log-and-permit);
  timeout (block | log-and-permit);
  too-many-requests (block | log-and-permit);
}
```

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name*]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Fallback settings tell the system how to handle errors.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

fallback-settings (Security Web Filtering Juniper Local)

Syntax

```
fallback-settings {
  default (block | log-and-permit);
  server-connectivity (block | log-and-permit);
  timeout (block | log-and-permit);
  too-many-requests (block | log-and-permit);
}
```

Hierarchy Level [edit security utm feature-profile web-filtering juniper–local profile *profile-name*]

Release Information Statement introduced in Junos OS Release 10.0 .

Description Fallback settings tell the system how to handle errors.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Web Filtering Overview on page 5764

Copyright © 2016, Juniper Networks, Inc.

fallback-settings (Security Web Filtering Websense Redirect)

Syntax

```
fallback-settings {
  default (block | log-and-permit);
  server-connectivity (block | log-and-permit);
  timeout (block | log-and-permit);
  too-many-requests (block | log-and-permit);
}
```

Hierarchy Level [edit security utm feature-profile web-filtering websense-redirect profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Fallback settings tell the system how to handle errors.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Redirect Web Filtering on page 5805

Copyright © 2016, Juniper Networks, Inc.

feature-profile

Syntax

```
feature-profile {
  anti-spam {
    address-blacklist list-name;
    address-whitelist list-name;
    sbl {
      profile profile-name {
        custom-tag-string [string];
        (sbl-default-server | no-sbl-default-server);
        spam-action (block | tag-header | tag-subject);
      }
    }
    traceoptions flag flag;
  }
  anti-virus {
    juniper-express-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile profile-name {
        fallback-options {
          content-size (block | log-and-permit);
          default (block | log-and-permit);
          engine-not-ready (block | log-and-permit);
          out-of-resources (block | (log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        notification-options {
          fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
          fallback-non-block {
            custom-message message;
```


Copyright © 2016, Juniper Networks, Inc.

```
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
          }
          virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
        }
        scan-options {
          content-size-limit value;
          (intelligent-prescreening | no-intelligent-prescreening);
          timeout value;
        }
        trickling {
          timeout value;
        }
      }
    }
    kaspersky-lab-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile profile-name {
        fallback-options {
          content-size (block | log-and-permit);
          corrupt-file (block | log-and-permit);
          decompress-layer (block | log-and-permit);
          default (block | log-and-permit);
          engine-not-ready (block | log-and-permit);
          out-of-resources (block | (log-and-permit);
          password-file (block | (log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        notification-options {
          fallback-block {
            administrator-email email-address;
            allow-email;
            custom-message message;
            custom-message-subject message-subject;
            display-host;
```

```
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
          fallback-non-block {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-recipient | no-notify-mail-recipient);
          }
          virus-detection {
            custom-message message;
            custom-message-subject message-subject;
            (notify-mail-sender | no-notify-mail-sender);
            type (message | protocol-only);
          }
        }
        scan-options {
          content-size-limit value;
          decompress-layer-limit value;
          (intelligent-prescreening | no-intelligent-prescreening);
          scan-extension filename;
          scan-mode (all | by-extension);
          timeout value;
        }
        trickling {
          timeout value;
        }
      }
    }
    mime-whitelist {
      exception listname;
      list listname {
        exception listname;
      }
    }
    sophos-engine {
      pattern-update {
        email-notify {
          admin-email email-address;
          custom-message message;
          custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
          password password-string;
          port port-number;
          server address-or-url;
          username name;
        }
        url url;
      }
      profile <name> {
        fallback-options {
          content-size (block | log-and-permit | permit);
          default (block | log-and-permit | permit);
          engine-not-ready (block | log-and-permit | permit);
```


Copyright © 2016, Juniper Networks, Inc.

```
            out-of-resources (block | log-and-permit | permit);
            timeout (block | log-and-permit | permit);
            too-many-requests (block | log-and-permit | permit);
          }
          notification-options {
            fallback-block {
              administrator-email email-address;
              allow-email;
              custom-message message;
              custom-message-subject message-subject;
              display-host;
              (notify-mail-sender | no-notify-mail-sender);
              type (message | protocol-only);
            }
            fallback-non-block {
              custom-message message;
              custom-message-subject message-subject;
              (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
              custom-message message;
              custom-message-subject message-subject;
              (notify-mail-sender | no-notify-mail-sender);
              type (message | protocol-only);
            }
          }
          scan-options {
            content-size-limit value;
            (no-uri-check | uri-check);
            timeout value;
          }
          trickling {
            timeout value;
          }
        }
        sxl-retry value;
        sxl-timeout seconds;
      }
      traceoptions flag flag;
      type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
      url-whitelist listname;
    }
    content-filtering {
      profile profile-name {
        block-command protocol-command-list;
        block-content-type (activex | exe | http-cookie | java-applet | zip);
        block-extension extension-list;
        block-mime {
          exception list-name;
          list list-name;
        }
        notification-options {
          custom-message message;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
```

```
        permit-command protocol-command-list;
      }
      traceoptions flag flag;
    }
    web-filtering {
      url-whitelist custwhitelist;
      url-blacklist custblacklist;
      http-reassemble;
      type juniper-enhanced;
      juniper-enhanced {
        cache {
          timeout 1800;
          size 500;
        }
        server {
          host rp.cloud.threatseeker.com;
          port 80;
        }
        profile junos-wf-enhanced-default {
          category {
            Enhanced_Hacking {
              action log-and-permit;
              }
              Enhanced_Government {
              action quarantine;
            }
          }
          site-reputation-action {
            very-safe permit;
            moderately-safe log-and-permit;
            fairly-safe log-and-permit;
            harmful block;
            suspicious block;
          }
          default block;
          custom-block-message "***access denied ***";
          fallback-settings {
            default block;
            server-connectivity block;
            timeout block;
            too-many-requests block;
          }
          timeout 10;
          no-safe-search;
        }
        utm-policy mypolicy {
          web-filtering {
        http-profile my_ewfprofile01;
      }
      }
    }
  }
```

Hierarchy Level [edit security utm]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Release 9.5 .

Description Configure UTM features, antivirus, antispam, content-filtering, and web-filtering by creating feature profiles.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

filename-extension

Syntax

```
filename-extension object-name {
  value [value];
}
```

Hierarchy Level [edit security utm custom-objects]

Release Information Statement introduced in Junos OS Release 9.5 .

Description When scanning content, you can use a file extension list to define a set of file extensions that are used in file extension scan mode (scan-by-extension).

Options
- ***object-name***—Name of the extension-list object.
- **value *value***—Value of the extension-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

flag (SMTP)

Syntax

```
flag {
  all;
  configuration;
  IPC;
  protocol-exchange;
  send-request;
}
```

Hierarchy Level [edit smtp traceoptions]

Release Information Statement added in Junos OS Release 10.0.

Description Set flag for the SMTP traceoptions.

Options The following flag options are supported:

- **IPC**—Trace interprocesss communication.
- **all**—Trace everything.
- **configuration**—Trace configuration event.
- **protocol-exchange**—Trace SMTP protocol exchanges.
- **send-request**—Trace send mail request event.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- SMTP Configuration Statement Hierarchy on page 5825

Copyright © 2016, Juniper Networks, Inc.

format (Security Log Stream)

Syntax format (binary | sd-syslog | syslog | welf)

Hierarchy Level [edit security log stream *stream-name*]

Release Information Statement introduced in Junos OS Release 10.0 . Updated in Junos OS Release 12.1 .

Description Set the format for remote security message logging to **binary**, **syslog** (system log), **sd-syslog** (structured system log), or **welf**. Note that for the WELF format, the category must be set to **content-security** (see category (Security Logging)).

Options

- **binary**—Binary encoded text to conserve resources.
- **sd-syslog**—Structured system log file.
- **syslog**—Traditional system log file.
- **welf**—Web Trends Extended Log Format.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *AppSecure Services Feature Guide for Security Devices*
- *Logical Systems Feature Guide for Security Devices*

Copyright © 2016, Juniper Networks, Inc.

from-zone (Security Policies)

Syntax

```
from-zone zone-name to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```


Copyright © 2016, Juniper Networks, Inc.

```
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            initial-tcp-mss mss-value;
            reverse-tcp-mss mss-value;
            sequence-check-required;
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

Hierarchy Level [edit security policies]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** option added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **description** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10. Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description Specify a source zone and destination zone to be associated with the security policy.

Options

- **from-zone** *zone-name*—Name of the source zone.
- **to-zone** *zone-name*—Name of the destination zone.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

ftp (UTM Policy Anti-Virus)

Syntax

```
ftp {
  download-profile profile-name;
  upload-profile profile-name;
}
```

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus FTP protocol and attach this policy to a security profile to implement it.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

Copyright © 2016, Juniper Networks, Inc.

ftp (UTM Policy Content Filtering)

Syntax

```
ftp {
  download-profile profile-name;
  upload-profile profile-name;
}
```

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering FTP protocol and attach this policy to a security profile to implement it.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

host (Security Web Filtering)

Syntax host *host-name*;

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated server]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* server]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* server]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Set server host parameters by entering the server name or IP address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

http-profile (Security Antivirus)

Syntax http-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

http-profile (Security Content Filtering)

Syntax http-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

http-profile (Security Web Filtering)

Syntax http-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* web-filtering]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a UTM policy for the Web-filtering HTTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Web Filtering Overview on page 5764

imap-profile (Security UTM Policy Antivirus)

Syntax imap-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus IMAP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

imap-profile (Security UTM Policy Content Filtering)

Syntax imap-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering IMAP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

intelligent-prescreening

Syntax intelligent-prescreening;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Enable intelligent prescreening.

Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if the scan engine finds that it is unlikely that the file is infected, it then determines that it is safe to bypass the normal scanning procedure.

You can disable intelligent prescreening with the **no-intelligent-prescreening** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

interval (Security Antivirus)

Syntax interval *value*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]

Release Information Statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1 .

Description Set the pattern data files auto-update interval. You can choose to leave the default interval value or you can change it by using this command. You can also force a manual update, if necessary.

NOTE: The data files used with Sophos are not typical virus pattern files; they are small files that help guide virus scanning logic. The full virus pattern database is stored on an external Sophos server called the Sophos Extensible List (SXL) server.

Options *value*—Pattern data files auto-update interval in minutes.

Range: 10 through 10,080 minutes (10 minutes through 7 days)

Default: For Juniper Express engine and Kaspersky Lab engine, 60 minutes; for Sophos engine, 1440 minutes (every 24 hours)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

ipc

Syntax

```
ipc {
  traceoptions flag flag;
}
```

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure trace options for IPC.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Enable trace for all IPC trace options.
  - **basic**—Trace basic IPC related information.
  - **connection-manager**—Trace IPC connection manager information.
  - **connection-status**—Trace IPC connection status information.
  - **detail**—Trace IPC related detailed information.
  - **pfe**—Trace communication with PFE.
  - **utm-realtime**—Trace IPC realtime-thread information.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

juniper-enhanced

Syntax

```
juniper-enhanced {
  cache {
    size value;
    timeout value;
  }
  profile profile-name {
    category customurl-list name {
      action (block | log-and-permit | permit | quarantine);
    }
    custom-block-message value;
    custom-quarantine-message value;
    default (block | log-and-permit | permit | quarantine);
    fallback-settings {
      default (block | log-and-permit);
      server-connectivity (block | log-and-permit);
      timeout (block | log-and-permit);
      too-many-requests (block | log-and-permit);
    }
    no-safe-search;
    site-reputation-action {
      fairly-safe (block | log-and-permit | permit | quarantine);
      harmful (block | log-and-permit | permit | quarantine);
      moderately-safe (block | log-and-permit | permit | quarantine);
      suspicious (block | log-and-permit | permit | quarantine);
      very-safe (block | log-and-permit | permit | quarantine);
      }
    timeout value;
  }
  server {
    host host-name;
    port number;
  }
}
```

Hierarchy Level [set security utm feature-profile web-filtering]

Release Information Statement introduced in Junos OS Release 11.4 .

Description Configure the UTM Enhanced Web Filtering feature.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Web Filtering Overview on page 5764

Copyright © 2016, Juniper Networks, Inc.

juniper-express-engine

Syntax

```
juniper-express-engine {
  pattern-update {
    email-notify {
      admin-email email-address;
      custom-message message;
      custom-message-subject message-subject;
    }
    interval value;
    no-autoupdate;
    proxy {
      password password-string;
      port port-number;
      server address-or-url;
      username name;
    }
    url url;
  }
  profile profile-name {
    fallback-options {
      content-size (block | log-and-permit);
      default (block | log-and-permit);
      engine-not-ready (block | log-and-permit);
      out-of-resources (block | (log-and-permit);
      timeout (block | log-and-permit);
      too-many-requests (block | log-and-permit);
    }
    notification-options {
      fallback-block {
        administrator-email email-address;
        allow-email;
        custom-message message;
        custom-message-subject message-subject;
        display-host;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
      }
      fallback-non-block {
        custom-message message;
        custom-message-subject message-subject;
        (notify-mail-recipient | no-notify-mail-recipient);
      }
      virus-detection {
        custom-message message;
        custom-message-subject message-subject;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
      }
    }
    scan-options {
      content-size-limit value;
      (intelligent-prescreening | no-intelligent-prescreening);
      timeout value;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        trickling {
          timeout value;
        }
      }
    }
```

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the UTM express antivirus feature.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Express Antivirus Configuration Overview on page 5623

Copyright © 2016, Juniper Networks, Inc.

juniper-local

Syntax

```
juniper-local {
  profile profile-name {
    custom-block-message value;
    default (block | log-and-permit | permit);
    fallback-settings {
      default (block | log-and-permit);
      server-connectivity (block | log-and-permit);
      timeout (block | log-and-permit);
      too-many-requests (block | log-and-permit);
    }
    timeout value;
  }
}
```

Hierarchy Level [set security utm feature-profile web-filtering]

Release Information Statement introduced in Junos OS Release 10.0 .

Description Configure the UTM Web-filtering local feature.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

limit (UTM Policy)

Syntax limit *value*;

Hierarchy Level [edit security utm utm-policy *policy-name* traffic-options sessions-per-client]

Release Information Statement introduced in Junos OS Release 9.5 .

Description In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle to limit sessions.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

list (Security Antivirus Mime Whitelist)

Syntax

```
list listname {
    exception listname;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus mime-whitelist]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the antivirus scanner to use MIME bypass lists (custom objects). If you want to have exceptions to the whitelist, create a mime-pattern list with the **exception** statement in addition to the **list** statement.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

list (Security Content Filtering Block Mime)

Syntax

```
list list-name;
```

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* block-mime]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the content filter to use MIME block lists (custom objects).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

log (Security)

Syntax

```
log {
  cache {
    exclude exclude-name {
      destination-address destination-address;
      destination-port destination-port;
      event-id event-id;
      failure;
      interface-name interface-name;
      policy-name policy-name;
      process process-name;
      protocol protocol;
      source-address source-address;
      source-port source-port;
      success;
      user-name user-name;
    }
    limit value;
  }
  disable;
  event-rate rate;
  facility-override;
  file {
    files max-file-number;
    name file-name;
    path binary-log-file-path;
    size maximum-file-size;
  }
  format (binary | sd-syslog | syslog);
  mode (event | stream);
  rate-cap rate-cap-value;
  (source-address source-address | source-interface interface-name);
  stream stream-name {
    category (all | content-security);
    format (binary | sd-syslog | syslog | welf);
    host {
      ip-address;
      port port-number;
    }
    severity (alert | critical | debug | emergency | error | info | notice | warning);
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
  transport {
```


Copyright © 2016, Juniper Networks, Inc.

```
    protocol (udp | tcp | tls);
    tls-profile tls-profile-name;
    tcp-connections tcp-connections;
  }
  utc-time-stamp;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.2.

Description You can set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options

- **disable**—Disable the security logging for the device.
- **event-rate** *rate*—Limits the rate (0 through 1500) at which logs will be streamed per second.
- **rate-cap** *rate-cap-value*—Works with event mode only. Limits the rate (0 through 5000) at which data plane logs will be generated per second.
- **source-address** *source-address*—Specify a source IP address or IP address used when exporting security logs.
- **source-interface** *interface-name*—Specify a source interface name, which is mandatory to configure **stream**.

NOTE: The **source-address** and **source-interface** are alternate values. Using one of the options is mandatory.

- **utc-time-stamp**—Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

mime-pattern

Syntax

```
mime-pattern object-name {
  value [value];
}
```

Hierarchy Level [edit security utm custom-objects]

Release Information Statement introduced in Junos OS Release 9.5.

Description The gateway device uses MIME (Multipurpose Internet Mail Extension) types to decide which traffic is allowed to bypass various types of scanning.

Options

- ***object-name***—Name of the MIME object.
- **value *value***—Value of the MIME object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

 Copyright © 2016, Juniper Networks, Inc.

mime-whitelist

Syntax

```
mime-whitelist {
  exception listname;
  list listname {
    exception listname;
  }
}
```

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 9.5. Statement updated for Sophos antivirus support in Junos OS Release 11.1.

Description Configure the antivirus scanner to use MIME bypass lists and exception lists. You can use your own custom object lists, or you can use the default list that ships with the device called junos-default-bypass-mime.

WARNING: When you configure the MIME whitelist feature, be aware that, because header information in HTTP traffic can be spoofed, you cannot always trust HTTP headers to be legitimate. When a Web browser is determining the appropriate action for a given file type, it detects the file type without checking the MIME header contents. However, the MIME whitelist feature does refer to the MIME encoding in the HTTP header. For these reasons, it is possible in certain cases for a malicious website to provide an invalid HTTP header. For example, a network administrator might inadvertently add a malicious website to a MIME whitelist, and, because the site is in the whitelist, it will not be blocked by Sophos even though Sophos has identified the site as malicious in its database. Internal hosts would then be able to reach this site and could become infected.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

no-autoupdate

Syntax no-autoupdate;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]

Release Information Statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1 .

Description Turn off automatic data file (pattern file) update for the Kaspersky Lab, Juniper Express, or Sophos engines.

NOTE: The data files used with Sophos are not typical virus pattern files; they are small files that help guide virus scanning logic. The full virus pattern database is stored on an external Sophos server called the Sophos Extensible List (SXL) server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

no-intelligent-prescreening

Syntax no-intelligent-prescreening;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Disables intelligent prescreening.

Intelligent prescreening tells the antivirus module to begin scanning a file much earlier. In this case, the scan engine uses the first packet or the first several packets to determine if a file could possibly contain malicious code. The scan engine does a quick check on these first packets and if the scan engine finds that it is unlikely that the file is infected, it then determines that it is safe to bypass the normal scanning procedure.

You can enable intelligent prescreening with the **intelligent-prescreening** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

no-notify-mail-recipient

Syntax no-notify-mail-recipient;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-non-block]

Release Information Statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Description Do not notify the e-mail recipient about errors returned by the antivirus scan engine when a fallback nonblocking action occurs.

You can specify that the e-mail recipient is to be notified with the **notify-mail-recipient** statement.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

no-notify-mail-sender (Security Content Filtering Notification Options)

Syntax no-notify-mail-sender;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* notification-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Do not notify the e-mail sender.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

no-notify-mail-sender (Security Fallback Block)

Syntax no-notify-mail-sender;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Do not notify the e-mail sender about errors returned by the antivirus scan engine when a fallback action occurs.

You can specify that the e-mail sender is to be notified with the **notify-mail-sender** statement.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

no-notify-mail-sender (Security Virus Detection)

Syntax no-notify-mail-sender;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description Do not notify the e-mail sender when a virus is detected by the antivirus engine.

You can specify that the e-mail sender is to be notified with the **notify-mail-sender** statement.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

no-sbl-default-server

Syntax no-sbl-default-server;

Hierarchy Level [edit security utm feature-profile anti-spam sbl profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Disable the default SBL server lookup.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Antispam Filtering Overview on page 5599

Copyright © 2016, Juniper Networks, Inc.

notification-options (Security Antivirus)

Syntax

```
notification-options {
  fallback-block {
    administrator-email email-address;
    allow-email;
    custom-message message;
    custom-message-subject message-subject;
    display-host;
    (notify-mail-sender | no-notify-mail-sender);
    type (message | protocol-only);
  }
  fallback-non-block {
    custom-message message;
    custom-message-subject message-subject;
    (notify-mail-recipient | no-notify-mail-recipient);
  }
  virus-detection {
    custom-message message;
    custom-message-subject message-subject;
    (notify-mail-sender | no-notify-mail-sender);
    type (message | protocol-only);
  }
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name*]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name*]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description

There are multiple notification options you can configure to trigger when a virus is detected.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

notification-options (Security Content Filtering)

Syntax

```
notification-options {
  custom-message message;
  (notify-mail-sender | no-notify-mail-sender);
  type (message | protocol-only);
}
```

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description You can configure a message notification to trigger when a content filter is matched.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

notify-mail-recipient

Syntax notify-mail-recipient;

Hierarchy Level

```
[edit security utm feature-profile anti-virus juniper-express-engine profile profile-name
  notification-options fallback-non-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile profile-name
  notification-options fallback-non-block]
[edit security utm feature-profile anti-virus sophos-engine profile profile-name
  notification-options fallback-non-block]
```

Release Information Statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Description Notify the e-mail recipient about errors returned by the antivirus scan engine when a fallback nonblocking action occurs.

You can specify that the e-mail recipient is not to be notified with the **no-notify-mail-recipient** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

notify-mail-sender (Security Content Filtering Notification Options)

Syntax notify-mail-sender;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* notification-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Notify the e-mail sender.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

notify-mail-sender (Security Fallback Block)

Syntax notify-mail-sender;

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description E-mail notification is used to notify the sender or the recipient about the errors returned by either the scan engine or the scan manager when a fallback action occurs.

You can specify that the sender is not to be notified with the **no-notify-mail-sender** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

notify-mail-sender (Security Virus Detection)

Syntax notify-mail-sender;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description E-mail notification is used to notify the sender or the recipient about the detected viruses or the scanning errors. When a virus is detected, an e-mail is sent to the sender upon virus detection.

You can specify that the sender is not to be notified with the **no-notify-mail-sender** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

no-uri-check

Syntax no-uri-check;

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Do not perform Sophos antivirus Uniform Resource Identifier (URI) checking. URI checking is performed by analyzing HTTP traffic URI content against a remote Sophos database server to identify malware or malicious content. URI checking is on by default.

You can enable Sophos antivirus URI checking with the **uri-check** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Sophos Antivirus Configuration Overview on page 5720

Copyright © 2016, Juniper Networks, Inc.

out-of-resources

Syntax out-of-resources (block | (log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. When out-of-resources occurs, scanning is aborted. The default action is block.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

out-of-resources (Security Antivirus Sophos Engine)

Syntax default (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Release 11.1 .

Description Virus scanning requires a great deal of memory and CPU resources. Due to resource constraints, memory allocation requests can be denied by the system. When out-of-resources occurs, scanning is aborted.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Sophos Antivirus Configuration Overview on page 5720

over-limit

Syntax over-limit (block | log-and-permit);

Hierarchy Level [edit security utm utm-policy *policy-name* traffic-options sessions-per-client]

Release Information Statement introduced in Junos OS Release 9.5 .

Description In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle to limit sessions and configure an action to occur when the limit is exceeded.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

packet-filter

Syntax

```
packet-filter packet-filter-name {
  action-profile (profile-name | default);
  destination-port (port-range | protocol-name);
  destination-prefix destination-prefix;
  interface logical-interface-name;
  protocol (protocol-number | protocol-name;
  source-port (port-range | protocol- name);
  source-prefix source-prefix;
  }
```

Hierarchy Level [edit security datapath-debug]

Release Information Command introduced in Junos OS Release 9.6 ; Support for IPv6 addresses for the **destination-prefix** and **source-prefix** options added in Junos OS Release 10.4. Support for IPv6 filter for the **interface** option added in Junos OS Release 10.4.

Description Set packet filter for taking the datapath-debug action. A maximum of four filters are supported at the same time.

Options

- **action-profile (*profile-name* | default)**—Identify the action profile to use. You can specify the name of the action profile to use or select default action profile.
- **destination-port (*port-range* | *protocol name*)**—Specify a destination port to match TCP/UDP destination port.
- **destination-prefix *destination-prefix***—Specify a destination IPv4/IPv6 address prefix.
- **interface *logical-interface-name***—Specify a logical interface name.
- **protocol (*protocol-number* | *protocol-name***—Match IP protocol type.
- **source-port (*port-range* | *protocol-name***—Match TCP/UDP source port.
- **source-prefix *source-prefix***—Specify a source IP address prefix.

Required Privilege Level security—To view this statement in the configuration
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603

Copyright © 2016, Juniper Networks, Inc.

password (Security Antivirus)

Syntax password *password-string*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]

Release Information Statement introduced in Release 11.2 .

Description Set the password for the proxy server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

password-file

Syntax password-file (block | (log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Description Password protected file is the error returned by the scan engine when the scanned file is protected by a password. The default action is log-and-permit.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Full Antivirus Configuration Overview on page 5648

Copyright © 2016, Juniper Networks, Inc.

pattern-update (Security Antivirus)

Syntax

```
pattern-update {
  email-notify {
    admin-email email-address;
    custom-message message;
    custom-message-subject message-subject;
  }
  interval value;
  no-autoupdate;
  proxy {
    password password-string;
    port port-number;
    server address-or-url;
    username name;
  }
  url url;
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine]
[edit security utm feature-profile anti-virus kaspersky-lab-engine]
[edit security utm feature-profile anti-virus sophos-engine]

Release Information

Statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1 .

Description

Updates to the pattern file are added as new viruses are discovered. You can configure the security device to regularly update the pattern file automatically, or you can update the file manually.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

permit-command

Syntax permit-command *protocol-command-list*;

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Apply protocol permit command custom-objects to the content-filtering profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

policies

Syntax

```
policies {
  default-policy (deny-all | permit-all);
  from-zone zone-name to-zone zone-name {
    policy policy-name {
      description description;
      match {
        application {
          [application];
          any;
        }
        destination-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-address {
          [address];
          any;
          any-ipv4;
          any-ipv6;
        }
        source-identity {
          [role-name];
          any;
          authenticated-user;
          unauthenticated-user;
          unknown-user;
        }
      }
      scheduler-name scheduler-name;
      then {
        count {
          alarm {
            per-minute-threshold number;
            per-second-threshold number;
          }
        }
        deny;
        log {
          session-close;
          session-init;
        }
        permit {
          application-services {
            application-firewall {
              rule-set rule-set-name;
            }
            application-traffic-control {
              rule-set rule-set-name;
            }
            gprs-gtp-profile profile-name;
```

Copyright © 2016, Juniper Networks, Inc.

```
            gprs-sctp-profile profile-name;
            idp;
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            user-firewall {
              access-profile profile-name;
              domain domain-name
              ssl-termination-profile profile-name;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
}
global {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
```


Copyright © 2016, Juniper Networks, Inc.

```
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    from-zone {
      [zone-name];
      any;
    }
    source-address {
      [address];
      any;
      any-ipv4;
      any-ipv6;
    }
    source-identity {
      [role-name];
      any;
      authenticated-user;
      unauthenticated-user;
      unknown-user;
    }
    to-zone {
      [zone-name];
      any;
    }
  }
  scheduler-name scheduler-name;
  then {
    count {
      alarm {
        per-minute-threshold number;
        per-second-threshold number;
      }
    }
    deny;
    log {
      session-close;
      session-init;
    }
    permit {
      application-services {
        application-firewall {
          rule-set rule-set-name;
        }
        application-traffic-control {
          rule-set rule-set-name;
        }
        gprs-gtp-profile profile-name;
        gprs-sctp-profile profile-name;
        idp;
        redirect-wx | reverse-redirect-wx;
        ssl-proxy {
          profile-name profile-name;
        }
        uac-policy {
```

```
                captive-portal captive-portal;
              }
              utm-policy policy-name;
            }
            destination-address {
              drop-translated;
              drop-untranslated;
            }
            firewall-authentication {
              pass-through {
                access-profile profile-name;
                client-match user-or-group-name;
                ssl-termination-profile profile-name;
                web-redirect;
                web-redirect-to-https;
              }
              web-authentication {
                client-match user-or-group-name;
              }
            }
            services-offload;
            tcp-options {
              initial-tcp-mss mss-value;
              reverse-tcp-mss mss-value;
              sequence-check-required;
              syn-check-required;
            }
          }
          reject;
        }
      }
    }
    policy-rematch;
    policy-stats {
      system-wide (disable | enable) ;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
    }
  }
```

Hierarchy Level [edit security]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 8.5.
Support for the **services-offload** option added in Junos OS Release 11.4.
Support for the **source-identitiy** option added in Junos OS Release 12.1.
Support for the **description** option added in Junos OS Release 12.1.
Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.
Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.
Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description Configure network security policies.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- Security Policies Overview on page 1085
- [edit security policies] Hierarchy Level on page 320

pop3-profile (Security UTM Policy Antivirus)

Syntax pop3-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus POP3 protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

pop3-profile (Security UTM Policy Content Filtering)

Syntax pop3-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content filtering POP3 protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

port (Security Antivirus)

Syntax port *port-number*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]

Release Information Statement introduced in Junos OS Release 11.2 .

Description Set the port number for the proxy server.

Options **Range:** 0 through 65,535

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

port (Security Web Filtering Server)

Syntax port *number*;

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated server]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* server]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* server]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Enter the port number for communicating with the server. (Default ports are 80, 8080, and 8081.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

primary-server

Syntax
```
primary-server {
  address ipv4-address;
  login sender-email-address {
    password password;
  }
}
```

Hierarchy Level [edit smtp]

Release Information Statement added in Junos OS Release 10.0.

Description Configure Simple Mail Transfer Protocol (SMTP) primary server for access authorization for SMTP requests.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- SMTP Configuration Statement Hierarchy on page 5825

Copyright © 2016, Juniper Networks, Inc.

profile (Security Antispam SBL)

Syntax

```
profile profile-name {
  custom-tag-string [string];
  (sbl-default-server | no-sbl-default-server);
  spam-action (block | tag-header | tag-subject);
}
```

Hierarchy Level [edit security utm feature-profile anti-spam sbl]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Create a profile for the antispam sbl feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

profile (Security Antivirus Juniper Express Engine)

Syntax

```
profile profile-name {
  fallback-options {
    content-size (block | log-and-permit);
    default (block | log-and-permit);
    engine-not-ready (block | log-and-permit);
    out-of-resources (block | (log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  notification-options {
    fallback-block {
      administrator-email email-address;
      allow-email;
      custom-message message;
      custom-message-subject message-subject;
      display-host;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
    fallback-non-block {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-recipient | no-notify-mail-recipient);
    }
    virus-detection {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
  }
  scan-options {
    content-size-limit value;
    (intelligent-prescreening | no-intelligent-prescreening);
    timeout value;
  }
  trickling {
    timeout value;
  }
}
```

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Create a profile for the Juniper express engine. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Express Antivirus Configuration Overview on page 5623

Copyright © 2016, Juniper Networks, Inc.

profile (Security Antivirus Kaspersky Lab Engine)

Syntax

```
profile profile-name {
  fallback-options {
    content-size (block | log-and-permit);
    corrupt-file (block | log-and-permit);
    decompress-layer (block | log-and-permit);
    default (block | log-and-permit);
    engine-not-ready (block | log-and-permit);
    out-of-resources (block | (log-and-permit);
    password-file (block | (log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  notification-options {
    fallback-block {
      administrator-email email-address;
      allow-email;
      custom-message message;
      custom-message-subject message-subject;
      display-host;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
    fallback-non-block {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-recipient | no-notify-mail-recipient);
    }
    virus-detection {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
  }
  scan-options {
    content-size-limit value;
    decompress-layer-limit value;
    (intelligent-prescreening | no-intelligent-prescreening);
    scan-extension filename;
    scan-mode (all | by-extension);
    timeout value;
  }
  trickling {
    timeout value;
  }
}
```

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine]

Release Information Statement introduced in Junos OS Release 9.5 .

Copyright © 2016, Juniper Networks, Inc.

Description Create a profile for the Kaspersky Lab engine. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- *kaspersky-lab-engine*

profile (Security Content Filtering)

Syntax

```
profile profile-name {
  block-command protocol-command-list;
  block-content-type (activex | exe | http-cookie | java-applet | zip);
  block-extension extension-list;
  block-mime {
    exception list-name;
    list list-name;
  }
  notification-options {
    custom-message message;
    (notify-mail-sender | no-notify-mail-sender);
    type (message | protocol-only);
  }
  permit-command protocol-command-list;
}
```

Hierarchy Level [edit security utm feature-profile content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Create a profile for the content-filtering feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Content Filtering Overview on page 5745

Copyright © 2016, Juniper Networks, Inc.

profile (Security Sophos Engine Antivirus)

Syntax

```
profile <name> {
  fallback-options {
    content-size (block | log-and-permit | permit);
    default (block | log-and-permit | permit);
    engine-not-ready (block | log-and-permit | permit);
    out-of-resources (block | log-and-permit | permit);
    timeout (block | log-and-permit | permit);
    too-many-requests (block | log-and-permit | permit);
  }
  notification-options {
    fallback-block {
      administrator-email email-address;
      allow-email;
      custom-message message;
      custom-message-subject message-subject;
      display-host;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
    fallback-non-block {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-recipient | no-notify-mail-recipient);
    }
    virus-detection {
      custom-message message;
      custom-message-subject message-subject;
      (notify-mail-sender | no-notify-mail-sender);
      type (message | protocol-only);
    }
  }
  scan-options {
    content-size-limit value;
    (no-uri-check | uri-check);
    timeout value;
  }
  trickling {
    timeout value;
  }
}
```

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Create a profile for the Sophos antivirus engine. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830
- Sophos Antivirus Configuration Overview on page 5720

profile (Security Web Filtering Juniper Enhanced)

Syntax

```
profile profile-name {
  category customurl-list name {
    action (block | log-and-permit | permit | quarantine);
  }
  custom-block-message value;
  custom-quarantine-message value;
  default (block | log-and-permit | permit | quarantine);
  fallback-settings {
    default (block | log-and-permit);
    server-connectivity (block | log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  no-safe-search;
  site-reputation-action {
    fairly-safe (block | log-and-permit | permit | quarantine);
    harmful (block | log-and-permit | permit | quarantine);
    moderately-safe (block | log-and-permit | permit | quarantine);
    suspicious (block | log-and-permit | permit | quarantine);
    very-safe (block | log-and-permit | permit | quarantine);
  }
  timeout value;
}
```

Hierarchy Level
[edit security utm feature-profile web-filtering juniper-enhanced]

Release Information
Statement introduced in Junos OS Release 11.4 .

Description
Create a profile for the juniper-enhanced feature. This profile includes all subsequent configuration options.

Options
The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Monitoring Web Filtering Configurations on page 5815

Copyright © 2016, Juniper Networks, Inc.

profile (Security Web Filtering Juniper Local)

Syntax

```
profile profile-name {
  custom-block-message value;
  default (block | log-and-permit | permit);
  fallback-settings {
    default (block | log-and-permit);
    server-connectivity (block | log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  timeout value;
}
```

Hierarchy Level [edit security utm feature-profile web-filtering juniper-local]

Release Information Statement introduced in Junos OS Release 10.0 .

Description Create a profile for the web-filtering juniper–local feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Monitoring Web Filtering Configurations on page 5815
- Example: Configuring Local Web Filtering on page 5799

Copyright © 2016, Juniper Networks, Inc.

profile (Security Web Filtering Surf Control Integrated)

Syntax

```
profile profile-name {
  category customurl-list name {
    action (block | log-and-permit | permit);
  }
  custom-block-message value;
  default (block | log-and-permit | permit);
  fallback-settings {
    default (block | log-and-permit);
    server-connectivity (block | log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  timeout value;
}
```

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Create a profile for the web-filtering surf-control-integrated feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Monitoring Web Filtering Configurations on page 5815

Copyright © 2016, Juniper Networks, Inc.

profile (Security Web Filtering Websense Redirect)

Syntax

```
profile profile-name {
  account value;
  custom-block-message value;
  fallback-settings {
    default (block | log-and-permit);
    server-connectivity (block | log-and-permit);
    timeout (block | log-and-permit);
    too-many-requests (block | log-and-permit);
  }
  server {
    host host-name;
    port number;
  }
  sockets value;
  timeout value;
}
```

Hierarchy Level [security utm feature-profile web-filtering websense-redirect]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Create a profile for the web-filtering web-sense feature. This profile includes all subsequent configuration options.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Monitoring Web Filtering Configurations on page 5815

Copyright © 2016, Juniper Networks, Inc.

protocol-command

Syntax

```
protocol-command object-name {
  value [value];
}
```

Hierarchy Level [edit security utm custom-objects]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level.

Options

- ***object-name***—Name of the command-list object.
- **value *value***—Value of the command-list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding UTM Custom Objects on page 5585

proxy (Security Antivirus)

Syntax

```
proxy {
  password password-string;
  port port-number;
  server address-or-url;
  username name;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]

Release Information Statement introduced in Junos OS Release 11.2 .

Description Update the pattern file on the proxy server.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

quarantine-message (Security UTM)

Syntax

```
quarantine-message {
  type {
    custom-redirect-url;
  }
  url url;
}
```

Hierarchy Level [edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X44-D10 for Enhanced Web Filtering.

Description Configure Juniper enhanced quarantine message settings.

Options

- **type**—Specify the following type of the quarantine message:
  - **custom-redirect-url**—Specify Custom redirect URL server.
- **url *url***—Specify an URL of the quarantine message.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

sbl

Syntax

```
sbl {
  profile profile-name {
    custom-tag-string [string];
    (sbl-default-server | no-sbl-default-server);
    spam-action (block | tag-header | tag-subject);
  }
}
```

Hierarchy Level [edit security utm feature-profile anti-spam]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure UTM server-based antispam features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

sbl-default-server

Syntax sbl-default-server;

Hierarchy Level [edit security utm feature-profile anti-spam sbl profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Enable the default SBL server lookup. You should enable this feature if you are using server-based spam filtering. (The SBL server is predefined on the device. It ships with the name and address of the SBL server.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

scan-extension

Syntax scan-extension *filename*;

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description For antivirus file extension scanning, configure the scan extension setting by specifying the name of the defined file extension list.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

scan-mode

Syntax scan-mode (all | by-extension);

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description You can scan all content or scan content with specific file extensions. You can use a file extension list to define a set of file extensions that are used in file extension scan mode. The antivirus module can then only scan files with extensions on the scan-extension list.

Options
- **all**—Scan all files.
- **by-extension**—Scan only files with extensions specified in a file extension list custom object.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

scan-options (Security Antivirus Juniper Express Engine)

Syntax

```
scan-options {
  content-size-limit value;
  (intelligent-prescreening | no-intelligent-prescreening);
  timeout value;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

scan-options (Security Antivirus Kaspersky Lab Engine)

Syntax

```
scan-options {
  content-size-limit value;
  decompress-layer-limit value;
  (intelligent-prescreening | no-intelligent-prescreening);
  scan-extension filename;
  scan-mode (all | by-extension);
  timeout value;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

scan-options (Security Antivirus Sophos Engine)

Syntax

```
scan-options {
  content-size-limit value;
  (no-uri-check | uri-check);
  timeout value;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure the antivirus feature to scan specific types of traffic based on various scanning configuration parameters.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

secondary-server

Syntax

```
secondary-server {
  address ipv4-address;
  login sender-email-address {
    password password;
  }
}
```

Hierarchy Level [edit smtp]

Release Information Statement added in Junos OS Release 10.0.

Description Configure Simple Mail Transfer Protocol (SMTP) secondary server for access authorization for SMTP requests.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- SMTP Configuration Statement Hierarchy on page 5825

Copyright © 2016, Juniper Networks, Inc.

server (Security Antivirus)

Syntax server *address-or-url*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]

Release Information Statement introduced in Junos OS Release 11.2 .

Description Set the IP address or URL for the proxy server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

server (Security Web Filtering)

Syntax
```
server {
    host host-name;
    port number;
}
```

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name*]
[edit security utm feature-profile web-filtering juniper-enhanced]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Set server parameters by entering the server name or IP address.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

server-connectivity

Syntax server-connectivity (block | log-and-permit);

Hierarchy Level
[edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* fallback-settings]

Release Information
Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Fallback settings tell the system how to handle errors. This is the action that occurs when a request fails for this reason.

Options
- block—Log the error and deny the traffic
- log-and-permit—Log the error and permit the traffic

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

sessions-per-client

Syntax

```
sessions-per-client {
  limit value;
  over-limit (block | log-and-permit);
}
```

Hierarchy Level [edit security utm utm-policy *policy-name* traffic-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle.

NOTE: The sessions-per-client limit command supports the antispam, content filtering, and antivirus UTM features. It does not support Web filtering.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

site-reputation-action

Syntax

```
site-reputation-action {
  harmful (block | log-and-permit | permit | quarantine);
  fairly-safe (block | log-and-permit | permit | quarantine);
  moderately-safe (block | log-and-permit | permit | quarantine);
  suspicious (block | log-and-permit | permit | quarantine);
  very-safe (block | log-and-permit | permit | quarantine);
}
```

Hierarchy Level [edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* category *category-name*]

Release Information Statement introduced in Junos OS Release 11.4 .

Description Specify the action to be taken depending on the site reputation returned for all types of URLs whether it is categorized or uncategorized.

Options

fairly-safe —Permit, log-and-permit, block or quarantine a request if a site-reputation of 70 through 79 is returned.

harmful—Permit, log-and-permit, block or quarantine a request if a site-reputation of zero through 59 is returned.

moderately-safe—Permit, log-and-permit, block or quarantine a request if a site-reputation of 80 through 89 is returned.

suspicious—Permit, log-and-permit, block or quarantine a request if a site-reputation of 60 through 69 is returned.

very-safe—Permit, log-and-permit, block or quarantine a request if a site-reputation of 90 through 100 is returned.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

size (Security Web Filtering Cache)

Syntax size *value*;

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated cache]
[edit security utm feature-profile web-filtering juniper-enhanced cache]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Set the cache size parameters for Web filtering.

Options **Range:** 0 through 4096 kilobytes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

smtp-profile (Security UTM Policy Antispam)

Syntax smtp-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-spam]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antispam SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

smtp-profile (Security UTM Policy Antivirus)

Syntax smtp-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

smtp-profile (Security UTM Policy Content Filtering)

Syntax smtp-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering SMTP protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

sockets

Syntax sockets *value*;

Hierarchy Level [edit security utm feature-profile web-filtering websense-redirect profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Enter the number of sockets used for communicating between the client and server. The default is 1.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

sophos-engine

Syntax

```
sophos-engine {
  pattern-update {
    email-notify {
      admin-email email-address;
      custom-message message;
      custom-message-subject message-subject;
    }
    interval value;
    no-autoupdate;
    proxy {
      password password-string;
      port port-number;
      server address-or-url;
      username name;
    }
    url url;
  }
  profile <name> {
    fallback-options {
      content-size (block | log-and-permit | permit);
      default (block | log-and-permit | permit);
      engine-not-ready (block | log-and-permit | permit);
      out-of-resources (block | log-and-permit | permit);
      timeout (block | log-and-permit | permit);
      too-many-requests (block | log-and-permit | permit);
    }
    notification-options {
      fallback-block {
        administrator-email email-address;
        allow-email;
        custom-message message;
        custom-message-subject message-subject;
        display-host;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
      }
      fallback-non-block {
        custom-message message;
        custom-message-subject message-subject;
        (notify-mail-recipient | no-notify-mail-recipient);
      }
      virus-detection {
        custom-message message;
        custom-message-subject message-subject;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
      }
    }
    scan-options {
      content-size-limit value;
      (no-uri-check | uri-check);
      timeout value;
```

```
    }
    trickling {
      timeout value;
    }
  }
  sxl-retry value;
  sxl-timeout seconds;
}
```

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure the UTM Sophos antivirus feature.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

spam-action

Syntax spam-action (block | tag-header | tag-subject);

Hierarchy Level [edit security utm feature-profile anti-spam sbl profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the action to be taken by the device when spam is detected.

Options
- **block**— Block e-mail.
- **tag-header**—Tag header of e-mail.
- **tag-subject**—Tag subject of e-mail.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Server-Based Antispam Filtering on page 5603
- Example: Configuring Local List Antispam Filtering on page 5610

Copyright © 2016, Juniper Networks, Inc.

surf-control-integrated

Syntax

```
surf-control-integrated {
  cache {
    size value;
    timeout value;
  }
  profile profile-name {
      category customurl-list name {
        action (block | log-and-permit | permit);
      }
    custom-block-message value;
    default (block | log-and-permit | permit);
    fallback-settings {
      default (block | log-and-permit);
      server-connectivity (block | log-and-permit);
      timeout (block | log-and-permit);
      too-many-requests (block | log-and-permit);
      }
      timeout value;
    }
  }
  server {
    host host-name;
    port number;
  }
}
```

Hierarchy Level [set security utm feature-profile web-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the UTM web-filtering integrated feature.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

sxl-retry

Syntax sxl-retry *value;*

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure the number of retry attempts to the remote Sophos Extensible List (SXL) server when a request timeout occurs.

Options ***value*** —Number of retries.
Range: 0 through 5
Default: 1

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

sxl-timeout

Syntax sxl-timeout *seconds;*

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Configure the timeout value for responses to a Sophos checksum or URI query.

Options ***seconds*** —Number of seconds before timeout occurs.
Range: 1 through 5 seconds
Default: 2 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

timeout (Security Antivirus Fallback Options)

Syntax timeout (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is aborted and the content is either passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option. The default action is block.

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

timeout (Security Antivirus Fallback Options Sophos Engine)

Syntax default (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Scanning a complex file could consume resources and time. If the time it is taking to scan exceeds the timeout setting in the antivirus profile, the processing is aborted and the content is either passed or blocked without completing the virus checking. The decision is made based on the timeout fallback option.

Options

- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

timeout (Security Antivirus Scan Options)

Syntax timeout *value*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* scan-options]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description The scanning timeout value includes the time frame from when the scan request is generated to when the scan result is returned by the scan engine. The time range can be 1 to 1800 seconds. By default, it is 180 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

timeout (Security Web Filtering)

Syntax timeout *value*;

Hierarchy Level [edit security utm feature-profile web-filtering websense-redirect profile *profile-name*]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Enter a timeout limit for requests. Once this limit is reached, fail mode settings are applied. The default here is 15 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

timeout (Security Web Filtering Cache)

Syntax timeout *value*;

Hierarchy Level [edit security utm feature-profile web-filtering surf-control-integrated cache]
[edit security utm feature-profile web-filtering juniper-enhanced cache]

Release Information Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Set the cache timeout parameters for surf-control-integrated web filtering (24 hours is the default and the maximum allowed life span of cached items).

Options **Range:** 1 through 1800 minutes.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

timeout (Security Web Filtering Fallback Settings)

Syntax timeout (block | log-and-permit);

Hierarchy Level
[edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* fallback-settings]

Release Information
Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Fallback settings tell the system how to handle errors.

Options
- log-and-permit—Log the error and permit the traffic
- block—Log the error and deny the traffic

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

too-many-requests (Security Antivirus Fallback Options)

Syntax too-many-requests (block | log-and-permit);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* fallback-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 9.5 .

Description If the total number of messages received concurrently exceeds 4000, the content is either passed or blocked depending on the too-many-request fallback option. The default action is block. (The allowed request limit is not configurable.)

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

too-many-requests (Security Antivirus Fallback Options Sophos Engine)

Syntax default (block | log-and-permit | permit);

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* fallback-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description If the total number of messages received concurrently exceeds the device limits, the content is either passed or blocked depending on the too-many-request fallback option. (The allowed request limit is not configurable.)

Options
- **block**—Log the error and deny the traffic
- **log-and-permit**—Log the error and permit the traffic
- **permit**—Permit the traffic

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

too-many-requests (Security Web Filtering Fallback Settings)

Syntax too-many-requests (block | log-and-permit);

Hierarchy Level
[edit security utm feature-profile web-filtering surf-control-integrated profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering websense-redirect profile *profile-name* fallback-settings]
[edit security utm feature-profile web-filtering juniper-enhanced profile *profile-name* fallback-settings]

Release Information
Statement introduced in Junos OS Release 9.5 .
Statement introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description If the total number of messages received concurrently exceeds the device limits, the content is either passed or blocked depending on the too-many-request fallback option. The default action is BLOCK. (The allowed request limit is not configurable.)

Options
- block—Log the error and deny the traffic
- log-and-permit—Log the error and permit the traffic

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

to-zone (Security Policies)

Syntax

```
to-zone zone-name {
  policy policy-name {
    description description;
    match {
      application {
        [application];
        any;
      }
      destination-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-address {
        [address];
        any;
        any-ipv4;
        any-ipv6;
      }
      source-identity {
        [role-name];
        any;
        authenticated-user;
        unauthenticated-user;
        unknown-user;
      }
    }
    scheduler-name scheduler-name;
    then {
      count {
        alarm {
          per-minute-threshold number;
          per-second-threshold number;
        }
      }
      deny;
      log {
        session-close;
        session-init;
      }
      permit {
        application-services {
          application-firewall {
            rule-set rule-set-name;
          }
          application-traffic-control {
            rule-set rule-set-name;
          }
          gprs-gtp-profile profile-name;
          gprs-sctp-profile profile-name;
          idp;
```

Copyright © 2016, Juniper Networks, Inc.

```
            redirect-wx | reverse-redirect-wx;
            ssl-proxy {
              profile-name profile-name;
            }
            uac-policy {
              captive-portal captive-portal;
            }
            utm-policy policy-name;
          }
          destination-address {
            drop-translated;
            drop-untranslated;
          }
          firewall-authentication {
            pass-through {
              access-profile profile-name;
              client-match user-or-group-name;
              ssl-termination-profile profile-name;
              web-redirect;
              web-redirect-to-https;
            }
            web-authentication {
              client-match user-or-group-name;
            }
          }
          services-offload;
          tcp-options {
            sequence-check-required;
            syn-check-required;
          }
          tunnel {
            ipsec-group-vpn group-vpn;
            ipsec-vpn vpn-name;
            pair-policy pair-policy;
          }
        }
        reject;
      }
    }
  }
```

Hierarchy Level [edit security policies from-zone *zone-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **services-offload** and **junos-host** options added in Junos OS Release 11.4. Support for the **source-identity** option added in Junos OS Release 12.1. Support for the **ssl-termination-profile** and **web-redirect-to-https** options added in Junos OS Release 12.1X44-D10.

Description Specify a destination zone to be associated with the security policy.

Options

- **zone-name**—Name of the destination zone object.
- **junos-host**—Default security zone for self-traffic of the device.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

traceoptions (Security Antispam)

Syntax
traceoptions flag *flag*;

Hierarchy Level
[edit security utm feature-profile anti-spam]

Release Information
Statement introduced in Junos OS Release 9.5 .

Description
Define tracing operations for UTM antispam features.

Options
- *flag*:
  - **all**—Enable all antispam trace flags.
  - **manager** —Trace antispam manager information.
  - **sbl**—Trace SBL server information.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Antivirus)

Syntax traceoptions flag *flag*;

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Define tracing operations for UTM antivirus features.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Enable trace all antivirus trace options.
  - **basic**—Trace antivirus module generic basic information.
  - **detail**—Trace antivirus module generic detail information.
  - **engine**—Trace scan engine information.
  - **event**—Trace communication events between routing engine side processes.
  - **ipc**—Trace communication events with Packet Forwarding Engine.
  - **manager**—Trace antivirus manager process activities.
  - **pattern**—Trace detail information of pattern loading.
  - **sendmail**—Trace mail notifying process activities.
  - **statistics**—Trace statistics information.
  - **updater**—Trace pattern updater process activities.
  - **worker**—Trace antivirus worker process activities.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Application Proxy)

Syntax

```
traceoptions {
  flag flag;
}
```

Hierarchy Level [edit security utm application-proxy]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure tracing options for application proxy.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **abort**—Trace aborted sessions for application proxy.
  - **all**—Trace with all flags enabled.
  - **anti-virus**—Trace anti-virus information.
  - **application-objects**—Trace application-proxy objects information.
  - **basic**—Trace application-proxy related basic information.
  - **buffer**— Trace application-proxy data buffer information.
  - **connection-rating**—Trace connection rating information.
  - **detail**—Trace application-proxy related detailed information.
  - **express-anti-virus**—Trace anti-virus express engine information.
  - **ftp-control**—Trace FTP control connection information.
  - **ftp-data**—Trace FTP data connection information.
  - **http**—Trace HTTP protocol information.
  - **imap**—Trace IMAP protocol information.
  - **memory**—Trace memory usage.
  - **mime**—Trace MIME parser information.
  - **parser**— Trace protocol parser information.
  - **pfe**—Trace communication with PFE.
  - **pop3**—Trace POP3 protocol information.
  - **queue**—Trace queue information.
  - **regex-engine**—Trace Pattern Match Engine (PME) information.
  - **smtp**—Trace SMTP protocol information.
  - **sophos-anti-virus**—Trace anti-virus sophos engine information.
  - **tcp**—Trace TCP level information.

Copyright © 2016, Juniper Networks, Inc.

- **timer**—Trace timer processing.
- **utm-realtime**—Trace application-proxy realtime-thread information

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

traceoptions (Security Content Filtering)

Syntax
traceoptions flag *flag*;

Hierarchy Level
[edit security utm feature-profile content-filtering]

Release Information
Statement introduced in Junos OS Release 9.5 .

Description
Define tracing options for content filtering features.

Options
- *flag*:
  - **all**—Enable all content filtering trace flags.
  - **basic** —Trace content filtering basic information.
  - **detail**—Trace content filtering detailed information.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security UTM)

Syntax traceoptions flag *flag*;

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Define tracing operations for UTM features.

Options

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Enable trace for all UTM trace options.
  - **cli**—Trace CLI configuration activity and command changes.
  - **daemon**—Trace daemon information.
  - **ipc**—Trace communication events with Packet Forwarding Engine (PFE).
  - **pfe**—Trace PFE information.

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Web Filtering)

Syntax traceoptions flag *flag*;

Hierarchy Level [edit security utm feature-profile web-filtering]

Release Information Command introduced in Junos OS Release 10.1 .
Command introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Define tracing operations for individual Web filtering modules. To specify more than one tracing operation, include multiple flag statements.

Options

- *flag*:
  - **all**—Enable all Web filtering trace flags.
  - **basic** —Trace basic information on the Web filtering module.
  - **cache**—Enable Web filtering flags for the Web filtering cache maintained on the Web filtering module.
  - **enhanced**—Enable Web filtering flags for processing through Enhanced Web Filtering.
  - **heartbeat**—Trace connectivity information with Web filter server.
  - **ipc**—Trace Web filtering IPC messages.
  - **packet**—Trace packet information from session management.
  - **profile**—Trace profile configuration information.
  - **requests**—Trace requests sent to Web filter server.
  - **response**—Trace response received from Web filter server.
  - **session manager**—Trace session management information.
  - **socket**—Trace the communication socket with Web filter server.
  - **timer**—Trace aging information for requests sent to server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

 Copyright © 2016, Juniper Networks, Inc.

traceoptions (SMTP)

Syntax

```
traceoptions {
  flag {
    all;
    configuration;
    IPC;
    protocol-exchange;
    send-request;
  }
}
```

Hierarchy Level [edit smtp]

Release Information Statement added in Junos OS Release 10.0.

Description Set the Simple Mail Transfer Protocol (SMTP) traceoptions.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- SMTP Configuration Statement Hierarchy on page 5825

traffic-options

Syntax

```
traffic-options {
  sessions-per-client {
    limit value;
    over-limit (block | log-and-permit);
  }
}
```

Hierarchy Level [edit security utm utm-policy *policy-name*]

Release Information Statement introduced in Junos OS Release 9.5 .

Description In an attempt to consume all available resources and hinder the ability of the device, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose a session throttle.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

trickling

Syntax

```
trickling {
    timeout value;
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name*]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name*]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.5. Statement updated for Sophos support in Junos OS Release 11.1.

Description

HTTP trickling is a mechanism used to prevent the HTTP client or server from timing-out during a file transfer or during antivirus scanning. HTTP Trickling is time-based and there is only one parameter to configure for this feature, which is the timeout Interval. By default, trickling is disabled.

WARNING: When you enable the trickling option, keep in mind that trickling might send part of a file to the client during its antivirus scan. It is therefore possible that some of the content could be received by the client before the file has been fully scanned.

Options

value—Timeout interval in seconds.

Range: 0 through 600 seconds

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

type (Security Antivirus Feature Profile)

Syntax type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 . Statement updated for Sophos in Junos OS Release 11.1 .

Description Set the antivirus engine that will be used on the device. You can only have one engine type running and you must restart the device if you change engines.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

type (Security Content Filtering Notification Options)

Syntax type (message | protocol-only);

Hierarchy Level [edit security utm feature-profile content-filtering profile *profile-name* notification-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description When content is blocked, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code might be returned to the client.

Options
- **message**—Send a generic notification.
- **protocol-only**—Send a protocol-specific notification.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

type (Security Fallback Block)

Syntax type (message | protocol-only);

Hierarchy Level
[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options fallback-block]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options fallback-block]

Release Information Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description You can configure notifications for both fallback blocking and fallback nonblocking actions. With protocol-only notifications, a protocol-specific error code may be returned to the client.

Options
- message—Send a generic notification.
- protocol-only—Send a protocol-specific notification.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

type (Security Virus Detection)

Syntax type (message | protocol-only);

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options virus-detection]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options virus-detection]

Release Information Statement introduced in Junos OS Release 9.5. Support for Sophos engine added in Junos OS Release 11.1.

Description When content is blocked because a virus is found or a scan error occurs, the client generally still receives a successful response code but with modified content (file replacement) containing a warning message. But with protocol-only notifications, a protocol-specific error code might be returned to the client.

Options
- message—Send a generic notification.
- protocol-only—Send a protocol-specific notification.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

type (Security Web Filtering)

Syntax type (juniper-enhanced | juniper-local | surf-control-integrated | websense-redirect);

Hierarchy Level [edit security utm feature-profile web-filtering]

Release Information Command introduced in Junos OS Release 9.5 .
Command introduced in Junos OS Release 11.4 for Enhanced Web Filtering.

Description Define the type of Web filtering solution or URL filtering solution used by the device.

Options

- **juniper-enhanced**—Enable Enhanced Web Filtering on the device.
- **juniper-local** —Enable Juniper Networks local URL filtering on the device.
- **surf-control-integrated**—Enable integrated Web filtering on the device.
- **websense-redirect**—Redirect the URL to the Websense server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

upload-profile (Security Antivirus FTP)

Syntax upload-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* anti-virus ftp]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the antivirus FTP (upload) protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

upload-profile (Security Content Filtering FTP)

Syntax upload-profile *profile-name*;

Hierarchy Level [edit security utm utm-policy *policy-name* content-filtering ftp]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure a UTM policy for the content-filtering FTP (upload) protocol and attach this policy to a security profile to implement it.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

uri-check

Syntax uri-check;

Hierarchy Level [edit security utm feature-profile anti-virus sophos-engine profile *profile-name* scan-options]

Release Information Statement introduced in Junos OS Release 11.1 .

Description Perform Sophos antivirus Uniform Resource Identifier (URI) checking. URI checking is a way of analyzing URI content in HTTP traffic against a remote Sophos database to identify malware or malicious content. URI checking is on by default.

You can disable Sophos antivirus URI checking with the **no-uri-check** statement.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

url (Security Antivirus)

Syntax url *url*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update]
[edit security utm feature-profile anti-virus sophos-engine pattern-update]

Release Information Statement introduced in Junos OS Release 9.5 for Juniper Express engine and Kaspersky Lab engine. Support for Sophos engine added in Junos OS Release 11.1 .

Description Specify the URL for the pattern database. You should not change the default URL unless you are experiencing problems with it and have called for support.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

url-blacklist

Syntax url-blacklist *listname*;

Hierarchy Level [edit security utm feature-profile web-filtering]

Release Information Statement introduced in Junos OS Release 9.5.

Description This is a global blacklist category, blocking content for Web filtering.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

url-pattern

Syntax

```
url-pattern object-name {
  value [value];
}
```

Hierarchy Level [edit security utm custom-objects]

Release Information Statement introduced in Junos OS Release 9.5.

Description Use URL pattern lists to create custom URL category lists. These are lists of patterns that bypass scanning.

WARNING: Custom category does not take precedence over predefined categories when it has the same name as one of the predefined categories. We do not recommend having a custom category name be the same as the predefined category name.

Options

- ***object-name***—Name of the URL list object.
- **value *value***—Value of the URL list object. You can configure multiple values separated by spaces and enclosed in square brackets.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *Building Blocks Feature Guide for Security Devices*

url-whitelist (Security Antivirus)

Syntax url-whitelist *listname*;

Hierarchy Level [edit security utm feature-profile anti-virus]

Release Information Statement introduced in Junos OS Release 9.5 .

Description A URL whitelist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for scanning.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

url-whitelist (Security Web Filtering)

Syntax url-whitelist *listname*;

Hierarchy Level [edit security utm feature-profile web-filtering]

Description A URL whitelist is a unique custom list that you define in which all the URLs or IP addresses in that list for a specified category are always bypassed for filtering

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

username (Security Antivirus)

Syntax username *name*;

Hierarchy Level [edit security utm feature-profile anti-virus juniper-express-engine pattern-update proxy]
[edit security utm feature-profile anti-virus kaspersky-lab-engine pattern-update proxy]
[edit security utm feature-profile anti-virus sophos-engine pattern-update proxy]

Release Information Statement introduced in Junos OS Release 11.2 .

Description Set the username for the proxy server.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

utm

Syntax

```
utm {
  application-proxy {
    traceoptions {
      flag flag;
    }
  }
  custom-objects {
    custom-url-category object-name {
      value [value];
    }
    filename-extension object-name {
      value [value];
    }
    mime-pattern object-name {
      value [value];
    }
    protocol-command object-name {
      value [value];
    }
    url-pattern object-name {
      value [value];
    }
  }
  feature-profile {
    anti-spam {
      address-blacklist list-name;
      address-whitelist list-name;
      sbl {
        profile profile-name {
          custom-tag-string [string];
          (sbl-default-server | no-sbl-default-server);
          spam-action (block | tag-header | tag-subject);
        }
      }
      traceoptions {
        flag flag;
      }
    }
    anti-virus {
      juniper-express-engine {
        pattern-update {
          email-notify {
            admin-email email-address;
            custom-message message;
            custom-message-subject message-subject;
          }
          interval value;
          no-autoupdate;
          proxy {
            password password-string;
            port port-number;
            server address-or-url;
```

Copyright © 2016, Juniper Networks, Inc.

```
            username name;
        }
        url url;
    }
    profile profile-name {
        fallback-options {
            content-size (block | log-and-permit);
            default (block | log-and-permit);
            engine-not-ready (block | log-and-permit);
            out-of-resources (block | (log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
        }
        notification-options {
            fallback-block {
                administrator-email email-address;
                allow-email;
                custom-message message;
                custom-message-subject message-subject;
                display-host;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
            }
            fallback-non-block {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
                custom-message message;
                custom-message-subject message-subject;
                (notify-mail-sender | no-notify-mail-sender);
                type (message | protocol-only);
            }
        }
        scan-options {
            content-size-limit value;
            (intelligent-prescreening | no-intelligent-prescreening);
            timeout value;
        }
        trickling {
            timeout value;
        }
    }
}
kaspersky-lab-engine {
    pattern-update {
        email-notify {
            admin-email email-address;
            custom-message message;
            custom-message-subject message-subject;
        }
        interval value;
        no-autoupdate;
        proxy {
            password password-string;
```


Copyright © 2016, Juniper Networks, Inc.

```
            port port-number;
            server address-or-url;
            username name;
          }
          url url;
        }
        profile profile-name {
          fallback-options {
            content-size (block | log-and-permit);
            corrupt-file (block | log-and-permit);
            decompress-layer (block | log-and-permit);
            default (block | log-and-permit);
            engine-not-ready (block | log-and-permit);
            out-of-resources (block | (log-and-permit);
            password-file (block | (log-and-permit);
            timeout (block | log-and-permit);
            too-many-requests (block | log-and-permit);
          }
          notification-options {
            fallback-block {
              administrator-email email-address;
              allow-email;
              custom-message message;
              custom-message-subject message-subject;
              display-host;
              (notify-mail-sender | no-notify-mail-sender);
              type (message | protocol-only);
            }
            fallback-non-block {
              custom-message message;
              custom-message-subject message-subject;
              (notify-mail-recipient | no-notify-mail-recipient);
            }
            virus-detection {
              custom-message message;
              custom-message-subject message-subject;
              (notify-mail-sender | no-notify-mail-sender);
              type (message | protocol-only);
            }
          }
          scan-options {
            content-size-limit value;
            decompress-layer-limit value;
            (intelligent-prescreening | no-intelligent-prescreening);
            scan-extension filename;
            scan-mode (all | by-extension);
            timeout value;
          }
          trickling {
            timeout value;
          }
        }
      }
      mime-whitelist {
        exception listname;
        list listname {
```

```
      exception listname;
    }
  }
  sophos-engine {
    pattern-update {
      email-notify {
        admin-email email-address;
        custom-message message;
        custom-message-subject message-subject;
      }
      interval value;
      no-autoupdate;
      proxy {
        password password-string;
        port port-number;
        server address-or-url;
        username name;
      }
      url url;
    }
    profile <name> {
      fallback-options {
        content-size (block | log-and-permit | permit);
        default (block | log-and-permit | permit);
        engine-not-ready (block | log-and-permit | permit);
        out-of-resources (block | log-and-permit | permit);
        timeout (block | log-and-permit | permit);
        too-many-requests (block | log-and-permit | permit);
      }
      notification-options {
        fallback-block {
          administrator-email email-address;
          allow-email;
          custom-message message;
          custom-message-subject message-subject;
          display-host;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
        fallback-non-block {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-recipient | no-notify-mail-recipient);
        }
        virus-detection {
          custom-message message;
          custom-message-subject message-subject;
          (notify-mail-sender | no-notify-mail-sender);
          type (message | protocol-only);
        }
      }
      scan-options {
        content-size-limit value;
        (no-uri-check | uri-check);
        timeout value;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
        trickling {
          timeout value;
        }
      }
      sxl-retry value;
      sxl-timeout seconds;
    }
    traceoptions {
      flag flag;
    }
    type (juniper-express-engine | kaspersky-lab-engine | sophos-engine);
    url-whitelist listname;
  }
  content-filtering {
    profile profile-name {
      block-command protocol-command-list;
      block-content-type (activex | exe | http-cookie | java-applet | zip);
      block-extension extension-list;
      block-mime {
        exception list-name;
        list list-name;
      }
      notification-options {
        custom-message message;
        (notify-mail-sender | no-notify-mail-sender);
        type (message | protocol-only);
      }
      permit-command protocol-command-list;
    }
    traceoptions {
      flag flag;
    }
  }
  web-filtering {
    juniper-enhanced {
      cache {
        size value;
        timeout value;
      }
      profile profile-name {
        block-message {
          type {
            custom-redirect-url;
          }
          url url;
          }
        quarantine-message {
          type {
            custom-redirect-url;
          }
            url url;
            }
        category customurl-list name {
          action (block | log-and-permit | permit | quarantine);
        }
        custom-block-message value;
```

```
        custom-quarantine-message value;
        default (block | log-and-permit | permit | quarantine);
        fallback-settings {
          default (block | log-and-permit);
          server-connectivity (block | log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        no-safe-search;
        site-reputation-action {
          fairly-safe (block | log-and-permit | permit | quarantine);
          harmful (block | log-and-permit | permit | quarantine);
          moderately-safe (block | log-and-permit | permit | quarantine);
          suspicious (block | log-and-permit | permit | quarantine);
          very-safe (block | log-and-permit | permit |quarantine);
        }
        timeout value;
      }
      server {
        host host-name;
        port number;
      }
    }
    juniper-local {
      profile profile-name {
        custom-block-message value;
        default (block | log-and-permit | permit);
        fallback-settings {
          default (block | log-and-permit);
          server-connectivity (block | log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        timeout value;
      }
    }
    surf-control-integrated {
      cache {
        size value;
        timeout value;
      }
      profile profile-name {
        category customurl-list name {
          action (block | log-and-permit | permit);
        }
        custom-block-message value;
        default (block | log-and-permit | permit);
        fallback-settings {
          default (block | log-and-permit);
          server-connectivity (block | log-and-permit);
          timeout (block | log-and-permit);
          too-many-requests (block | log-and-permit);
        }
        timeout value;
      }
      server {
```


Copyright © 2016, Juniper Networks, Inc.

```
            host host-name;
            port number;
          }
        }
        traceoptions {
          flag flag;
        }
        type (juniper-enhanced | juniper-local | surf-control-integrated | websense-redirect);
        url-blacklist listname;
        url-whitelist listname;
        websense-redirect {
          profile profile-name {
            account value;
            custom-block-message value;
            fallback-settings {
              default (block | log-and-permit);
              server-connectivity (block | log-and-permit);
              timeout (block | log-and-permit);
              too-many-requests (block | log-and-permit);
            }
            server {
              host host-name;
              port number;
            }
            sockets value;
            timeout value;
          }
        }
      }
    }
    ipc {
      traceoptions flag flag;
    }
    traceoptions {
      flag flag;
    }
    utm-policy policy-name {
      anti-spam {
        smtp-profile profile-name;
      }
      anti-virus {
        ftp {
          download-profile profile-name;
          upload-profile profile-name;
        }
        http-profile profile-name;
        imap-profile profile-name;
        pop3-profile profile-name;
        smtp-profile profile-name;
      }
      content-filtering {
        ftp {
          download-profile profile-name;
          upload-profile profile-name;
        }
        http-profile profile-name;
```

```
        imap-profile profile-name;
        pop3-profile profile-name;
        smtp-profile profile-name;
      }
      traffic-options {
        sessions-per-client {
          limit value;
          over-limit (block | log-and-permit);
        }
      }
      web-filtering {
        http-profile profile-name;
      }
    }
  }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure UTM features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

utm-policy

Syntax

```
utm-policy policy-name {
  anti-spam {
    smtp-profile profile-name;
  }
  anti-virus {
    ftp {
      download-profile profile-name;
      upload-profile profile-name;
    }
    http-profile profile-name;
    imap-profile profile-name;
    pop3-profile profile-name;
    smtp-profile profile-name;
  }
  content-filtering {
    ftp {
      download-profile profile-name;
      upload-profile profile-name;
    }
    http-profile profile-name;
    imap-profile profile-name;
    pop3-profile profile-name;
    smtp-profile profile-name;
  }
  traffic-options {
    sessions-per-client {
      limit value;
      over-limit (block | log-and-permit);
    }
  }
  web-filtering {
    http-profile profile-name;
  }
}
```

Hierarchy Level [edit security utm]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a UTM policy for antivirus, antispam, content-filtering, traffic-options, and web-filtering protocols and attach this policy to a security profile to implement it.

Options ***policy-name***—Specify name of the UTM policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Policies Overview on page 1085

Copyright © 2016, Juniper Networks, Inc.

- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091
- Security Configuration Statement Hierarchy on page 603

utm-policy (Application Services)

Syntax utm-policy *policy-name*;

Hierarchy Level [edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name* then permit application-services]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure a UTM policy for application services and attach this policy to a security profile to implement it.

Options ***policy-name***—Specify the name of the UTM policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

virus-detection (Security Antivirus)

Syntax

```
virus-detection {
    custom-message message;
    custom-message-subject message-subject;
    (notify-mail-sender | no-notify-mail-sender);
    type (message | protocol-only);
}
```

Hierarchy Level

[edit security utm feature-profile anti-virus juniper-express-engine profile *profile-name* notification-options]
[edit security utm feature-profile anti-virus kaspersky-lab-engine profile *profile-name* notification-options]
[edit security utm feature-profile anti-virus sophos-engine profile *profile-name* notification-options]

Release Information

Statement introduced in Junos OS Release 9.5 . Support for Sophos engine added in Junos OS Release 11.1 .

Description

Configure a notification to send when a virus is detected.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- [edit security utm] Hierarchy Level on page 5830

Copyright © 2016, Juniper Networks, Inc.

web-filtering

Syntax

```
web-filtering {
  url-whitelist custwhitelist;
  url-blacklist custblacklist;
  http-reassemble;
  http-persist;
  type juniper-enhanced;
  juniper-enhanced {
    cache {
      timeout 1800;
      size 500;
    }
    server {
      host rp.cloud.threatseeker.com;
      port 80;
    }
    profile junos-wf-enhanced-default {
      category {
        Enhanced_Hacking {
          action log-and-permit;
          }
          Enhanced_Government {
          action quarantine;
        }
      }
      site-reputation-action {
        very-safe permit;
        moderately-safe log-and-permit;
        fairly-safe log-and-permit;
        harmful block;
        suspicious block;
      }
      default block;
      custom-block-message "***access denied ***";
      fallback-settings {
        default block;
        server-connectivity block;
        timeout block;
        too-many-requests block;
      }
      timeout 10;
      no-safe-search;
    }
    utm-policy mypolicy {
      web-filtering {
    http-profile my_ewfprofile01;
  }
  }
}
```

Hierarchy Level [edit security utm feature-profile]

Release Information Statement introduced in Junos OS Release 9.5 .

Copyright © 2016, Juniper Networks, Inc.

Description Configure UTM web-filtering features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Local Web Filtering on page 5797
- Monitoring Web Filtering Configurations on page 5815

websense-redirect

Syntax

```
websense-redirect {
  profile profile-name {
    account value;
    custom-block-message value;
    fallback-settings {
      default (block | log-and-permit);
      server-connectivity (block | log-and-permit);
      timeout (block | log-and-permit);
      too-many-requests (block | log-and-permit);
    }
    server {
      host host-name;
      port number;
    }
    sockets value;
    timeout value;
  }
}
```

Hierarchy Level [edit security utm feature-profile web-filtering]

Release Information Statement introduced in Junos OS Release 9.5 .

Description Configure the websense redirect engine features.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Example: Enhancing Security by Configuring Redirect Web Filtering Using Custom Objects on page 5806

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 251

Operational Commands

- clear security utm anti-spam statistics
- clear security utm antivirus statistics
- clear security utm content-filtering statistics
- clear security utm session
- clear security utm web-filtering statistics
- request security utm anti-virus juniper-express-engine
- request security utm anti-virus kaspersky-lab-engine
- request security utm anti-virus sophos-engine
- request system license update
- show configuration smtp
- show groups junos-defaults
- show security log
- show security policies
- show security utm anti-spam statistics
- show security utm anti-spam status
- show security utm anti-virus statistics
- show security utm anti-virus status
- show security utm content-filtering statistics
- show security utm session
- show security utm status
- show security utm web-filtering statistics
- show security utm web-filtering status

Copyright © 2016, Juniper Networks, Inc.

clear security utm anti-spam statistics

Syntax clear security utm anti-spam statistics

Release Information Command introduced in Junos OS Release 9.5 . Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Clear antispam statistics information. With chassis cluster support for UTM, statistics from both the nodes is cleared.

Required Privilege Level clear

Related Documentation

- show security utm anti-spam statistics on page 6024
- show security utm anti-spam status on page 6025

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security utm antivirus statistics

Syntax clear security utm anti-virus statistics

Release Information Command introduced in Junos OS Release 9.5. Support for Sophos Antivirus added in Junos OS Release 11.1.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Clear antivirus statistics information. With chassis cluster support for UTM, statistics from both the nodes are cleared.

Required Privilege Level clear

Related Documentation

- show security utm anti-virus statistics on page 6026
- show security utm anti-virus status on page 6028
- request security utm anti-virus juniper-express-engine on page 6007
- request security utm anti-virus kaspersky-lab-engine on page 6008

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security utm content-filtering statistics

Syntax clear security utm content-filtering statistics

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Clear content-filtering statistics information. With chassis cluster support for UTM, statistics from both the nodes are cleared.

Required Privilege Level clear

Related Documentation
- show security utm content-filtering statistics on page 6030

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security utm session

Syntax clear security utm session

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Clear UTM session information. With chassis cluster support for UTM, sessions on both the nodes are cleared.

Required Privilege Level clear

Related Documentation
- show security utm session on page 6031
- show security utm status on page 6032

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security utm web-filtering statistics

Syntax clear security utm web-filtering statistics

Release Information Command introduced in Junos OS Release 9.5 .
Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Clear web filtering statistics information. With chassis cluster support for UTM, statistics from both the nodes is cleared.

Required Privilege Level clear

Related Documentation
- show security utm web-filtering statistics on page 6033
- show security utm web-filtering status on page 6036

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

request security utm anti-virus juniper-express-engine

Syntax request security utm anti-virus juniper-express-engine

Release Information Command introduced in Junos OS Release 9.5 .
Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Manually update the express antivirus pattern database using the command described. You can update the express antivirus pattern database automatically or manually. With full chassis cluster support for UTM this command is operational on both the nodes.

Options

- **pattern-delete** — Delete the current express antivirus pattern database.
- **pattern-reload** — Reload the express antivirus pattern database.
- **pattern-update** — Update the express antivirus pattern database with the latest signatures.

Required Privilege Level maintenance

Related Documentation

- request security utm anti-virus kaspersky-lab-engine on page 6008
- clear security utm antivirus statistics on page 6003
- show security utm anti-virus statistics on page 6026
- show security utm anti-virus status on page 6028

List of Sample Output request security utm anti-virus juniper-express-engine pattern-update on page 6007

Output Fields request security utm anti-virus juniper-express-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus juniper-express-engine pattern-update

```
user@host> request security utm anti-virus juniper-express-engine pattern-update
```

request security utm anti-virus kaspersky-lab-engine

Syntax request security utm anti-virus kaspersky-lab-engine

Release Information Command introduced in Junos OS Release 11.1 .
Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Manually update the full file-based antivirus pattern database using the commands described. You can update the full file-based antivirus pattern database automatically or manually. With full chassis cluster support for UTM this command is operational on both the nodes.

Options

- **pattern-delete** — Delete the current full file-based antivirus pattern database.
- **pattern-reload** — Reload the full file-based antivirus pattern database.
- **pattern-update** — Update the full file-based antivirus pattern database with the latest signatures.

Required Privilege Level maintenance

Related Documentation

- request security utm anti-virus juniper-express-engine on page 6007
- clear security utm antivirus statistics on page 6003
- show security utm anti-virus statistics on page 6026
- show security utm anti-virus status on page 6028

List of Sample Output request security utm anti-virus kaspersky-lab-engine pattern-update on page 6008

Output Fields request security utm anti-virus kaspersky-lab-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus kaspersky-lab-engine pattern-update

```
user@host> request security anti-virus kaspersky-lab-engine pattern-update
```


Copyright © 2016, Juniper Networks, Inc.

request security utm anti-virus sophos-engine

Syntax request security utm anti-virus sophos-engine

Release Information Command introduced in Junos OS Release 11.1 .
Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Manually update the Sophos antivirus pattern database using the command described. You can update the express antivirus pattern database automatically or manually. To update automatically you use the configuration statement **set security utm feature-profile anti-virus sophos-engine pattern-update interval *seconds***. With full chassis cluster support for UTM this command is operational on both the nodes.

Options
- **pattern-delete** — Delete the current Sophos antivirus pattern database.
- **pattern-reload** — Reload the Sophos antivirus pattern database.
- **pattern-update** — Update the Sophos antivirus pattern database with the latest signatures.

Required Privilege Level maintenance

Related Documentation
- clear security utm antivirus statistics on page 6003
- show security utm anti-virus statistics on page 6026
- show security utm anti-virus status on page 6028

List of Sample Output request security utm anti-virus sophos-engine pattern-update on page 6009

Output Fields request security utm anti-virus sophos-engine pattern-update

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security utm anti-virus sophos-engine pattern-update

```
user@host> request security utm anti-virus sophos-engine pattern-update
```

Copyright © 2016, Juniper Networks, Inc.

request system license update

Syntax request system license update

Release Information Command introduced in Junos OS Release 9.5.

Description Start autoupdating license keys from the LMS server.

Options **trial**—Starts autoupdating trial license keys from the LMS server.

Required Privilege Level maintenance

Related Documentation
- *show system license (View)*

List of Sample Output
request system license update on page 6010
request system license update trial on page 6010

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license update

```
user@host> request system license update

Request to automatically update license keys from https://ae1.juniper.net has
been sent, use show system license to check status.
```

request system license update trial

```
user@host> request system license update trial

Request to automatically update trial license keys from https://ae1.juniper.net
has been sent, use show system license to check status.
```


Copyright © 2016, Juniper Networks, Inc.

show configuration smtp

Syntax show configuration smtp

Release Information Command introduced in Junos OS Release 10.0 .

Description Display complete SMTP information.

Options
- apply-groups—Groups from which SMTP inherits configuration data.
- apply-groups-except—Groups from which SMTP restricts inheriting configuration data.

Required Privilege Level view

Related Documentation
- SMTP Configuration Statement Hierarchy on page 5825

List of Sample Output show configuration smtp on page 6011

Output Fields Table 497 describes the output fields for the **show configuration smtp** command.

Table 497: show configuration smtp

Field Name	Field Description	Level of Output
address	SMTP server's IPv4 address	All levels
login	Configure a mail sender account to the server	All levels
password	Default sender password for user authentication	All levels

Sample Output

show configuration smtp

```
user@host> show configuration smtp
primary-server {
    address 218.102.48.213;
    login "dayone@example.com" {
        password "$ABC123"; ## SECRET-DATA
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

show groups junos-defaults

Syntax show groups junos-defaults

Release Information Command introduced before Junos OS Release 7.4.

Description Display the full set of available preset statements from the Junos OS defaults group.

```
user@host# show groups junos-defaults
  groups {
    junos-defaults {
      applications {
        # File Transfer Protocol
        application junos-ftp {
          application-protocol ftp;
          protocol tcp;
          destination-port 21;
        }
        # Trivial File Transfer Protocol
        application junos-tftp {
          application-protocol tftp;
          protocol udp;
          destination-port 69;
        }
        # RPC port mapper on TCP
        application junos-rpc-portmap-tcp {
          application-protocol rpc-portmap;
          protocol tcp;
          destination-port 111;
        }
        # RPC port mapper on UDP
      }
    }
  }
```

Required Privilege Level view

Related Documentation

- *Using Junos OS Defaults Groups*

Copyright © 2016, Juniper Networks, Inc.

show security log

Syntax show security log {*all*| *destination-address*| *destination-port*| *event-id*| *failure*|*interface-name*| *newer-than*| *older-than*| *process*| *protocol*| *severity*| *sort-by*| *source-address*| *source-port*| *success*| *user*}

Release Information Command introduced in Junos OS Release 11.2 .

Description Display security event logs. This command continuously displays security events on the screen. To stop the display, press Ctrl+c.

Options

all—Displays all audit event logs stored in the device memory.

destination-address—Displays audit event logs with the specified destination address.

destination-port—Displays audit event logs with the specified destination port.

event-id—Displays audit event logs with the specified event identification number.

failure—Displays failed audit event logs.

interface-name—Displays audit event logs with the specified interface.

newer-than—Displays audit event logs newer than the specified date and time.

older-than—Displays audit event logs older than the specified date and time.

process—Displays audit event logs with the specified process that generated the event.

protocol—Displays audit event logs generated through the specified protocol.

severity—Displays audit event logs generated with the specified severity.

sort-by—Displays audit event logs generated sorted with the specified options.

source-address—Displays audit event logs with the specified source address.

source-port—Displays audit event logs with the specified source port.

success—Displays successful audit event logs.

username—Displays audit event logs generated for the specified user.

Required Privilege Level view

Related Documentation

- *exclude (Security Log)*
- *clear security log*

List of Sample Output show security log on page 6014

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 498 lists the output fields for the **show security log** command. Output fields are listed in the approximate order in which they appear.

Table 498: show security log Output Fields

Field Name	Field Description
Event time	The timestamp of the events received. On SRX Series devices, security logs were always timestamped using the UTC time zone by running **set system time-zone utc** and **set security log utc-timestamp** CLI commands. Now, time zone can be defined using the local time zone by running the **set system time-zone time-zone** command to specify the local time zone that the system should use when timestamping the security logs.
Message	Security events are listed.

Sample Output

show security log

```
user@host> show security log
Event time               Message
2010-10-22 13:28:37 CST  session created 1.1.1.2/1->2.2.2.2/1308 icmp
1.1.1.2/1->2.2.2.2/1308 None None 1 policy1 trustZone untrustZone 52 N/A(N/A)
ge-0/0/1.0
2010-10-22 13:28:38 CST  session created 1.1.1.2/2->2.2.2.2/1308 icmp
1.1.1.2/2->2.2.2.2/1308 None None 1 policy1 trustZone untrustZone 54 N/A(N/A)
ge-0/0/1.0

...

2010-10-22 13:36:12 CST  session denied 1.1.1.2/1->2.2.2.2/54812 icmp 1(8) policy1
 trustZone untrustZone N/A(N/A) ge-0/0/1.0
2010-10-22 13:36:14 CST  session denied 1.1.1.2/2->2.2.2.2/54812 icmp 1(8) policy1
 trustZone untrustZone N/A(N/A) ge-0/0/1.0

...

2010-10-27 15:50:11 CST  IP spoofing! source: 2.2.2.20, destination: 2.2.2.2,
protocol-id: 17, zone name: trustZone, interface name: ge-0/0/1.0, action: drop
2010-10-27 15:50:11 CST  IP spoofing! source: 2.2.2.20, destination: 2.2.2.2,
protocol-id: 17, zone name: trustZone, interface name: ge-0/0/1.0, action: drop

...

2011-02-18 15:53:34 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/common/certification-authority/ca-profile1-ca1.cert
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/common/crl/ca-profile1.crl
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/system-key-pair/system-generated.priv
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/system-cert/system-generated.cert
2011-02-18 15:53:35 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/common/key-pair/cert1.priv
2011-02-18 15:53:42 CST  PKID_PV_OBJECT_READ: A PKI object was read into memory
from /var/db/certs/common/key-pair/test2.priv

...
```


Copyright © 2016, Juniper Networks, Inc.

```
2011-03-14 23:00:40 PDT  IDP_COMMIT_COMPLETED: IDP policy commit is complete.
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]

 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]

 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]
                         IDP_POLICY_LOAD_FAILED: IDP policy loading failed ;poli
cy[/var/db/idpd/bins/.bin.gz.v], detector[/usr/libdata/libidp-detector.so.tgz.v]

 ,failure detail[Policy loading failed :: Policy file not found
2011-03-14 23:00:58 PDT  ]

...

Event time               Message
2011-03-21 14:21:49 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp
 9.9.9.1 source-address 6.6.6.1 '
2011-03-21 14:23:01 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp
 9.9.9.1 source-address 6.6.6.1 '
2011-03-21 14:23:05 CST  KMD_PM_SA_ESTABLISHED: Local gateway: 7.7.7.1, Remote
gateway: 8.8.8.1, Local ID: ipv4(any:0,[0..3]=6.6.6.1), Remote ID:
ipv4(any:0,[0..3]=9.9.9.1), Direction: inbound, SPI: 37a2a179, AUX-SPI: 0, Mode:
 tunnel, Type: dynamic
2011-03-21 14:23:05 CST  KMD_PM_SA_ESTABLISHED: Local gateway: 7.7.7.1, Remote
gateway: 8.8.8.1, Local ID: ipv4(any:0,[0..3]=6.6.6.1), Remote ID:
ipv4(any:0,[0..3]=9.9.9.1), Direction: outbound, SPI: b2231c1f, AUX-SPI: 0, Mode:
 tunnel, Type: dynamic
2011-03-21 14:23:08 CST  UI_CMDLINE_READ_LINE: User 'root', command 'set date ntp
 9.9.9.1 source-address 6.6.6.1 '
2011-03-21 14:23:13 CST  UI_CMDLINE_READ_LINE: User 'root', command 'show security
 log '
```

Copyright © 2016, Juniper Networks, Inc.

show security policies

Syntax

show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information

Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20.

Description

Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options

- **none**—Display basic information about all configured policies.
- **detail**—(Optional) Display a detailed view of all of the policies configured on the device.
- **policy-name *policy-name***—(Optional) Display information about the specified policy.
- **global**—Display information about global policies.

Required Privilege Level

view

Related Documentation

- Security Policies Overview on page 1085
- Understanding Security Policy Rules on page 1087
- Understanding Security Policy Elements on page 1091

List of Sample Output

show security policies on page 6019
show security policies policy-name p1 detail on page 6020
show security policies (services-offload) on page 6021
show security policies detail on page 6021
show security policies detail (TCP Options) on page 6022
show security policies policy-name p1 (Negated Address) on page 6022
show security policies policy-name p1 detail (Negated Address) on page 6023
show security policies global on page 6023

Copyright © 2016, Juniper Networks, Inc.

Output Fields Table 51 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 499: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Copyright © 2016, Juniper Networks, Inc.

Table 499: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.
Application Firewall	An application firewall includes the following: • **Rule-set**—Name of the rule set. • **Rule**—Name of the rule. ◦ **Dynamic applications**—Name of the applications. ◦ **Dynamic application groups**—Name of the application groups. ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following: ▪ **permit** ▪ **deny** • **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.
Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following: ◦ **permit** ◦ **firewall-authentication** ◦ **tunnel ipsec-vpn** *vpn-name* ◦ **pair-policy** *pair-policy-name* ◦ **source-nat pool** *pool-name* ◦ **pool-set** *pool-set-name* ◦ **interface** ◦ **destination-nat** *name* ◦ **deny** ◦ **reject** ◦ **services-offload**
Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Copyright © 2016, Juniper Networks, Inc.

Table 499: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```

Copyright © 2016, Juniper Networks, Inc.

```
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, application services, log, scheduled
    Application firewall : my_ruleset1
  Policy: p2, State: enabled, Index: 5, Sequence number: 2
    Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::22/32
    Destination addresses:
    da-1-ipv4:  198.51.100.12/24
    da-2-ipv6:  2001:db8:a0b:12f0::1/32
    da-3-ipv6:  2001:db8:a0b:12f0::9/32
    Source identities: role1, role4
    Applications: any
    Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 198.51.100.12/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Destination Address Translation: drop translated
  Application firewall :
  Rule-set: my_ruleset1
    Rule: rule1
      Dynamic Applications: junos:FACEBOOK, junos:YSMG
      Dynamic Application groups: junos:web, junos:chat
      Action: deny
    Default rule: permit
  Session log: at-create, at-close
  Scheduler name: sch20
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :               18144                 545 bps
      Initial direction:                9072                272 bps
      Reply direction  :                9072                 272 bps
    Output bytes       :               18144                 545 bps
      Initial direction:                9072                 272 bps
```


Copyright © 2016, Juniper Networks, Inc.

```
    Reply direction  :                9072                  272 bps
  Input  packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Output packets     :                 216                    6 pps
    Initial direction:                 108                    3 bps
    Reply direction  :                 108                    3 bps
  Session rate       :                 108                    3 sps
  Active sessions    :                  93
  Session deletions  :                  15
  Policy lookups     :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Policy statistics:
    Input  bytes       :                18144                  545 bps
      Initial direction:                 9072                 272 bps
      Reply direction  :                 9072                  272 bps
    Output bytes       :                18144                  545 bps
```

```
      Initial direction:                9072                  272 bps
      Reply direction  :                9072                  272 bps
    Input  packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Output packets     :                 216                    6 pps
      Initial direction:                 108                    3 bps
      Reply direction  :                 108                    3 bps
    Session rate       :                 108                    3 sps
    Active sessions    :                  93
    Session deletions  :                  15
    Policy lookups     :                 108
Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
 5, Scope Policy: 0
  Policy Type: Configured
  Description: The policy p2 is for the sales team
  Sequence number: 1
  From zone: untrust, To zone: trust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
 Source identities:
  role1
  role2
  role4
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global): 0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
  Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
user@host>show security policies policy-name p1
node0:
--------------------------------------------------------------------------
```


Copyright © 2016, Juniper Networks, Inc.

```
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses(excluded): as1
    Destination addresses(excluded): as2
    Applications: any
    Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
user@host>show security policies policy-name p1 detail
node0:
--------------------------------------------------------------------------
Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses(excluded):
    ad1(ad): 255.255.255.255/32
    ad2(ad):  1.1.1.1/32
    ad3(ad):  15.100.199.56 ~ 15.200.100.16
    ad4(ad):  15.100.196.0/22
    ad5(ad):  15.1.7.199 ~ 15.1.8.19
    ad6(ad):  15.1.8.0/21
    ad7(ad):  15.1.7.0/24
  Destination addresses(excluded):
    ad13(ad2): 20.1.7.0/24
    ad12(ad2):  20.1.4.1/32
    ad11(ad2):  20.1.7.199 ~ 20.1.8.19
    ad10(ad2):  50.1.4.0/22
    ad9(ad2):  20.1.1.11 ~ 50.1.5.199
    ad8(ad2):  2.1.1.1/32
  Application: any
    IP protocol: 0, ALG: 0, Inactivity timeout: 0
      Source port range: [0-0]
      Destination port range: [0-0]
  Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
user@host>show security policies global policy-name Pa
node0:
--------------------------------------------------------------------------
  Global policies:
  Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    From zones: zone1, zone2
    To zones: zone3, zone4
    Source addresses: any
    Destination addresses: any
    Applications: any
    Action: permit
```

Copyright © 2016, Juniper Networks, Inc.

show security utm anti-spam statistics

Syntax show security utm anti-spam statistics

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Display antispam statistics for connections including total e-mail scanned, tagged, and dropped connections.

Statistics from both the nodes (with full chassis cluster support for UTM) are displayed.

Required Privilege Level view

Related Documentation
- clear security utm anti-spam statistics on page 6002
- show security utm anti-spam status on page 6025

Output Fields show security utm anti-spam statistics

Output fields are listed in the approximate order in which they appear.

show security utm anti-spam statistics

```
user@host> show security utm anti-spam statistics
Total connections:     0
Denied connections:    0
Total greetings:       0
Denied greetings:      0
Total e-mail scanned:  0
White list hit:        0
Black list hit:        0
Spam total:            0
Spam tagged:           0
Spam dropped:          0
DNS errors:            0
Timeout errors:        0
Return errors:         0
Invalid parameter errors: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security utm anti-spam status

Syntax show security utm anti-spam status

Release Information Command introduced in Junos OS Release 9.5 .
Support for UTM in chassis cluster added in Junos OS Release 11.4 .

Description Display antispam status for connections including whitelist and blacklist server information. Status of both the nodes (with full chassis cluster support for UTM) is displayed.

Required Privilege Level view

Related Documentation
- clear security utm anti-spam statistics on page 6002
- show security utm anti-spam statistics on page 6024

Output Fields show security utm anti-spam status

Output fields are listed in the approximate order in which they appear.

show security utm anti-spam status

```
user@host> show security utm anti-spam status
SBL Whitelist Server:
SBL Blacklist Server:
    msgsecurity.example.net

DNS Server:
    Primary  :     1.2.3.4, Src Interface: ge-0/0/0
    Secondary:     0.0.0.0, Src Interface: ge-0/0/1
    Ternary  :     0.0.0.0, Src Interface: fe-0/0/2
```

Copyright © 2016, Juniper Networks, Inc.

show security utm anti-virus statistics

Syntax show security utm anti-virus statistics <fpc <fpc-slot *fpc-slot* pic-slot *pic-slot*>>

Release Information Command introduced in Junos OS Release 9.5. Support for Sophos Antivirus added in Junos OS Release 11.1.
Support for UTM in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Description Display antivirus statistics for connections including clean and infected files, scan engine status, and aggregated statistics from all FPCs and PICs. Statistics from both the nodes (with full chassis cluster support for UTM) are displayed.

Required Privilege Level view

Related Documentation
- clear security utm antivirus statistics on page 6003
- show security utm anti-virus status on page 6028
- request security utm anti-virus juniper-express-engine on page 6007
- request security utm anti-virus kaspersky-lab-engine on page 6008

List of Sample Output
show security utm anti-virus statistics on page 6026
show security utm anti-virus statistics fpc on page 6027
show security utm anti-virus statistics fpc fpc-slot 5 pic-slot 0 on page 6027

Output Fields show security utm anti-virus statistics

Output fields are listed in the approximate order in which they appear.

Sample Output

show security utm anti-virus statistics

```
user@host>show security utm anti-virus statistics
 UTM Anti Virus statistics:
 MIME-whitelist passed:                0
 URL-whitelist passed:                 0
 Scan Request:

  Total           Clean         Threat-found    Fallback
      0               0                0              0

 Fallback:
                                Log-and-Permit    Block             Permit
 Engine not ready:                  0               0                 0
 Out of resources:                  0               0                 0
 Timeout:                           0               0                 0
 Maximum content size:              0               0                 0
 Too many requests:                 0               0                 0
 Others:                            0               0                 0
```


Copyright © 2016, Juniper Networks, Inc.

show security utm anti-virus statistics fpc

```
user@host>show security utm anti-virus statistics fpc

 fpc-slot 5 pic-slot 0
 UTM Anti Virus statistics:
 MIME-whitelist passed:               0
 URL-whitelist passed:                0
 Scan Request:

  Total           Clean         Threat-found    Fallback
      0               0              0               0

 Fallback:
                                Log-and-Permit    Block             Permit
 Engine not ready:                  0                0                 0
 Out of resources:                  0                0                 0
 Timeout:                           0                0                 0
 Maximum content size:              0                0                 0
 Too many requests:                 0                0                 0
 Others:                            0                0                 0
```

show security utm anti-virus statistics fpc fpc-slot 5 pic-slot 0

```
user@host>show security utm anti-virus statistics fpc fpc-slot 5 pic-slot 0
 UTM Anti Virus statistics:
 MIME-whitelist passed:               0
 URL-whitelist passed:                0
 Scan Request:

  Total           Clean         Threat-found    Fallback
      0               0              0               0

 Fallback:
                                Log-and-Permit    Block             Permit
 Engine not ready:                  0                0                 0
 Out of resources:                  0                0                 0
 Timeout:                           0                0                 0
 Maximum content size:              0                0                 0
 Too many requests:                 0                0                 0
 Others:                            0                0                 0
```

Copyright © 2016, Juniper Networks, Inc.

show security utm anti-virus status

Syntax show security utm anti-virus status <fpc <fpc-slot *fpc-slot* pic-slot *pic-slot*>>

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Description Display antivirus status for connections including clean and infected files, scan engine status, and aggregated status from all FPCs and PICs. Status of both the nodes (with full chassis cluster support for UTM) is displayed.

Required Privilege Level view

Related Documentation
- clear security utm antivirus statistics on page 6003
- show security utm anti-virus statistics on page 6026
- request security utm anti-virus juniper-express-engine on page 6007
- request security utm anti-virus kaspersky-lab-engine on page 6008

List of Sample Output
show security utm anti-virus status on page 6028
show security utm anti-virus status fpc on page 6028
show security utm anti-virus status fpc fpc-slot 5 pic-slot 0 on page 6029

Output Fields show security utm anti-virus status

Output fields are listed in the approximate order in which they appear.

Sample Output

show security utm anti-virus status

```
user@host> show security utm anti-virus status
 UTM anti-virus status:

    Anti-virus key expire date: 2014-11-01 08:00:00
    Update server: http://update.juniper-updates.net/AV/device-name/
    Interval: 60 minutes
    Pattern update status: in process
    Last result: downloading signature files
    Anti-virus signature version: 12/18/2013 03:12 GMT, virus records: 574705
    Anti-virus signature compiler version: N/A
    Scan engine type: kaspersky-lab-engine
    Scan engine information: engine is not ready
```

show security utm anti-virus status fpc

```
user@host> show security utm anti-virus status fpc
 fpc-slot 5 pic-slot 0
 UTM anti-virus status:

    Anti-virus key expire date: license not installed
    Update server: http://update.juniper-updates.net/SAV/
```


Copyright © 2016, Juniper Networks, Inc.

```
        Interval: 1440 minutes
        Pattern update status: update disabled due to no license
        Last result: already have latest database
    Anti-virus signature version:  000000_00
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
```

show security utm anti-virus status fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm anti-virus status fpc fpc-slot 5 pic-slot 0
 UTM anti-virus status:

    Anti-virus key expire date: license not installed
    Update server: http://update.juniper-updates.net/SAV/
        Interval: 1440 minutes
        Pattern update status: update disabled due to no license
        Last result: already have latest database
    Anti-virus signature version: 000000_00
    Scan engine type: sophos-engine
    Scan engine information: last action result: No error
```

Copyright © 2016, Juniper Networks, Inc.

show security utm content-filtering statistics

Syntax show security utm content-filtering statistics

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Display content-filtering statistics for connections including lists of blocked files and the reasons for blocking. Statistics from both the nodes (with full chassis cluster support for UTM) are displayed.

Required Privilege Level view

Related Documentation
- clear security utm content-filtering statistics on page 6004

Output Fields show security utm content-filtering statistics

Output fields are listed in the approximate order in which they appear.

show security utm content-filtering statistics

```
user@host> show security utm content-filtering statistics
Content-filtering-statistic:         Blocked
     Base on command list:                   0
     Base on mime list:                      0
     Base on extension list:                 0
     ActiveX plugin:                         0
     Java applet:                            0
     EXE files:                              0
     ZIP files:                              0
     HTTP cookie:                            0
```


Copyright © 2016, Juniper Networks, Inc.

show security utm session

Syntax show security utm session

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Display general UTM session information including all allocated sessions and active sessions. Also, display information from both nodes in a chassis cluster.

Required Privilege Level view

Related Documentation
- clear security utm session on page 6005
- show security utm status on page 6032

Output Fields show security utm session

When you enter this command, you are provided feedback on the status of your request.

show security utm session

```
user@host> show security utm session
Maximum sessions:           4000
 Total allocated sessions:         0
 Total freed sessions:             0
 Active sessions:                  0
```

show security utm status

Syntax show security utm status

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4.

Description Display whether the UTM service is running or not and status of both the nodes (with full chassis cluster support for UTM).

Required Privilege Level view

Related Documentation
- clear security utm session on page 6005
- show security utm session on page 6031

Output Fields show security utm status

When you enter this command, you are provided feedback on the status of your request.

show security utm status

```
user@host> show security utm status
UTM service status: Running
```


Copyright © 2016, Juniper Networks, Inc.

show security utm web-filtering statistics

Syntax show security utm web-filtering statistics <fpc <fpc-slot *fpc-slot* pic-slot *pic-slot*>>

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC statistics added in Junos OS Release 12.1X46-D10.

Description Display Web filtering statistics for connections including whitelist and blacklist hits and custom category hits. The aggregated statistics from all FPCs and PICs and statistics from both the nodes (with full chassis cluster support for UTM) are also displayed.

Required Privilege Level view

Related Documentation
- clear security utm web-filtering statistics on page 6006
- show security utm web-filtering status on page 6036

List of Sample Output
show security utm web-filtering statistics on page 6033
show security utm web-filtering statistics fpc on page 6034
show security utm web-filtering statistics fpc fpc-slot 5 pic-slot 0 on page 6034

Output Fields show security utm web-filtering statistics

Output fields are listed in the approximate order in which they appear.

Sample Output

show security utm web-filtering statistics

```
user@host> show security utm web-filtering statistics
 UTM web-filtering statistics:
    Total requests:                    0
    white list hit:                    0
    Black list hit:                    0
    Queries to server:                 0
    Server reply permit:               0
    Server reply block:                0
    Server reply quarantine:           0
    Server reply quarantine block:     0
    Server reply quarantine permit:    0
    Custom category permit:            0
    Custom category block:             0
    Custom category quarantine:        0
    Custom category qurantine block:   0
    Custom category quarantine permit: 0
    Site reputation permit:            0
    Site reputation block:             0
    Site reputation quarantine:        0
    Site reputation quarantine block:  0
    Site reputation quarantine permit: 0
    Site reputation by Category        0
    Site reputation by Global          0
    Cache hit permit:                  0
```

Copyright © 2016, Juniper Networks, Inc.

```
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
    Safe-search redirect:               0
    Web-filtering sessions in total:    128000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit          block
          Default                                0              0
          Timeout                                0              0
     Connectivity                                0              0
Too-many-requests                                0              0
```

show security utm web-filtering statistics fpc

```
user@host> show security utm web-filtering statistics fpc

 fpc-slot 5 pic-slot 0
 UTM web-filtering statistics:
    Total requests:                     0
    white list hit:                     0
    Black list hit:                     0
    Queries to server:                  0
    Server reply permit:                0
    Server reply block:                 0
    Server reply quarantine:            0
    Server reply quarantine block:      0
    Server reply quarantine permit:     0
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
    Cache hit permit:                   0
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
    Safe-search redirect:               0
    Web-filtering sessions in total:    128000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit          block
          Default                                0              0
          Timeout                                0              0
     Connectivity                                0              0
Too-many-requests                                0              0
```

show security utm web-filtering statistics fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm web-filtering statistics fpc fpc-slot 5 pic-slot 0
 UTM web-filtering statistics:
    Total requests:                     0
    white list hit:                     0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Black list hit:                     0
    Queries to server:                  0
    Server reply permit:                0
    Server reply block:                 0
    Server reply quarantine:            0
    Server reply quarantine block:      0
    Server reply quarantine permit:     0
    Custom category permit:             0
    Custom category block:              0
    Custom category quarantine:         0
    Custom category qurantine block:    0
    Custom category quarantine permit:  0
    Site reputation permit:             0
    Site reputation block:              0
    Site reputation quarantine:         0
    Site reputation quarantine block:   0
    Site reputation quarantine permit:  0
    Site reputation by Category         0
    Site reputation by Global           0
    Cache hit permit:                   0
    Cache hit block:                    0
    Cache hit quarantine:               0
    Cache hit quarantine block:         0
    Cache hit quarantine permit:        0
    Safe-search redirect:               0
    Web-filtering sessions in total:    128000
    Web-filtering sessions in use:      0
    Fallback:                       log-and-permit           block
          Default                                0               0
          Timeout                                0               0
     Connectivity                                0               0
Too-many-requests                                0               0
```

show security utm web-filtering status

Syntax show security utm web-filtering status <fpc <fpc-slot *fpc-slot* pic-slot *pic-slot*>>

Release Information Command introduced in Junos OS Release 9.5.
Support for UTM in chassis cluster added in Junos OS Release 11.4. Support for Flexible PIC Concentrator (FPC) and PIC status added in Junos OS Release 12.1X46-D10.

Description Display whether the Web filtering server connection is up or not. The aggregated status from all FPCs and PICs and status of both the nodes (with full chassis cluster support for UTM) are also displayed.

Required Privilege Level view

Related Documentation
- clear security utm web-filtering statistics on page 6006
- show security utm web-filtering statistics on page 6033

List of Sample Output
show security utm web-filtering status on page 6036
show security utm web-filtering status fpc on page 6036
show security utm web-filtering status fpc fpc-slot 5 pic-slot 0 on page 6036

Output Fields show security utm web-filtering status

Output fields are listed in the approximate order in which they appear.

Sample Output

show security utm web-filtering status

```
user@host> show security utm web-filtering status
 UTM web-filtering status:
    Server status: Juniper Enhanced using Websense server UP
```

show security utm web-filtering status fpc

```
user@host> show security utm web-filtering status fpc
 UTM web-filtering status fpc:
 fpc-slot 5 pic-slot 0
 Connectivity status: UP
 fpc-slot 0 pic-slot 1
 Connectivity status: UP
```

show security utm web-filtering status fpc fpc-slot 5 pic-slot 0

```
user@host> show security utm web-filtering status fpc fpc-slot 5 pic-slot 0
 UTM web-filtering status:
 Connectivity status: UP
```


Copyright © 2016, Juniper Networks, Inc.

VPN Feature Guide for Security Devices

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 79

Overview

- Introduction to IPsec VPNs on page 6041
- Understanding VPN Tunnel Management on page 6067

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 252

Introduction to IPsec VPNs

- IPsec VPN Overview on page 6041
- Understanding IKE and IPsec Packet Processing on page 6048
- Understanding Phase 1 of IKE Tunnel Negotiation on page 6055
- Understanding Phase 2 of IKE Tunnel Negotiation on page 6057
- IPsec VPN with Autokey IKE Configuration Overview on page 6059
- IPsec VPN with Manual Keys Configuration Overview on page 6060
- Recommended Configuration Options for Site-to-Site VPN with Static IP Addresses on page 6060
- Recommended Configuration Options for Site-to-Site or Dialup VPNs with Dynamic IP Addresses on page 6061
- Configuring Remote IKE IDs for Site-to-Site VPNs on page 6062
- Configuring IPsec VPN Using the VPN Wizard on page 6063
- Understanding Suite B and PRIME Cryptographic Suites on page 6064

IPsec VPN Overview

A virtual private network (VPN) provides a means for securely communicating among remote computers across a public WAN such as the Internet.

A VPN connection can link two LANs (site-to-site VPN) or a remote dial-up user and a LAN. The traffic that flows between these two points passes through shared resources such as routers, switches, and other network equipment that make up the public WAN. To secure VPN communication while passing through the WAN, the two participants create an IP Security (IPsec) tunnel.

NOTE: The term *tunnel* does not denote tunnel mode (see "Packet Processing in Tunnel Mode" on page 6049). Instead, it refers to the IPsec connection.

IPsec is a suite of related protocols for cryptographically securing communications at the IP Packet Layer. IPsec also provides methods for the manual and automatic negotiation of security associations (SAs) and key distribution, all the attributes for which are gathered in a domain of interpretation (DOI). The IPsec DOI is a document containing

Copyright © 2016, Juniper Networks, Inc.

definitions for all the security parameters required for the successful negotiation of a VPN tunnel—essentially, all the attributes required for SA and IKE negotiations. See RFC 2407 and RFC 2408 for more information.

This topic includes the following sections:

- IPsec VPN Topologies on page 6042
- Comparison of Policy-Based VPNs and Route-Based VPNs on page 6042
- Security Associations on page 6043
- IPsec Key Management on page 6044
- IPsec Security Protocols on page 6046
- IPsec Tunnel Negotiation on page 6047

IPsec VPN Topologies

The following are some of the IPsec VPN topologies that Junos operating system (OS) supports:

- Site-to-site VPNs—Connects two sites in an organization together and allows secure communications between the sites.
- Hub-and-spoke VPNs—Connects branch offices to the corporate office in an enterprise network. You can also use this topology to connect spokes together by sending traffic through the hub.
- Remote access VPNs—Allows users working at home or traveling to connect to the corporate office and its resources. This topology is sometimes referred to as an *end-to-site tunnel*.

Comparison of Policy-Based VPNs and Route-Based VPNs

NOTE: **Policy-based VPNs are only supported on SRX5400, SRX5600, and SRX5800 devices.**

Table 500 summarizes the differences between policy-based VPNs and route-based VPNs.

Table 500: Comparison Between Policy-Based VPNs and Route-Based VPNs

Policy-Based VPNs	Route-Based VPNs
In policy-based VPNs, a tunnel is treated as an object that, together with source, destination, application, and action, constitutes a tunnel policy that permits VPN traffic.	In route-based VPNs, a policy does not specifically reference a VPN tunnel.
A tunnel policy specifically references a VPN tunnel by name.	A route determines which traffic is sent through the tunnel based on a destination IP address.

Copyright © 2016, Juniper Networks, Inc.

Table 500: Comparison Between Policy-Based VPNs and Route-Based VPNs *(continued)*

Policy-Based VPNs	Route-Based VPNs
The number of policy-based VPN tunnels that you can create is limited by the number of tunnels that the device supports.	The number of route-based VPN tunnels that you create is limited by the number of st0 interfaces (for point-to-point VPNs) or the number of tunnels that the device supports, whichever is lower.
With a policy-based VPN, although you can create numerous tunnel policies referencing the same VPN tunnel, each tunnel policy pair creates an individual IPsec SA with the remote peer. Each SA counts as an individual VPN tunnel.	Because the route, not the policy, determines which traffic goes through the tunnel, multiple policies can be supported with a single SA or VPN.
In a policy-based VPN, the action must be permit and must include a tunnel.	In a route-based VPN, the regulation of traffic is not coupled to the means of its delivery.
The exchange of dynamic routing information is not supported in policy-based VPNs.	Route-based VPNs support the exchange of dynamic routing information through VPN tunnels. You can enable an instance of a dynamic routing protocol, such as OSPF, on an st0 interface that is bound to a VPN tunnel.
If you need more granularity than a route can provide to specify the traffic sent to a tunnel, using a policy-based VPN with security policies is the best choice.	Route-based VPNs uses routes to specify the traffic sent to a tunnel; a policy does not specifically reference a VPN tunnel.
With a policy-based VPN tunnel, you can consider a tunnel as an element in the construction of a policy.	When the security device does a route lookup to find the interface through which it must send traffic to reach an address, it finds a route through a secure tunnel (st0) interface. With a route-based VPN tunnel, you can consider a tunnel as a means for delivering traffic, and can consider the policy as a method for either permitting or denying the delivery of that traffic.

Security Associations

A security association (SA) is a unidirectional agreement between the VPN participants regarding the methods and parameters to use in securing a communication channel. Full bidirectional communication requires at least two SAs, one for each direction. Through the SA, an IPsec tunnel can provide the following security functions:

- Privacy (through encryption)
- Content integrity (through data authentication)
- Sender authentication and—if using certificates—nonrepudiation (through data origin authentication)

The security functions you employ depend on your needs. If you need only to authenticate the IP packet source and content integrity, you can authenticate the packet without applying any encryption. On the other hand, if you are concerned only with preserving privacy, you can encrypt the packet without applying any authentication mechanisms. Optionally, you can both encrypt and authenticate the packet. Most network security designers choose to encrypt, authenticate, and replay-protect their VPN traffic.

Copyright © 2016, Juniper Networks, Inc.

An IPsec tunnel consists of a pair of unidirectional SAs—one SA for each direction of the tunnel—that specify the security parameter index (SPI), destination IP address, and security protocol (Authentication Header [AH] or Encapsulating Security Payload [ESP] employed. An SA groups together the following components for securing communications:

- Security algorithms and keys.
- Protocol mode, either transport or tunnel. Junos OS devices always use tunnel mode. (See "Packet Processing in Tunnel Mode" on page 6049.)
- Key-management method, either manual key or AutoKey IKE. (See "IPsec Key Management" on page 6044.)
- SA lifetime.

For inbound traffic, Junos OS looks up the SA by using the following triplet:

- Destination IP address.
- Security protocol, either AH or ESP. (See "IPsec Security Protocols" on page 6046.)
- Security parameter index (SPI) value.

For outbound VPN traffic, the policy invokes the SA associated with the VPN tunnel.

IPsec Key Management

The distribution and management of keys are critical to using VPNs successfully. Junos OS supports IPsec technology for creating VPN tunnels with three kinds of key creation mechanisms:

- Manual key
- AutoKey IKE with a preshared key or a certificate

You can choose your key creation mechanism—also called authentication method—during Phase 1 and Phase 2 proposal configuration. See "IPsec Tunnel Negotiation" on page 6047.

This topic includes the following sections:

- Manual Key on page 6044
- AutoKey IKE on page 6045
- Diffie-Hellman Exchange on page 6045

Manual Key

With manual keys, administrators at both ends of a tunnel configure all the security parameters. This is a viable technique for small, static networks where the distribution, maintenance, and tracking of keys are not difficult. However, safely distributing manual-key configurations across great distances poses security issues. Aside from passing the keys face-to-face, you cannot be completely sure that the keys have not been compromised while in transit. Also, whenever you want to change the key, you are faced with the same security issues as when you initially distributed it.

Copyright © 2016, Juniper Networks, Inc.

AutoKey IKE

When you need to create and manage numerous tunnels, you need a method that does not require you to configure every element manually. IPsec supports the automated generation and negotiation of keys and security associations using the Internet Key Exchange (IKE) protocol. Junos OS refers to such automated tunnel negotiation as AutoKey IKE and supports AutoKey IKE with preshared keys and AutoKey IKE with certificates.

- AutoKey IKE with preshared keys—Using AutoKey IKE with preshared keys to authenticate the participants in an IKE session, each side must configure and securely exchange the preshared key in advance. In this regard, the issue of secure key distribution is the same as that with manual keys. However, once distributed, an autokey, unlike a manual key, can automatically change its keys at predetermined intervals using the IKE protocol. Frequently changing keys greatly improves security, and automatically doing so greatly reduces key-management responsibilities. However, changing keys increases traffic overhead; therefore, changing keys too often can reduce data transmission efficiency.

NOTE: **A preshared key is a key for both encryption and decryption, which both participants must have before initiating communication.**

- AutoKey IKE with certificates—When using certificates to authenticate the participants during an AutoKey IKE negotiation, each side generates a public-private key pair and acquires a certificate. As long as the issuing certificate authority (CA) is trusted by both sides, the participants can retrieve the peer's public key and verify the peer's signature. There is no need to keep track of the keys and SAs; IKE does it automatically.

Diffie-Hellman Exchange

A Diffie-Hellman (DH) exchange allows participants to produce a shared secret value. The strength of the technique is that it allows participants to create the secret value over an unsecured medium without passing the secret value through the wire. The size of the prime modulus used in each group's calculation differs as follows:

- DH Group 1—768-bit modulus
- DH Group 2—1024-bit modulus
- DH Group 5—1536-bit modulus
- DH Group 14—2048-bit modulus
- DH Group 19—256-bit modulus elliptic curve
- DH Group 20—384-bit modulus elliptic curve
- DH Group 24—2048-bit modulus with 256-bit prime order subgroup

NOTE: **We do not recommend the use of DH groups 1, 2, and 5.**

Copyright © 2016, Juniper Networks, Inc.

Because the modulus for each DH group is a different size, the participants must agree to use the same group.

IPsec Security Protocols

IPsec uses two protocols to secure communications at the IP layer:

- Authentication Header (AH)—A security protocol for authenticating the source of an IP packet and verifying the integrity of its content
- Encapsulating Security Payload (ESP)—A security protocol for encrypting the entire IP packet (and authenticating its content)

You can choose your security protocols—also called *authentication and encryption algorithms*—during Phase 2 proposal configuration. See "IPsec Tunnel Negotiation" on page 6047.

For each VPN tunnel, both AH and ESP tunnel sessions are installed on Services Processing Units (SPUs) and the control plane. For branch SRX Series devices, tunnel sessions are updated with the negotiated protocol after negotiation is completed. For high-end SRX Series devices, tunnel sessions on anchor SPUs are updated with the negotiated protocol while non-anchor SPUs retain ESP and AH tunnel sessions. ESP and AH tunnel sessions are displayed in the outputs for the **show security flow session** and **show security flow cp-session** operational mode commands.

This topic includes the following sections:

- AH Protocol on page 6046
- ESP Protocol on page 6047

AH Protocol

The Authentication Header (AH) protocol provides a means to verify the authenticity and integrity of the content and origin of a packet. You can authenticate the packet by the checksum calculated through a Hash Message Authentication Code (HMAC) using a secret key and either MD5 or SHA-1 hash functions.

- Message Digest 5 (MD5)—An algorithm that produces a 128-bit hash (also called a *digital signature* or *message digest*) from a message of arbitrary length and a 16-byte key. The resulting hash is used, like a fingerprint of the input, to verify content and source authenticity and integrity.
- Secure Hash Algorithm (SHA-1)—An algorithm that produces a 160-bit hash from a message of arbitrary length and a 20-byte key. It is generally regarded as more secure than MD5 because of the larger hashes it produces. Because the computational processing is done in the ASIC, the performance cost is negligible.

NOTE: For more information on MD5 hashing algorithms, see RFC 1321 and RFC 2403. For more information on SHA hashing algorithms, see RFC 2404. For more information on HMAC, see RFC 2104.

Copyright © 2016, Juniper Networks, Inc.

ESP Protocol

The Encapsulating Security Payload (ESP) protocol provides a means to ensure privacy (encryption) and source authentication and content integrity (authentication). ESP in tunnel mode encapsulates the entire IP packet (header and payload) and then appends a new IP header to the now-encrypted packet. This new IP header contains the destination address needed to route the protected data through the network. (See "Packet Processing in Tunnel Mode" on page 6049.)

With ESP, you can both encrypt and authenticate, encrypt only, or authenticate only. For encryption, you can choose one of the following encryption algorithms:

- Data Encryption Standard (DES)—A cryptographic block algorithm with a 56-bit key.
- Triple DES (3DES)—A more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key. DES provides significant performance savings but is considered unacceptable for many classified or sensitive material transfers.
- Advanced Encryption Standard (AES)—An encryption standard which offers greater interoperability with other devices. Junos OS supports AES with 128-bit, 192-bit, and 256-bit keys.

For authentication, you can use either the MD5 or the SHA-1 algorithm.

NOTE: **Even though it is possible to select NULL for encryption, it has been demonstrated that IPsec might be vulnerable to attack under such circumstances. Therefore, we suggest that you choose an encryption algorithm for maximum security.**

IPsec Tunnel Negotiation

To establish an AutoKey IKE IPsec tunnel, two phases of negotiation are required:

- In Phase 1, the participants establish a secure channel in which to negotiate the IPsec security associations (SAs).
- In Phase 2, the participants negotiate the IPsec SAs for encrypting and authenticating the ensuing exchanges of user data.

For a manual key IPsec tunnel, because all the SA parameters have been previously defined, there is no need to negotiate which SAs to use. In essence, the tunnel has already been established. When traffic matches a policy using that manual key tunnel or when a route involves the tunnel, the Juniper Networks device simply encrypts and authenticates the data, as you determined, and forwards it to the destination gateway.

The remote IKE gateway address can be in any virtual routing (VR) instance. VR is determined during IKE Phase 1 and Phase 2 negotiation. VR does not have to be configured in the IKE proposals. If the IKE gateway interface is moved from one VR to another, the

Copyright © 2016, Juniper Networks, Inc.

existing IKE Phase 1 and Phase 2 negotiations for the IKE gateway are cleared, and new Phase 1 and Phase 2 negotiations are performed.

NOTE:

- On SRX Series devices, when you enable VPN, overlapping of IP addresses across virtual routers is supported with the following limitations:
  - An IKE external interface address cannot overlap with any other virtual router.
  - An internal or trust interface address can overlap across virtual routers.
  - An st0 interface address cannot overlap in route-based VPN in point-to-multipoint tunnel such as NHTB.
  - An st0 interface address can overlap in route-based VPN in point-to-point tunnel.
- The combinations of local IP addresses and remote gateway IP addresses of IP sec VPN tunnels configured across VRs have to be unique.
- When the loopback interface is used as the IKE gateway external interface, the physical interface for IKE negotiation should be in the same VR.

Related Documentation

- Example: Configuring a Policy-Based VPN on page 6222
- Example: Configuring a Route-Based VPN on page 6074
- Understanding IKE and IPsec Packet Processing on page 6048
- Understanding Phase 1 of IKE Tunnel Negotiation on page 6055
- Understanding Phase 2 of IKE Tunnel Negotiation on page 6057
- Understanding Hub-and-Spoke VPNs on page 6093

Understanding IKE and IPsec Packet Processing

An IPsec VPN tunnel consists of tunnel setup and applied security. During tunnel setup, the peers establish security associations (SAs), which define the parameters for securing traffic between themselves. (See "IPsec VPN Overview" on page 6041.) After the tunnel is established, IPsec protects the traffic sent between the two tunnel endpoints by applying the security parameters defined by the SAs during tunnel setup. Within the Junos OS implementation, IPsec is applied in tunnel mode, which supports the Encapsulating Security Payload (ESP) and Authentication Header (AH) protocols.

This topic includes the following sections:

- Packet Processing in Tunnel Mode on page 6049
- IKE Packet Processing on page 6050
- IPsec Packet Processing on page 6053

Copyright © 2016, Juniper Networks, Inc.

Packet Processing in Tunnel Mode

IPsec operates in one of two modes—transport or tunnel. When both ends of the tunnel are hosts, you can use either mode. When at least one of the endpoints of a tunnel is a security gateway, such as a Junos OS router or firewall, you must use tunnel mode. Juniper Networks devices always operate in tunnel mode for IPsec tunnels.

In tunnel mode, the entire original IP packet—payload and header—is encapsulated within another IP payload, and a new header is appended to it, as shown in Figure 240. The entire original packet can be encrypted, authenticated, or both. With the Authentication Header (AH) protocol, the AH and new headers are also authenticated. With the Encapsulating Security Payload (ESP) protocol, the ESP header can also be authenticated.

Figure 240: Tunnel Mode

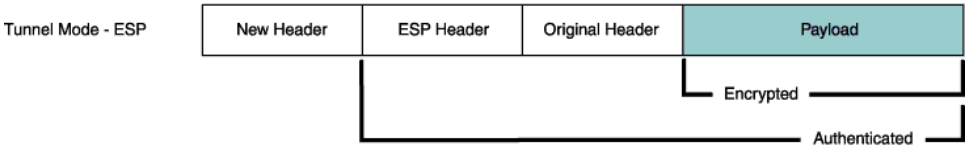

In a site-to-site VPN, the source and destination addresses used in the new header are the IP addresses of the outgoing interface. See Figure 241.

Figure 241: Site-to-Site VPN in Tunnel Mode

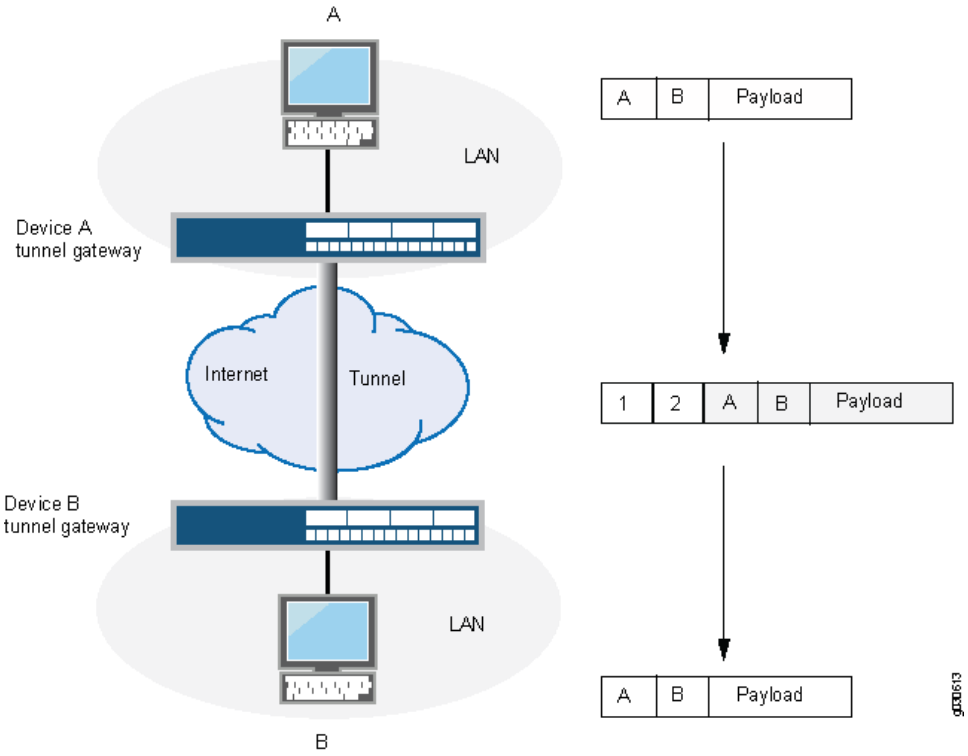

In a dial-up VPN, there is no tunnel gateway on the VPN dial-up client end of the tunnel; the tunnel extends directly to the client itself (see Figure 242). In this case, on packets

Copyright © 2016, Juniper Networks, Inc.

sent from the dial-up client, both the new header and the encapsulated original header have the same IP address: that of the client's computer.

NOTE: Some VPN clients, such as the dynamic VPN client and Netscreen-Remote, use a virtual inner IP address (also called a "sticky address"). Netscreen-Remote enables you to define the virtual IP address. The dynamic VPN client uses the virtual IP address assigned during the XAuth configuration exchange. In such cases, the virtual inner IP address is the source IP address in the original packet header of traffic originating from the client, and the IP address that the ISP dynamically assigns the dial-up client is the source IP address in the outer header.

Figure 242: Dial-Up VPN in Tunnel Mode

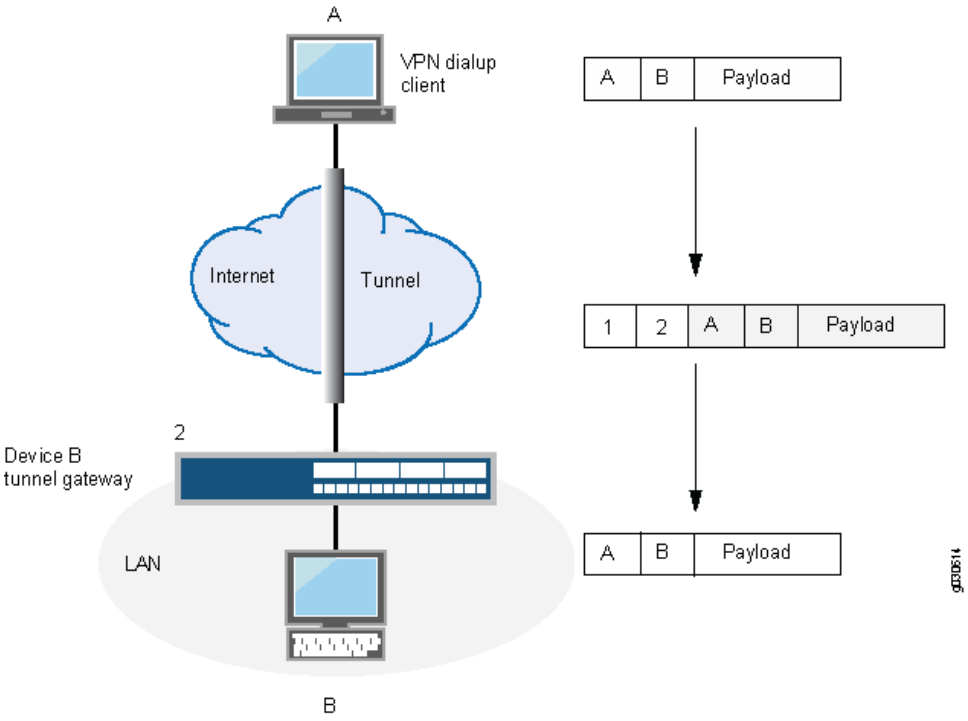

IKE Packet Processing

When a cleartext packet arrives on a Juniper Networks device that requires tunneling, and no active Phase 2 SA exists for that tunnel, Junos OS begins IKE negotiations and drops the packet. The source and destination addresses in the IP packet header are those of the local and remote IKE gateways, respectively. In the IP packet payload, there is a UDP segment encapsulating an ISAKMP (IKE) packet. The format for IKE packets is the same for Phase 1 and Phase 2. See Figure 243.

Meanwhile, the source host has sent the dropped packet again. Typically, by the time the second packet arrives, IKE negotiations are complete, and Junos OS protects the packet and all subsequent packets in the session—with IPsec before forwarding it.

Copyright © 2016, Juniper Networks, Inc.

Figure 243: IKE Packet for Phases 1 and 2

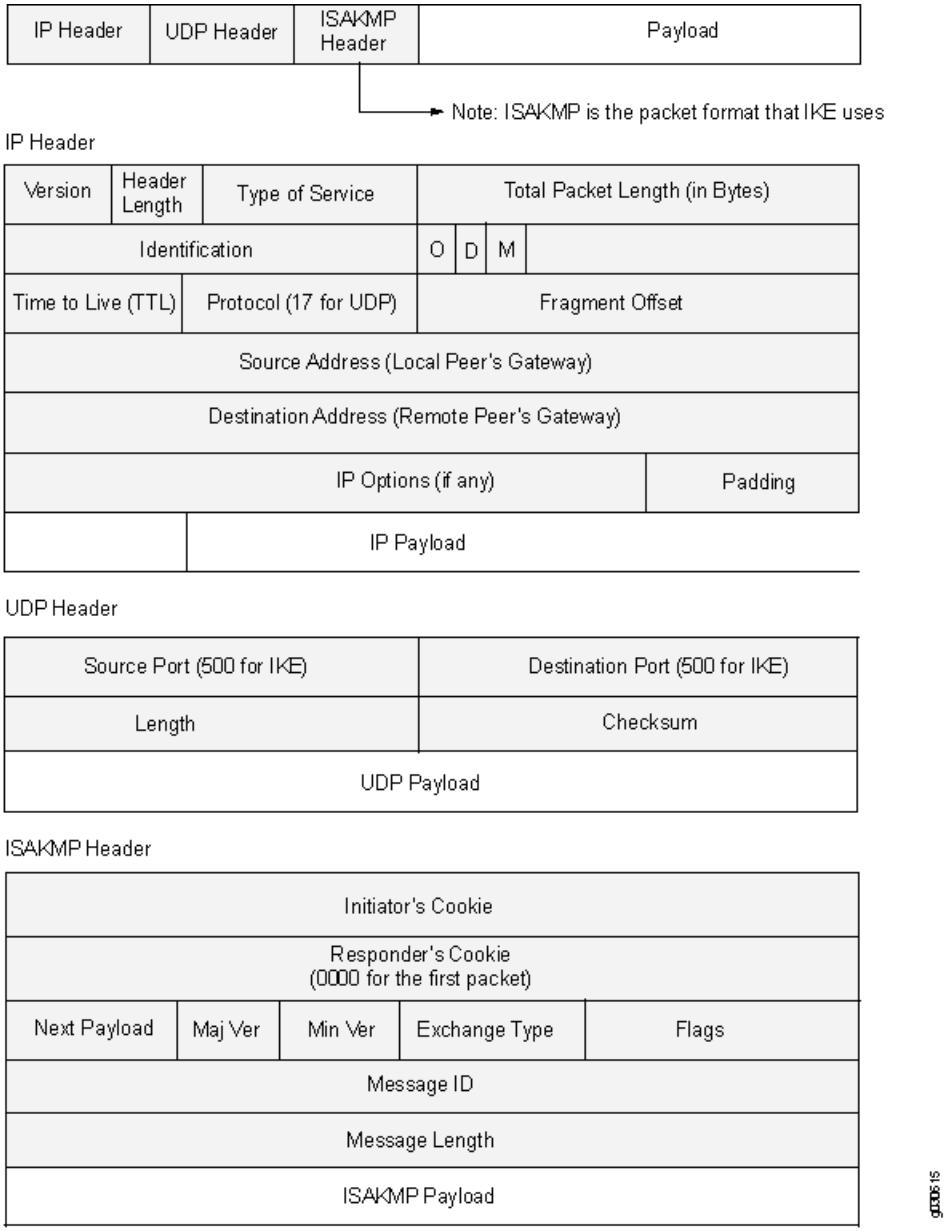

The Next Payload field contains a number indicating one of the following payload types:

- 0002—SA Negotiation Payload contains a definition for a Phase 1 or Phase 2 SA.
- 0004—Proposal Payload can be a Phase 1 or Phase 2 proposal.
- 0008—Transform Payload gets encapsulated in a proposal payload that gets encapsulated in an SA payload.
- 0010—Key Exchange (KE) Payload contains information necessary for performing a key exchange, such as a DH public value.

Copyright © 2016, Juniper Networks, Inc.

- 0020—Identification (IDx) Payload.
  - In Phase 1, IDii indicates the initiator ID, and IDir indicates the responder ID.
  - In Phase 2, IDui indicates the user initiator, and IDur indicates the user responder.

  The IDs are IKE ID types such as FQDN, U-FQDN, IP address, and ASN.1_DN.
- 0040—Certificate (CERT) Payload.
- 0080—Certificate Request (CERT_REQ) Payload.
- 0100—Hash (HASH) Payload contains the digest output of a particular hash function.
- 0200—Signature (SIG) Payload contains a digital signature.
- 0400—Nonce (Nx) Payload contains some pseudorandom information necessary for the exchange).
- 0800—Notify Payload.
- 1000—ISAKMP Delete Payload.
- 2000—Vendor ID (VID) Payload can be included anywhere in Phase 1 negotiations. Junos OS uses it to mark support for NAT-T.

Each ISAKMP payload begins with the same generic header, as shown in Figure 244.

Figure 244: Generic ISAKMP Payload Header

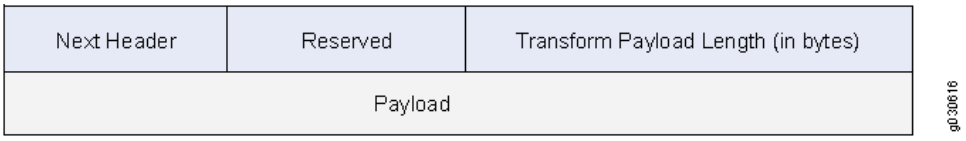

There can be multiple ISAKMP payloads chained together, with each subsequent payload type indicated by the value in the Next Header field. A value of **0000** indicates the last ISAKMP payload. See Figure 245 for an example.

Copyright © 2016, Juniper Networks, Inc.

Figure 245: ISAKMP Header with Generic ISAKMP Payloads

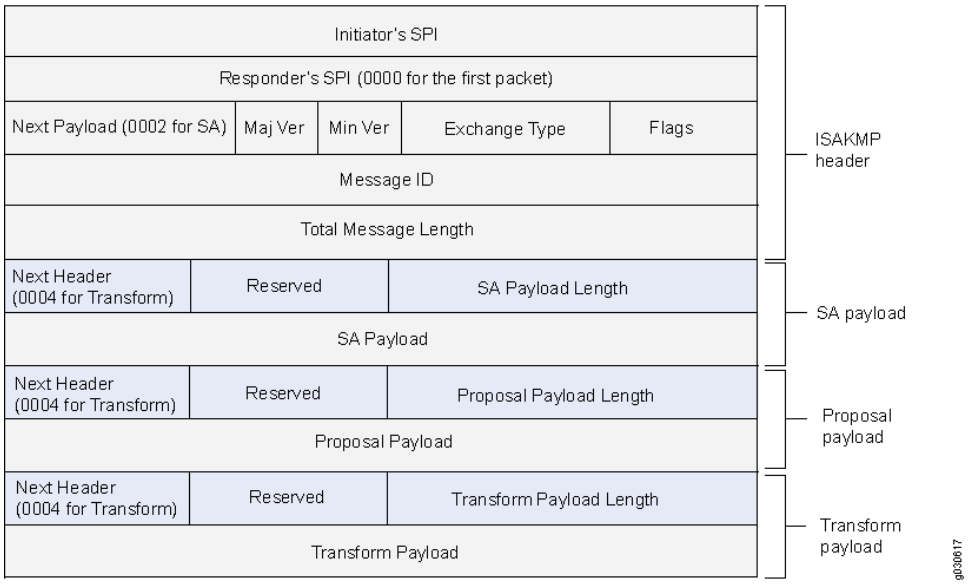

IPsec Packet Processing

After IKE negotiations complete and the two IKE gateways have established Phase 1 and Phase 2 security associations (SAs), all subsequent packets are forwarded using the tunnel. If the Phase 2 SA specifies the Encapsulating Security Protocol (ESP) in tunnel mode, the packet looks like the one shown in Figure 246. The device adds two additional headers to the original packet that the initiating host sends.

NOTE: For information about ESP, see "ESP Protocol" on page 6047. For information about tunnel mode, see "Packet Processing in Tunnel Mode" on page 6049.

As shown in Figure 246, the packet that the initiating host constructs includes the payload, the TCP header, and the inner IP header (IP1).

Figure 246: IPsec Packet—ESP in Tunnel Mode

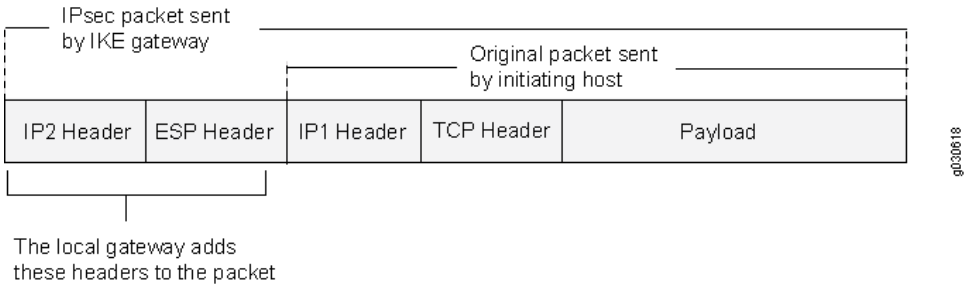

The router IP header (IP2), which Junos OS adds, contains the IP address of the remote gateway as the destination IP address and the IP address of the local router as the source IP address. Junos OS also adds an ESP header between the outer and inner IP headers.

Copyright © 2016, Juniper Networks, Inc.

The ESP header contains information that allows the remote peer to properly process the packet when it receives it. This is shown in Figure 247.

Figure 247: Outer IP Header (IP2) and ESP Header

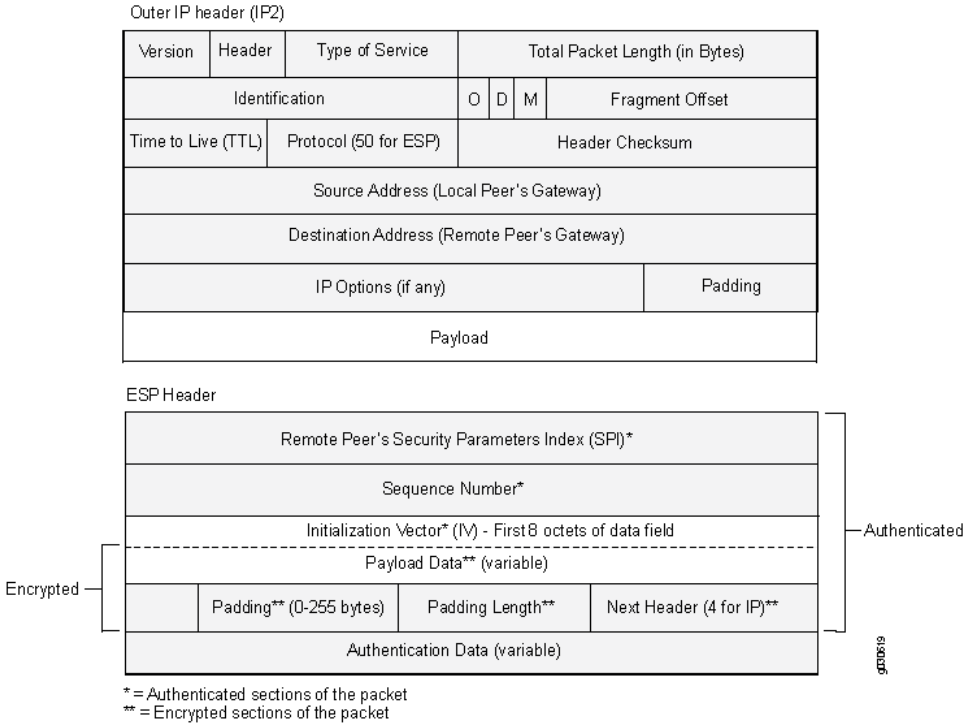

The Next Header field indicates the type of data in the payload field. In tunnel mode, this value is 4, indicating an IP packet is contained within the payload. See Figure 248.

Copyright © 2016, Juniper Networks, Inc.

Figure 248: Inner IP Header (IP1) and TCP Header

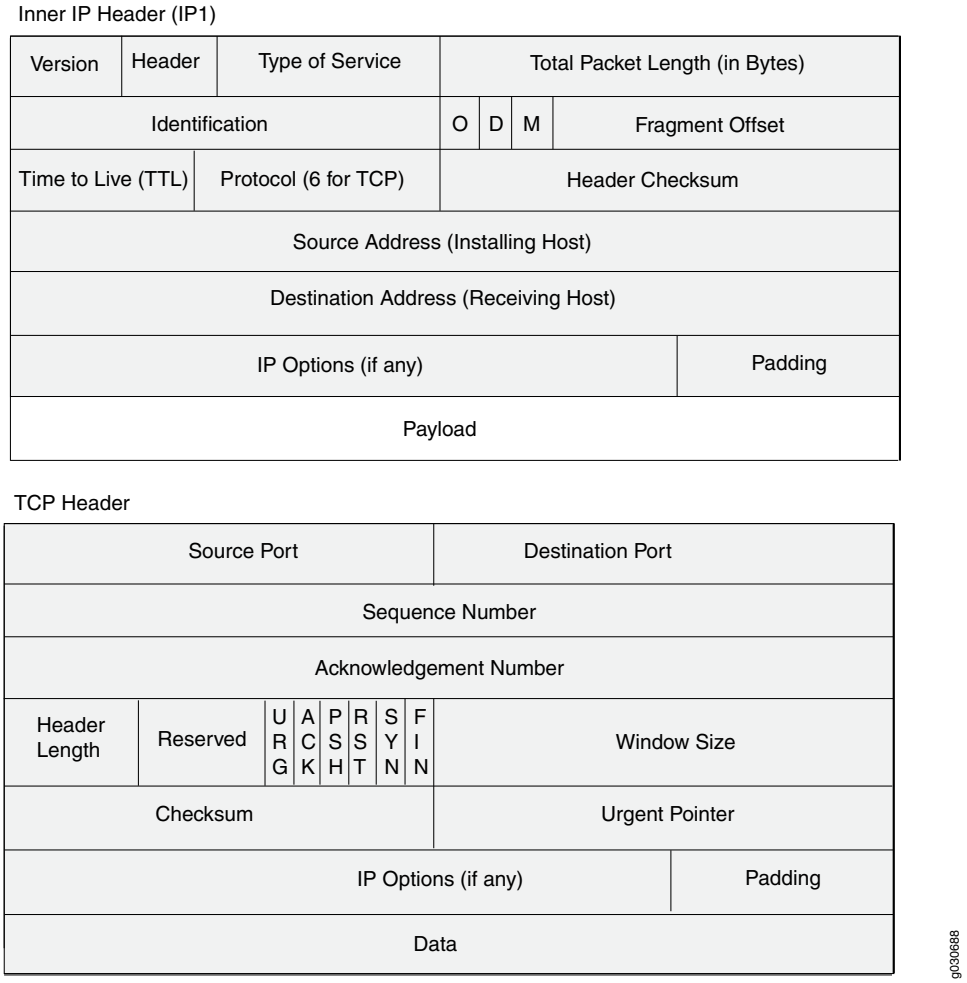

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding Phase 1 of IKE Tunnel Negotiation on page 6055
- Understanding Phase 2 of IKE Tunnel Negotiation on page 6057
- Understanding Hub-and-Spoke VPNs on page 6093
- Example: Configuring a Policy-Based VPN on page 6222
- Example: Configuring a Route-Based VPN on page 6074

Understanding Phase 1 of IKE Tunnel Negotiation

Phase 1 of an AutoKey Internet Key Exchange (IKE) tunnel negotiation consists of the exchange of proposals for how to authenticate and secure the channel. The participants exchange proposals for acceptable security services such as:

Copyright © 2016, Juniper Networks, Inc.

- Encryption algorithms—Data Encryption Standard (DES), triple Data Encryption Standard (3DES), and Advanced Encryption Standard (AES). (See "IPsec Security Protocols" on page 6046.)
- Authentication algorithms—Message Digest 5 (MD5) and Secure Hash Algorithm (SHA-1). (See "IPsec Security Protocols" on page 6046.)
- Diffie-Hellman (DH) group. (See "Diffie-Hellman Exchange" on page 6045.)
- Preshared key or RSA/DSA certificates. (See "IPsec Key Management" on page 6044.)

A successful Phase 1 negotiation concludes when both ends of the tunnel agree to accept at least one set of the Phase 1 security parameters proposed and then process them. Juniper Networks devices support up to four proposals for Phase 1 negotiations, allowing you to define how restrictive a range of security parameters for key negotiation you will accept.

Junos OS provides the following predefined Phase 1 proposals:

- Standard—pre-g2-aes128-sha and pre-g2-3des-sha
- Compatible—pre-g2-3des-sha, pre-g2-3des-md5, pre-g2-des-sha, and pre-g2-des-md5
- Basic—pre-g1-des-sha and pre-g1-des-md5

You can also define custom Phase 1 proposals.

Phase 1 exchanges can take place in either main mode or aggressive mode. You can choose your mode during IKE policy configuration.

This topic includes the following sections:

- Main Mode on page 6056
- Aggressive Mode on page 6057

Main Mode

In main mode, the initiator and recipient send three two-way exchanges (six messages total) to accomplish the following services:

- First exchange (messages 1 and 2)—Proposes and accepts the encryption and authentication algorithms.
- Second exchange (messages 3 and 4)—Executes a DH exchange, and the initiator and recipient each provide a pseudorandom number.
- Third exchange (messages 5 and 6)—Sends and verifies the identities of the initiator and recipient.

The information transmitted in the third exchange of messages is protected by the encryption algorithm established in the first two exchanges. Thus, the participants' identities are encrypted and therefore not transmitted "in the clear."

Copyright © 2016, Juniper Networks, Inc.

Aggressive Mode

In aggressive mode, the initiator and recipient accomplish the same objectives as with main mode, but in only two exchanges, with a total of three messages:

- First message—The initiator proposes the security association (SA), initiates a DH exchange, and sends a pseudorandom number and its IKE identity.

  NOTE: When configuring aggressive mode with multiple proposals for Phase 1 negotiations, use the same DH group in all proposals because the DH group cannot be negotiated. Up to four proposals can be configured.

- Second message—The recipient accepts the SA; authenticates the initiator; and sends a pseudorandom number, its IKE identity, and, if using certificates, the recipient's certificate.
- Third message—The initiator authenticates the recipient, confirms the exchange, and, if using certificates, sends the initiator's certificate.

Because the participants' identities are exchanged in the clear (in the first two messages), aggressive mode does not provide identity protection.

NOTE: When a dial-up VPN user negotiates an AutoKey IKE tunnel with a preshared key, aggressive mode must be used. Note also that a dial-up VPN user can use an e-mail address, a fully qualified domain name (FQDN), or an IP address as its IKE ID. A dynamic peer can use either an e-mail address or an FQDN, but not an IP address.

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding Phase 2 of IKE Tunnel Negotiation on page 6057
- Example: Configuring a Policy-Based VPN on page 6222
- Example: Configuring a Route-Based VPN on page 6074

Understanding Phase 2 of IKE Tunnel Negotiation

After the participants have established a secure and authenticated channel, they proceed through Phase 2, in which they negotiate security associations (SAs) to secure the data to be transmitted through the IPsec tunnel.

Similar to the process for Phase 1, the participants exchange proposals to determine which security parameters to employ in the SA. A Phase 2 proposal also includes a security protocol—either Encapsulating Security Payload (ESP) or Authentication Header (AH)—and selected encryption and authentication algorithms. The proposal can also specify a Diffie-Hellman (DH) group, if Perfect Forward Secrecy (PFS) is desired.

Copyright © 2016, Juniper Networks, Inc.

Regardless of the mode used in Phase 1, Phase 2 always operates in quick mode and involves the exchange of three messages.

Juniper Networks devices support up to four proposals for Phase 2 negotiations, allowing you to define how restrictive a range of tunnel parameters you will accept. Junos OS provides the following predefined Phase 2 proposals:

- Standard—g2-esp-3des-sha and g2-esp-aes128-sha
- Compatible—nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5
- Basic—nopfs-esp-des-sha and nopfs-esp-des-md5

You can also define custom Phase 2 proposals.

This topic includes the following sections:

- Proxy IDs on page 6058
- Perfect Forward Secrecy on page 6058
- Replay Protection on page 6058

Proxy IDs

In Phase 2, the peers exchange proxy IDs. A proxy ID consists of a local and remote IP address prefix. The proxy ID for both peers must match, which means that the local IP address specified for one peer must be the same as the remote IP address specified for the other peer.

Perfect Forward Secrecy

PFS is a method for deriving Phase 2 keys independent from and unrelated to the preceding keys. Alternatively, the Phase 1 proposal creates the key (the SKEYID_d key) from which all Phase 2 keys are derived. The SKEYID_d key can generate Phase 2 keys with a minimum of CPU processing. Unfortunately, if an unauthorized party gains access to the SKEYID_d key, all your encryption keys are compromised.

PFS addresses this security risk by forcing a new DH key exchange to occur for each Phase 2 tunnel. Using PFS is thus more secure, although the rekeying procedure in Phase 2 might take slightly longer with PFS enabled.

Replay Protection

A replay attack occurs when an unauthorized person intercepts a series of packets and uses them later either to flood the system, causing a denial of service (DoS), or to gain entry to the trusted network. Junos OS provides a replay protection feature that enables devices to check every IPsec packet to see if it has been received previously. If packets arrive outside a specified sequence range, Junos OS rejects them. Use of this feature does not require negotiation, because packets are always sent with sequence numbers. You simply have the option of checking or not checking the sequence numbers.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring a Policy-Based VPN on page 6222
- Example: Configuring a Route-Based VPN on page 6074

IPsec VPN with Autokey IKE Configuration Overview

IPsec VPN negotiation occurs in two phases. In Phase 1, participants establish a secure channel in which to negotiate the IPsec security association (SA). In Phase 2, participants negotiate the IPsec SA for authenticating traffic that will flow through the tunnel.

This overview describes the basic steps to configure a route-based or policy-based IPsec VPN using autokey IKE (preshared keys or certificates).

To configure a route-based or policy-based IPsec VPN using autokey IKE:

1. Configure interfaces, security zones, and address book information.

   (For route-based VPNs) Configure a secure tunnel **st0.*x*** interface. Configure routing on the device.

2. Configure Phase 1 of the IPsec VPN tunnel.

   a. (Optional) Configure a custom IKE Phase 1 proposal. This step is optional, as you can use a predefined IKE Phase 1 proposal set (Standard, Compatible, or Basic).

   b. Configure an IKE policy that references either your custom IKE Phase 1 proposal or a predefined IKE Phase 1 proposal set. Specify autokey IKE preshared key or certificate information. Specify the mode (main or aggressive) for the Phase 1 exchanges.

   c. Configure an IKE gateway that references the IKE policy. Specify the IKE IDs for the local and remote devices. If the IP address of the remote gateway is not known, specify how the remote gateway is to be identified.

3. Configure Phase 2 of the IPsec VPN tunnel.

   a. (Optional) Configure a custom IPsec Phase 2 proposal. This step is optional, as you can use a predefined IPsec Phase 2 proposal set (Standard, Compatible, or Basic).

   b. Configure an IPsec policy that references either your custom IPsec Phase 2 proposal or a predefined IPsec Phase 2 proposal set. Specify perfect forward secrecy (PFS) keys.

   c. Configure an IPsec VPN tunnel that references both the IKE gateway and the IPsec policy. Specify the proxy IDs to be used in Phase 2 negotiations.

      (For route-based VPNs) Bind the secure tunnel interface **st0.*x*** to the IPsec VPN tunnel.

4. Configure a security policy to permit traffic from the source zone to the destination zone.

Copyright © 2016, Juniper Networks, Inc.

(For policy-based VPNs) Specify the security policy action **tunnel ipsec-vpn** with the name of the IPsec VPN tunnel that you configured.

5. Update your global VPN settings. See "Example: Configuring Global SPI and VPN Monitoring Features" on page 6786.

Related Documentation

- Understanding Route-Based IPsec VPNs on page 6073
- Understanding Policy-Based IPsec VPNs on page 6221
- Configuring IPsec VPN Using the VPN Wizard on page 6063

IPsec VPN with Manual Keys Configuration Overview

This overview describes the basic steps to configure a route-based or policy-based IPsec VPN using manual keys.

To configure a route-based or policy-based IPsec VPN using manual keys:

1. Configure interfaces, security zones, and address book information.

   (For route-based VPNs) Configure routing. Configure a secure tunnel **st0.*x*** interface.

2. Configure an IPsec VPN tunnel by specifying the following parameters:

   - Authentication algorithm and key
   - Encryption algorithm and key
   - Outgoing interface
   - IP address of the peer
   - IPsec protocol for the security association
   - Security parameter index

   (For route-based VPNs) Bind the secure tunnel interface **st0.*x*** to the IPsec VPN tunnel.

3. Configure security policy to permit traffic from the source zone to the destination zone.

   (For policy-based VPNs) Specify the security policy action **tunnel ipsec-vpn** with the name of the IPsec VPN tunnel that you configured.

Related Documentation

- Understanding Route-Based IPsec VPNs on page 6073
- Understanding Policy-Based IPsec VPNs on page 6221
- Example: Configuring an IPv6 IPsec Manual VPN on page 6341

Recommended Configuration Options for Site-to-Site VPN with Static IP Addresses

Table 501 lists the configuration options for a generic site-to-site VPN between two security devices with static IP addresses. The VPN can be either route-based or policy-based.

Copyright © 2016, Juniper Networks, Inc.

Table 501: Recommended Configuration for Site-to-Site VPN with Static IP Addresses

Configuration Option	Comment
IKE configuration:	
Autokey IKE with certificates	Manual key is not recommended.
Main mode	Used when peers have static IP addresses.
RSA or DSA certificates	RSA or DSA certificates can be used on the local device. Specify the type of certificate (PKCS7 or X.509) on the peer.
Diffie-Hellman (DH) group 2	DH group 2 provides more security than DH group 1 and incurs less processing overhead than DH groups 5 and 14.
Advanced Encryption Standard (AES) encryption	AES is cryptographically stronger than Data Encryption Standard (DES) and Triple DES (3DES) when key lengths are equal. Approved encryption algorithm for Federal Information Processing Standards (FIPS) and Common Criteria EAL4 standards.
Secure Hash Algorithm (SHA-1) authentication	SHA-1 provides more cryptographic security than Message Digest 5 (MD5) and incurs less processing overhead than SHA-256.
Local and remote IKE identifiers	IP addresses are used by default if local and remote peers have static IP addresses. Specify the peer's address. IP addresses can be used with a certificate if the IP address appears in the SubjectAltName field.
IPsec configuration:	
Perfect Forward Secrecy (PFS) DH group 2	PFS DH group 2 provides increased security because the peers perform a second DH exchange to produce the key used for IPsec encryption and decryption.
Encapsulating Security Payload (ESP) protocol	ESP provides both confidentiality through encryption and encapsulation of the original IP packet and integrity through authentication.
AES encryption	AES is cryptographically stronger than DES and 3DES when key lengths are equal. Approved encryption algorithm for FIPS and Common Criteria EAL4 standards.
SHA-1 authentication	SHA-1 provides more cryptographic security than MD5.
Anti-replay protection	Enabled by default. Disabling this feature might resolve compatibility issues with third-party peers.

Related Documentation

- IPsec VPN Overview on page 6041

Recommended Configuration Options for Site-to-Site or Dialup VPNs with Dynamic IP Addresses

Table 502 lists the configuration options for a generic site-to-site or dialup VPN, where the peer devices have dynamic IP addresses.

Copyright © 2016, Juniper Networks, Inc.

Table 502: Recommended Configuration for Site-to-Site or Dialup VPNs with Dynamic IP Addresses

Configuration Option	Comment
IKE configuration:	
Autokey IKE with certificates	Manual key is not recommended.
Aggressive mode	Required when the IP address of one or both of the IPsec peers is dynamically assigned.
1024-bit certificates	RSA or DSA certificates can be used. Specify the certificate to be used on the local device. Specify the type of certificate (PKCS7 or X.509) on the peer.
Diffie-Hellman (DH) group 2	DH group 2 provides more security than DH group 1 and incurs less processing overhead than DH groups 5 and 14.
Advanced Encryption Standard (AES) encryption	AES is cryptographically stronger than Data Encryption Standard (DES) and Triple DES (3DES) when key lengths are equal. Approved encryption algorithm for Federal Information Processing Standards (FIPS) and Common Criteria EAL4 standards.
Secure Hash Algorithm (SHA-1) authentication	SHA-1 provides more cryptographic security than Message Digest 5 (MD5) and incurs less processing overhead than SHA-256.
Local and remote IKE IDs	User-Fully Qualified Domain Name (U-FQDN) is an e-mail address that can be used with a certificate if the U-FQDN appears in the SubjectAltName field.
IPsec configuration:	
Perfect Forward Secrecy (PFS) DH group 2	PFS DH group 2 provides increased security because the peers perform a second DH exchange to produce the key used for IPsec encryption and decryption.
Encapsulating Security Payload (ESP) protocol	ESP provides both confidentiality through encryption and encapsulation of the original IP packet and integrity through authentication.
AES encryption	AES is cryptographically stronger than DES and 3DES when key lengths are equal. Approved encryption algorithm for FIPS and Common Criteria EAL4 standards.
SHA-1 authentication	SHA-1 provides more cryptographic security than MD5.
Anti-replay protection	Enabled by default. Disabling this might resolve compatibility issues with third-party peers.

Related Documentation

- IPsec VPN Overview on page 6041

Configuring Remote IKE IDs for Site-to-Site VPNs

As of Junos OS Release 11.4, checks are performed to validate the IKE ID received from the VPN peer device. By default, SRX Series devices validate the IKE ID received from the peer with the IP address configured for the IKE gateway. In certain network setups, the

Copyright © 2016, Juniper Networks, Inc.

IKE ID received from the peer (which can be an IPv4 or IPv6 address, fully qualified domain name [FQDN], distinguished name, or e-mail address) does not match the IKE gateway configured on the SRX Series device. This can lead to a Phase 1 validation failure.

To modify the configuration of the SRX Series device or the peer device for the IKE ID that is used:

- On the SRX Series device, configure the **remote-identity** statement at the [**edit security ike gateway *gateway-name***] hierarchy level to match the IKE ID that is received from the peer. Values can be an IPv4 or IPv6 address, FQDN, distinguished name, or e-mail address.

  NOTE: If you do not configure remote-identity, the device uses the IPv4 or IPv6 address that corresponds to the remote peer by default.

- On the peer device, ensure that the IKE ID is the same as the **remote-identity** configured on the SRX Series device. If the peer device is an SRX Series device, configure the **local-identity** statement at the [**edit security ike gateway *gateway-name***] hierarchy level. Values can be an IPv4 or IPv6 address, FQDN, distinguished name, or e-mail address.

Related Documentation

- Understanding NAT-T on page 6243
- Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device on page 6244
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272

Configuring IPsec VPN Using the VPN Wizard

The VPN Wizard enables you to perform basic IPsec VPN configuration, including both Phase 1 and Phase 2. For more advanced configuration, use the J-Web interface or the CLI.

To configure IPsec VPN using the VPN Wizard:

1. Select **Configure>Tasks>Configure VPN** in the J-Web interface.
2. Click the Launch VPN Wizard button.
3. Follow the wizard prompts.

The upper left area of the wizard page shows where you are in the configuration process. The lower left area of the page shows field-sensitive help. When you click a link under the Resources heading, the document opens in your browser. If the document opens in a new tab, be sure to close only the tab (not the browser window) when you close the document.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding Phase 1 of IKE Tunnel Negotiation on page 6055
- Understanding Phase 2 of IKE Tunnel Negotiation on page 6057

Understanding Suite B and PRIME Cryptographic Suites

Suite B is a set of cryptographic algorithms designated by the U.S. National Security Agency to allow commercial products to protect traffic that is classified at secret or top secret levels. Suite B protocols are defined in RFC 6379, *Suite B Cryptographic Suites for IPsec*. The Suite B cryptographic suites provide Encapsulating Security Payload (ESP) integrity and confidentiality and should be used when ESP integrity protection and encryption are both required. Protocol Requirements for IP Modular Encryption (PRIME), an IPsec profile defined for public sector networks in the United Kingdom, is based on the Suite B cryptographic suite, but uses AES-GCM rather than AES-CBC for IKEv2 negotiations.

The following cryptographic suites are supported:

- Suite-B-GCM-128
  - ESP: Advanced Encryption Standard (AES) encryption with 128-bit keys and 16-octet integrity check value (ICV) in Galois Counter Mode (GCM).
  - IKE: AES encryption with 128-bit keys in cipher block chaining (CBC) mode, integrity using SHA-256 authentication, key establishment using Diffie-Hellman (DH) group 19, and authentication using Elliptic Curve Digital Signature Algorithm (ECDSA) 256-bit elliptic curve signatures.
- Suite-B-GCM-256
  - ESP: AES encryption with 256-bit keys and 16-octet ICV in GCM for ESP.
  - IKE: AES encryption with 256-bit keys in CBC mode, integrity using SHA-384 authentication, key establishment using DH group 20, and authentication using ECDSA 384-bit elliptic curve signatures.
- PRIME-128
  - ESP: AES encryption with 128-bit keys and 16-octet ICV in GCM.
  - IKE: AES encryption with 128-bit keys in GCM, key establishment using DH group 19, and authentication using ECDSA 256-bit elliptic curve signatures.
- PRIME-256
  - ESP: AES encryption with 256-bit keys and 16-octet ICV in GCM for ESP.
  - IKE: AES encryption with 256-bit keys in GCM, key establishment using DH group 20, and authentication using ECDSA 384-bit elliptic curve signatures.

Suite-B cryptographic suites support IKEv1 and IKEv2. PRIME cryptographic suites only support IKEv2.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Suite B and PRIME are not fully supported on SRX5400, SRX5600 and SRX5800 devices that do not have the SPC2 (SRX5K-SPC-4-14-320). You can configure IKE with Suite B options on these devices, but AES-GCM options are not supported. If you configure IKE with Suite B options on these devices, VPN establishment is slower because the devices do not have the hardware processors that can accelerate Suite B algorithm processing.

NOTE: Suite B and PRIME are not supported with the Group VPNv2 feature.

CLI options support Suite B and PRIME compliance in IKE and IPsec proposal configuration:

- For IKE proposals configured at the [**edit security ike proposal** ***proposal-name***] hierarchy level:
  - **authentication-algorithm** options include **sha-256** and **sha-384**.
  - **authentication-method** options include **ecdsa-signatures-256** and **ecdsa-signatures-384**.
  - **dh-group** options include **group19** and **group20**.
  - **encryption-algorithm** options for PRIME include **aes-128-gcm** and **aes-256-gcm**.
- For IPsec proposals configured at the [**edit security ipsec proposal** ***proposal-name***] hierarchy level, **encryption-algorithm** options include **aes-128-gcm**, **aes-192-gcm**, and **aes-256-gcm**.
- For IPsec policies configured at the [**edit security ipsec policy** ***policy-name***] hierarchy level, the **perfect-forward-secrecy keys** options include **group19** and **group20**.
- For convenience, predefined proposals that provide compliance with Suite B (**suiteb-gcm-128** and **suiteb-gcm-256**) and PRIME (**prime-128** and **prime-256**) are available at the [**edit security ike policy** ***policy-name***] and [**edit security ipsec policy** ***policy-name***] hierarchy levels.

NOTE: VPN monitoring and cryptographic configuration options **ecdsa-signatures-384** (for IKE authentication) and DH group 20 consume considerable CPU resources. If VPN monitoring and the **ecdsa-signatures-384** and **group20** options are used on an SRX Series device with a large number of tunnels configured, the SRX Series device must have the SPC2 installed.

Related Documentation

- IPsec VPN Overview on page 6041
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 253

Understanding VPN Tunnel Management

- Understanding Distributed VPNs in SRX Series Services Gateways on page 6067
- Understanding VPN Support for Inserting Services Processing Cards on page 6068

Understanding Distributed VPNs in SRX Series Services Gateways

In the SRX5000 lines, the IKE provides tunnel management for IPsec and authenticates end entities. The IKE performs a Diffie-Hellman (DH) key exchange to generate an IPsec tunnel between network devices. The IPsec tunnels generated by IKE are used to encrypt, decrypt, and authenticate user traffic between the network devices at the IP layer.

The VPN is created by distributing the IKE and IPsec workload among the multiple Services Processing Units (SPUs) of the platform. For site-to-site tunnels, the least-loaded SPU is chosen as the anchor SPU. If multiple SPUs have the same smallest load, any of them can be chosen as an anchor SPU. Here, load corresponds to the number of site-to-site gateways or manual VPN tunnels anchored on an SPU. For dynamic tunnels, the newly established dynamic tunnels employ a round-robin algorithm to select the SPU.

In IPsec, the workload is distributed by the same algorithm that distributes the IKE. The Phase 2 SA for a given VPN tunnel termination points pair is exclusively owned by a particular SPU, and all IPsec packets belonging to this Phase 2 SA are forwarded to the anchoring SPU of that SA for IPsec processing.

Multiple IPsec sessions (Phase 2 SA) can operate over one or more IKE sessions. The SPU that is selected for anchoring the IPsec session is based on the SPU that is anchoring the underlying IKE session. Therefore, all IPsec sessions that run over a single IKE gateway are serviced by the same SPU and are not load-balanced across several SPUs.

Table 503 shows an example of load balancing on an SRX5000 line device with three SPUs running eight IPsec tunnels over four IKE gateways. Note that SPU 0, SPU 1, or SPU 2 could be selected for IKE gateway 4 because all three SPUs have an equal load of one gateway each.

Copyright © 2016, Juniper Networks, Inc.

Table 503: Load Balancing Across SPUs

SPU	IKE Gateway	IPsec Tunnel
SPU0	IKE-1	IPsec-1
		IPsec-2
		IPsec-3
SPU1	IKE-2	IPsec-4
		IPsec-5
		IPsec-6
SPU2	IKE-3	IPsec-7
SPU0, SPU1, or SPU2	IKE-4	IPsec-8

Setting up and tearing down existing IPsec tunnels does not affect the underlying IKE session or existing IPsec tunnels.

Use the following **show** command to view the current tunnel count per SPU: **show security ike tunnel-map**.

Use the **summary** option of the command to view the anchor points of each gateway: s**how security ike tunnel-map summary**.

Understanding VPN Support for Inserting Services Processing Cards

High-end SRX Series devices have a chassis-based distributed processor architecture. The flow processing power is shared and is based on the number of Services Processing Cards (SPCs). You can scale the processing power of the device by installing a new SPC.

In a high-end SRX Series chassis cluster, you can insert an SPC on a device without affecting or disrupting the traffic on existing VPN tunnels. However, existing tunnels cannot use the processing power of the Services Processing Units (SPUs) in the new SPC. A new SPU can anchor newly established site-to-site and dynamic tunnels, although newly configured tunnels are not guaranteed to be anchored on the new SPU.

Site-to-site tunnels are anchored on different SPUs based on a load-balancing algorithm. For a new site-to-site tunnel, the SPU with the smallest load is chosen as the anchor SPU. (The load corresponds to the number of site-to-site gateways or manual VPN tunnels anchored on an SPU.) If multiple SPUs have the same smallest load, then any of the SPUs can be chosen as the anchor SPU. A newly configured site-to-site tunnel is guaranteed to be anchored on the new SPU only if the loads of the existing SPUs are all greater than 0.

Dynamic tunnels are anchored on different SPUs based on a round-robin algorithm. Newly configured dynamic tunnels are not guaranteed to be anchored on the new SPC.

Copyright © 2016, Juniper Networks, Inc.

You can view the tunnel mapping on different SPUs using the **show security ike tunnel-map** command.

Related Documentation

- show security ike tunnel-map on page 7053
- Understanding Distributed VPNs in SRX Series Services Gateways on page 6067

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 80

Configuring Route-Based IPsec VPNs

- Configuring Route-Based VPNs on page 6073
- Configuring Hub-and-Spoke VPNs on page 6093
- Configuring VPNs for IKEv2 on page 6127
- Configuring Secure Tunnel Interface in a Virtual Router on page 6175
- Configuring Dual Stack Tunnels over an External Interface on page 6181
- Configuring Traffic Selectors in Route-Based VPNs on page 6195

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 254

Configuring Route-Based VPNs

- Understanding Route-Based IPsec VPNs on page 6073
- Example: Configuring a Route-Based VPN on page 6074

Understanding Route-Based IPsec VPNs

With route-based VPNs, you can configure dozens of security policies to regulate traffic flowing through a single VPN tunnel between two sites, and there is just one set of IKE and IPsec SAs at work. Unlike policy-based VPNs, for route-based VPNs, a policy refers to a destination address, not a VPN tunnel. When Junos OS looks up a route to find the interface to use to send traffic to the packet's destination address, it finds a route through a secure tunnel interface (st0.*x*). The tunnel interface is bound to a specific VPN tunnel, and the traffic is routed to the tunnel if the policy action is permit.

NOTE: **A secure tunnel (st0) interface supports only one IPv4 address and one IPv6 address at the same time. This applies to all route-based VPNs.**

Examples of where route-based VPNs can be used:

- There are overlapping subnets or IP addresses between the two LANs.
- A hub-and-spoke VPN topology is used in the network, and spoke-to-spoke traffic is required.
- Primary and backup VPNs are required.
- A dynamic routing protocol (for example, OSPF, RIP, or BGP) is running across the VPN.

NOTE: **Configuring RIP demand circuits over VPN interfaces is not supported.**

We recommend that you use route-based VPN when you want to configure VPN between multiple remote sites. Route-based VPN allows for routing between the spokes between multiple remote sites; it is easier to configure, monitor, and troubleshoot.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Hub-and-Spoke VPN on page 6094
- Example: Configuring a Policy-Based VPN on page 6222

Example: Configuring a Route-Based VPN

This example shows how to configure a route-based IPsec VPN to allow data to be securely transferred between a branch office and the corporate office.

- Requirements on page 6074
- Overview on page 6074
- Configuration on page 6078
- Verification on page 6087

Requirements

This example uses the following hardware:

- SRX240 device
- SSG140 device

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

In this example, you configure a route-based VPN for a branch office in Chicago, Illinois, because you want to conserve tunnel resources but still get granular restrictions on VPN traffic. Users in the Chicago office will use the VPN to connect to their corporate headquarters in Sunnyvale, California.

Figure 249 shows an example of a route-based VPN topology. In this topology, the SRX Series device is located in Sunnyvale, and an SSG Series device (or a third-party device) is located in Chicago.

Copyright © 2016, Juniper Networks, Inc.

Figure 249: Route-Based VPN Topology

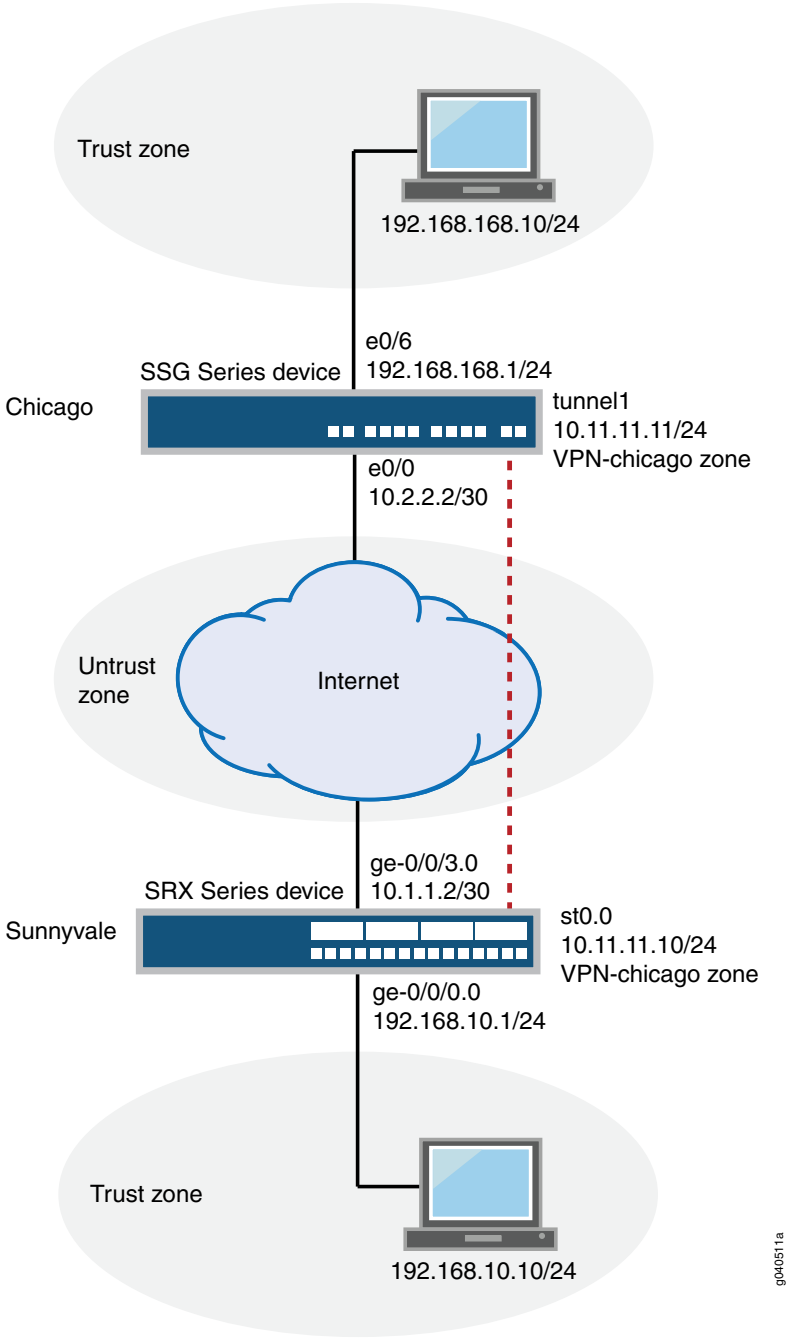

In this example, you configure interfaces, an IPv4 default route, security zones, and address books. Then you configure IKE Phase 1, IPsec Phase 2, security policy, and TCP-MSS parameters. See Table 504 through Table 508 for specific configuration parameters used in this example.

Copyright © 2016, Juniper Networks, Inc.

Table 504: Interface, Static Route, Security Zone, and Address Book Information

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/0.0	192.168.10.1/24
	ge-0/0/3.0	10.1.1.2/30
	st0.0 (tunnel interface)	10.11.11.10/24
Static routes	0.0.0.0/0 (default route)	The next hop is 10.1.1.2.
	192.168.168.0/24	The next hop is st0.0.
Security zones	trust	• All system services are allowed. • The ge-0/0/0.0 interface is bound to this zone.
	untrust	• IKE is the only allowed system service. • The ge-0/0/3.0 interface is bound to this zone.
	vpn-chicago	The st0.0 interface is bound to this zone.
Address book entries	sunnyvale	• This address is an entry in the address book **book1**, which is attached to a zone called **trust**. • The address for this address book entry is 192.168.10.0/24.
	chicago	• This address is an entry in the address book **book2**, which is attached to a zone called **vpn-chicago**. • The address for this address book entry is 192.168.168.0/24.

Table 505: IKE Phase 1 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ike-phase1-proposal	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: aes-128-cbc
Policy	ike-phase1-policy	• Mode: main • Proposal reference: ike-phase1-proposal • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gw-chicago	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/3.0 • Gateway address: 10.2.2.2

Copyright © 2016, Juniper Networks, Inc.

Table 506: IPsec Phase 2 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ipsec-phase2-proposal	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: aes-128-cbc
Policy	ipsec-phase2-policy	• Proposal reference: ipsec-phase2-proposal • PFS: Diffie-Hellman group2
VPN	ike-vpn-chicago	• IKE gateway reference: gw-chicago • IPsec policy reference: ipsec-phase2-policy • Bind to interface: st0.0

Table 507: Security Policy Configuration Parameters

Purpose	Name	Configuration Parameters

The security policy permits traffic from the trust zone to the vpn-chicago zone.	vpn-tr-chi	• Match criteria:
		  ◦ source-address sunnyvale
		  ◦ destination-address chicago
		  ◦ application any
		• Action: permit

The security policy permits traffic from the vpn-chicago zone to the trust zone.	vpn-chi-tr	• Match criteria:
		  ◦ source-address chicago
		  ◦ destination-address sunnyvale
		  ◦ application any
		• Action: permit

Table 508: TCP-MSS Configuration Parameters

Purpose	Configuration Parameters
TCP-MSS is negotiated as part of the TCP three-way handshake and limits the maximum size of a TCP segment to better fit the MTU limits on a network. For VPN traffic, the IPsec encapsulation overhead, along with the IP and frame overhead, can cause the resulting ESP packet to exceed the MTU of the physical interface, which causes fragmentation. Fragmentation increases bandwidth and device resources. NOTE: We recommend a value of 1350 as the starting point for most Ethernet-based networks with an MTU of 1500 or greater. You might need to experiment with different TCP-MSS values to obtain optimal performance. For example, you might need to change the value if any device in the path has a lower MTU, or if there is any additional overhead such as PPP or Frame Relay.	MSS value: 1350

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Configuring Interface, Static Route, Security Zone, and Address Book Information on page 6078
- Configuring IKE on page 6081
- Configuring IPsec on page 6083
- Configuring Security Policies on page 6084
- Configuring TCP-MSS on page 6086
- Configuring the SSG Series Device on page 6086

Configuring Interface, Static Route, Security Zone, and Address Book Information

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
set interfaces st0 unit 0 family inet address 10.11.11.10/24
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.2
set routing-options static route 192.168.168.0/24 next-hop st0.0
set security zones security-zone untrust interfaces ge-0/0/3.0
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone vpn-chicago interfaces st0.0
set security address-book book1 address sunnyvale 192.168.10.0/24
set security address-book book1 attach zone trust
set security address-book book2 address chicago 192.168.168.0/24
set security address-book book2 attach zone vpn-chicago
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interface, static route, security zone, and address book information:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
   user@host# set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
   user@host# set interfaces st0 unit 0 family inet address 10.11.11.10/24
   ```

2. Configure static route information.

   ```
   [edit]
   user@host# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.2
   user@host# set routing-options static route 192.168.168.0/24 next-hop st0.0
   ```

3. Configure the untrust security zone.

   ```
   [edit ]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# edit security zones security-zone untrust
```

4. Assign an interface to the security zone.

```
[edit security zones security-zone untrust]
user@host# set interfaces ge-0/0/3.0
```

5. Specify allowed system services for the security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@host# edit security zones security-zone trust
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/0.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

9. Configure an address book and attach a zone to it.

```
[edit security address-book book1]
user@host# set address sunnyvale 192.168.10.0/24
user@host# set attach zone trust
```

10. Configure the vpn-chicago security zone.

```
[edit]
user@host# edit security zones security-zone vpn-chicago
```

11. Assign an interface to the security zone.

```
[edit security zones security-zone vpn-chicago]
user@host# set interfaces st0.0
```

12. Configure another address book and attach a zone to it.

```
[edit security address-book book2]
user@host# set address chicago 192.168.168.0/24
user@host# set attach zone vpn-chicago
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, and **show security address-book** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.10.1/24;
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 10.1.1.2/30
    }
  }
}
st0{
  unit 0 {
    family inet {
      address 10.11.11.10/24
    }
  }
}
```

```
[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 10.1.1.2;
  route 192.168.168.0/24 next-hop st0.0;
}
```

```
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      ike;
    }
  }
  interfaces {
    ge-0/0/3.0;
  }
}
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
  }
  interfaces {
    ge-0/0/0.0;
  }
}
security-zone vpn-chicago {
  interfaces {
    st0.0;
  }
}
[edit]
user@host# show security address-book
book1 {
  address sunnyvale 192.168.10.0/24;
  attach {
    zone trust;
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
  book2 {
    address chicago 192.168.168.0/24;
    attach {
      zone vpn-chicago;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text "$ABC123"
set security ike gateway gw-chicago external-interface ge-0/0/3.0
set security ike gateway gw-chicago ike-policy ike-phase1-policy
set security ike gateway gw-chicago address 10.2.2.2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set encryption-algorithm aes-128-cbc
   ```

Copyright © 2016, Juniper Networks, Inc.

6. Create an IKE Phase 1 policy.

```
[edit security ike]
user@host# set policy ike-phase1-policy
```

7. Set the IKE Phase 1 policy mode.

```
[edit security ike policy ike-phase1-policy]
user@host# set mode main
```

8. Specify a reference to the IKE proposal.

```
[edit security ike policy ike-phase1-policy]
user@host# set proposals ike-phase1-proposal
```

9. Define the IKE Phase 1 policy authentication method.

```
[edit security ike policy ike-phase1-policy]
user@host# set pre-shared-key ascii-text “$ABC123”
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike]
user@host# set gateway gw-chicago external-interface ge-0/0/3.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw-chicago]
user@host# set ike-policy ike-phase1-policy
```

12. Define the IKE Phase 1 gateway address.

```
[edit security ike gateway gw-chicago]
user@host# set address 10.2.2.2
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security ike
proposal ike-phase1-proposal {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm sha1;
   encryption-algorithm aes-128-cbc;
}
policy ike-phase1-policy {
   mode main;
   proposals ike-phase1-proposal;
   pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway gw-chicago {
   ike-policy ike-phase1-policy;
   address 10.2.2.2;
   external-interface ge-0/0/3.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring IPsec

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy ipsec-phase2-policy proposals ipsec-phase2-proposal
set security ipsec policy ipsec-phase2-policy perfect-forward-secrecy keys group2
set security ipsec vpn ike-vpn-chicago ike gateway gw-chicago
set security ipsec vpn ike-vpn-chicago ike ipsec-policy ipsec-phase2-policy
set security ipsec vpn ike-vpn-chicago bind-interface st0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec-phase2-proposal
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set authentication-algorithm hmac-sha1-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set encryption-algorithm aes-128-cbc
   ```

5. Create the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@host# set policy ipsec-phase2-policy
   ```

6. Specify the IPsec Phase 2 proposal reference.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@host# set proposals ipsec-phase2-proposal
   ```

7. Specify IPsec Phase 2 PFS to use Diffie-Hellman group 2.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@host# set perfect-forward-secrecy keys group2
   ```

8. Specify the IKE gateway.

   ```
   [edit security ipsec]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set vpn ike-vpn-chicago ike gateway gw-chicago
```

9. Specify the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set vpn ike-vpn-chicago ike ipsec-policy ipsec-phase2-policy
```

10. Specify the interface to bind.

```
[edit security ipsec]
user@host# set vpn ike-vpn-chicago bind-interface st0.0
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security ipsec
proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
}
policy ipsec-phase2-policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec-phase2-proposal;
}
vpn ike-vpn-chicago {
    bind-interface st0.0;
    ike {
        gateway gw-chicago;
        ipsec-policy ipsec-phase2-policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    source-address sunnyvale
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    destination-address chicago
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    application any
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi then permit
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
    source-address chicago
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
  destination-address sunnyvale
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
  application any
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies:

1. Create the security policy to permit traffic from the trust zone to the vpn-chicago zone.

```
[edit security policies from-zone trust to-zone vpn-chicago]
user@host# set policy vpn-tr-chi match source-address sunnyvale
user@host# set policy vpn-tr-chi match destination-address chicago
user@host# set policy vpn-tr-chi match application any
user@host# set policy vpn-tr-chi then permit
```

2. Create the security policy to permit traffic from the vpn-chicago zone to the trust zone.

```
[edit security policies from-zone vpn-chicago to-zone trust]
user@host# set policy vpn-chi-tr match source-address chicago
user@host# set policy vpn-chi-tr match destination-address sunnyvale
user@host# set policy vpn-chi-tr match application any
user@host# set policy vpn-chi-tr then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
from-zone trust to-zone vpn-chicago {
  policy vpn-tr-vpn {
    match {
      source-address sunnyvale;
      destination-address chicago;
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone vpn-chicago to-zone trust {
  policy vpn-tr-vpn {
    match {
      source-address chicago;
      destination-address sunnyvale;
      application any;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
        then {
            permit;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring TCP-MSS

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security flow tcp-mss ipsec-vpn mss 1350
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure TCP-MSS information:

1. Configure TCP-MSS information.

```
[edit]
user@host# set security flow tcp-mss ipsec-vpn mss 1350
```

Results

From configuration mode, confirm your configuration by entering the **show security flow** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security flow
tcp-mss {
    ipsec-vpn {
        mss 1350;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the SSG Series Device

CLI Quick Configuration

For reference, the configuration for the SSG Series device is provided. For information about configuring SSG Series devices, see the *Concepts and Examples ScreenOS Reference Guide*, which is located at http://www.juniper.net/techpubs.

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set zone name vpn-chicago
```


Copyright © 2016, Juniper Networks, Inc.

```
set interface ethernet0/6 zone Trust
set interface ethernet0/0 zone Untrust
set interface tunnel.1 zone vpn-chicago
set interface ethernet0/6 ip 192.168.168.1/24
set interface ethernet0/6 route
set interface ethernet0/0 ip 10.2.2.2/30
set interface ethernet0/0 route
set interface tunnel.1 ip 10.11.11.11/24
set flow tcp-mss 1350
set address Trust “168-net” 192.168.168.0 255.255.255.0
set address vpn-chicago "10-net" 192.168.10.0 255.255.255.0
set ike gateway corp-ike address 10.1.1.2 Main outgoing-interface ethernet0/0 preshare
  $ABC123 sec-level standard
set vpn corp-vpn gateway corp-ike replay tunnel idletime 0 sec-level standard
set vpn corp-vpn monitor optimized rekey
set vpn corp-vpn bind interface tunnel.1
set policy from Trust to Untrust “ANY” “ANY” “ANY” nat src permit
set policy from Trust to vpn-chicago “192.168.168-net” “10-net” “ANY” permit
set policy from vpn-chicago to Trust “10.10.10-net” “168-net” “ANY” permit
set route 192.168.10.0/24 interface tunnel.1
set route 0.0.0.0/0 interface ethernet0/0 gateway 10.2.2.1
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the IKE Phase 1 Status on page 6087
- Verifying the IPsec Phase 2 Status on page 6089
- Reviewing Statistics and Errors for an IPsec Security Association on page 6090
- Testing Traffic Flow Across the VPN on page 6091

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action

NOTE: Before starting the verification process, you need to send traffic from a host in the 192.168.10/24 network to a host in the 192.168.168/24 network. For route-based VPNs, traffic can be initiated by the SRX Series device through the tunnel. We recommend that when testing IPsec tunnels, test traffic be sent from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate a ping from 192.168.10.10 to 192.168.168.10.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode
1       10.2.2.2         UP     744a594d957dd513  1e1307db82f58387  Main
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host> show security ike security-associations index 1 detail
IKE peer 10.2.2.2, Index 1,
  Role: Responder, State: UP
  Initiator cookie: 744a594d957dd513, Responder cookie: 1e1307db82f58387
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.1.1.2:500, Remote: 10.2.2.2:500
  Lifetime: Expires in 28570 seconds
  Algorithms:
   Authentication        : sha1
   Encryption            : aes-cbc (128 bits)
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input bytes    :                 852
   Output bytes   :                 940
   Input packets  :                   5
   Output packets :                   5
  Flags: Caller notification sent
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote Address—Verify that the remote IP address is correct.
- State
  - UP—The Phase 1 SA has been established.
  - DOWN—There was a problem establishing the Phase 1 SA.
- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations index 1 detail** command lists additional information about the security association with an index number of 1:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

 Copyright © 2016, Juniper Networks, Inc.

NOTE: Troubleshooting is best performed on the peer using the responder role.

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying the IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
  total configured sa: 2
  ID    Gateway          Port  Algorithm          SPI      Life:sec/kb  Mon vsys

  <16384 10.2.2.2         500   ESP:aes-128/sha1   76d64d1d 3363/ unlim   -   0
  >16384 10.2.2.2         500   ESP:aes-128/sha1   a1024ee2 3363/ unlim   -   0

user@host> show security ipsec security-associations index 16384 detail
  Virtual-system: Root
  Local Gateway: 10.1.1.2, Remote Gateway: 10.2.2.2
  Local Identity: ipv4_subnet(any:0,[0..7]=192.168.10.0/24)
  Remote Identity: ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
    DF-bit: clear

    Direction: inbound, SPI: 1993755933, AUX-SPI: 0
    Hard lifetime: Expires in 3352 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 2775 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)

    Anti-replay service: enabled, Replay window size: 32

    Direction: outbound, SPI: 2701283042, AUX-SPI: 0
    Hard lifetime: Expires in 3352 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 2775 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc
 (128 bits)
    Anti-replay service: enabled, Replay window size: 32
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

Copyright © 2016, Juniper Networks, Inc.

- The ID number is 16384. Use this value with the **show security ipsec security-associations index** command to get more information about this particular SA.
- There is one IPsec SA pair using port 500, which indicates that no NAT-traversal is implemented. (NAT-traversal uses port 4500 or another random high-number port.)
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3363/ unlim value indicates that the Phase 2 lifetime expires in 3363 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

The output from the **show security ipsec security-associations index 16384 detail** command lists the following information:

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, are correct for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be specified. For some third-party vendors, the proxy ID must be manually entered to match.
- Another common reason for Phase 2 failure is not specifying the ST interface binding. If IPsec cannot complete, check the kmd log or set traceoptions.

Reviewing Statistics and Errors for an IPsec Security Association

Purpose Review ESP and authentication header counters and errors for an IPsec security association.

Action From operational mode, enter the **show security ipsec statistics index** *index_number* command, using the index number of the VPN for which you want to see statistics.

```
user@host> show security ipsec statistics index 16384
ESP Statistics:
  Encrypted bytes:              920
  Decrypted bytes:             6208
  Encrypted packets:              5
  Decrypted packets:             87
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

You can also use the **show security ipsec statistics** command to review statistics and errors for all SAs.

To clear all IPsec statistics, use the **clear security ipsec statistics** command.

Meaning If you see packet loss issues across a VPN, you can run the **show security ipsec statistics** or **show security ipsec statistics detail** command several times to confirm that the encrypted and decrypted packet counters are incrementing. You should also check whether the other error counters are incrementing.

Testing Traffic Flow Across the VPN

Purpose Verify the traffic flow across the VPN.

Action You can use the **ping** command from the SRX Series device to test traffic flow to a remote host PC. Make sure that you specify the source interface so that the route lookup is correct and the appropriate security zones are referenced during policy lookup.

From operational mode, enter the **ping** command.

```
ssg-> ping 192.168.168.10 interface ge-0/0/0 count 5
PING 192.168.168.10 (192.168.168.10): 56 data bytes
64 bytes from 192.168.168.10: icmp_seq=0 ttl=127 time=8.287 ms
64 bytes from 192.168.168.10: icmp_seq=1 ttl=127 time=4.119 ms
64 bytes from 192.168.168.10: icmp_seq=2 ttl=127 time=5.399 ms
64 bytes from 192.168.168.10: icmp_seq=3 ttl=127 time=4.361 ms
64 bytes from 192.168.168.10: icmp_seq=4 ttl=127 time=5.137 ms

--- 192.168.168.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.119/5.461/8.287/1.490 ms
```

You can also use the **ping** command from the SSG Series device.

```
user@host> ping 192.168.10.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.10.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/4/5 ms
```

Meaning If the **ping** command fails from the SRX Series or SSG Series device, there might be a problem with the routing, security policies, end host, or encryption and decryption of ESP packets.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Hub-and-Spoke VPN on page 6094
- Example: Configuring a Policy-Based VPN on page 6222

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 255

Configuring Hub-and-Spoke VPNs

- Understanding Hub-and-Spoke VPNs on page 6093
- Example: Configuring a Hub-and-Spoke VPN on page 6094

Understanding Hub-and-Spoke VPNs

If you create two VPN tunnels that terminate at a device, you can set up a pair of routes so that the device directs traffic exiting one tunnel to the other tunnel. You also need to create a policy to permit the traffic to pass from one tunnel to the other. Such an arrangement is known as *hub-and-spoke VPN*. (See Figure 250.)

You can also configure multiple VPNs and route traffic between any two tunnels.

NOTE: SRX Series devices support only the route-based hub-and-spoke feature.

Figure 250: Multiple Tunnels in a Hub-and-Spoke VPN Configuration

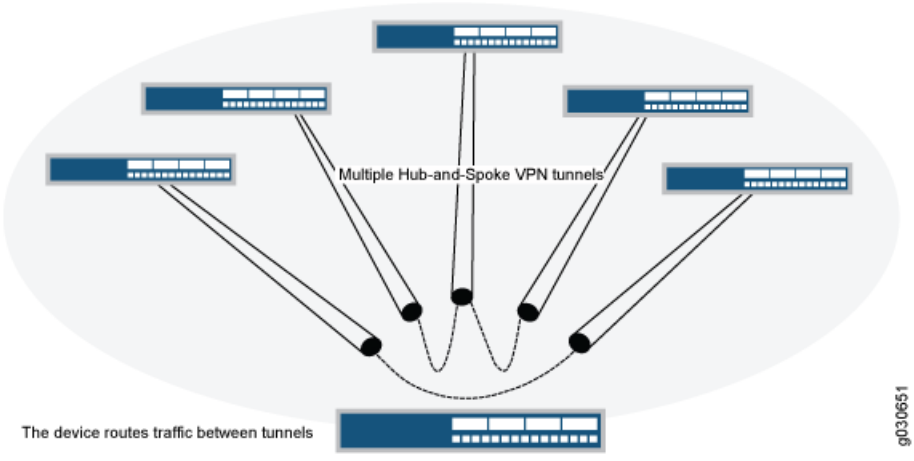

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Hub-and-Spoke VPN on page 6094

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring a Hub-and-Spoke VPN

This example shows how to configure a hub-and-spoke IPsec VPN for an enterprise-class deployment.

- Requirements on page 6094
- Overview on page 6094
- Configuration on page 6100
- Verification on page 6120

Requirements

This example uses the following hardware:

- SRX240 device
- SRX5800 device
- SSG140 device

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

This example describes how to configure a hub-and-spoke VPN typically found in branch deployments. The hub is the corporate office, and there are two spokes—a branch office in Sunnyvale, California, and a branch office in Westford, Massachusetts. Users in the branch offices will use the VPN to securely transfer data with the corporate office.

Figure 251 shows an example of a hub-and-spoke VPN topology. In this topology, an SRX5800 device is located at the corporate office. An SRX240 device is located at the Westford branch, and an SSG140 device is located at the Sunnyvale branch.

Copyright © 2016, Juniper Networks, Inc.

Figure 251: Hub-and-Spoke VPN Topology

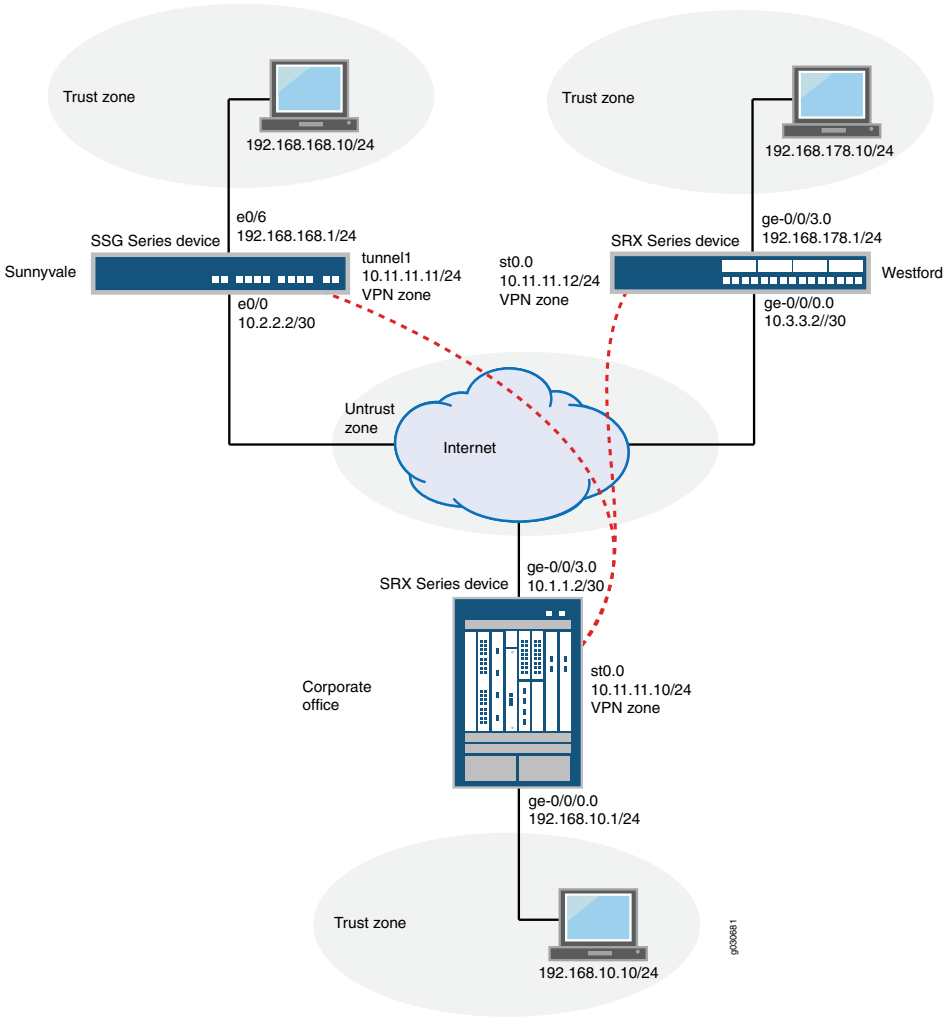

In this example, you configure the corporate office hub, the Westford spoke, and the Sunnyvale spoke. First you configure interfaces, IPv4 static and default routes, security zones, and address books. Then you configure IKE Phase 1 and IPsec Phase 2 parameters, and bind the st0.0 interface to the IPsec VPN. On the hub, you configure st0.0 for multipoint and add a static NHTB table entry for the Sunnyvale spoke. Finally, you configure security policy and TCP-MSS parameters. See Table 509 through Table 513 for specific configuration parameters used in this example.

Table 509: Interface, Security Zone, and Address Book Information

Hub or Spoke	Feature	Name	Configuration Parameters
Hub	Interfaces	ge-0/0/0.0	192.168.10.1/24
		ge-0/0/3.0	10.1.1.2/30

Copyright © 2016, Juniper Networks, Inc.

Table 509: Interface, Security Zone, and Address Book Information *(continued)*

Hub or Spoke	Feature	Name	Configuration Parameters
		st0	10.11.11.10/24
Spoke	Interfaces	ge-0/0/0.0	10.3.3.2/30
		ge-0/0/3.0	192.168.178.1/24
		st0	10.11.11.12/24
Hub	Security zones	trust	• All system services are allowed. • The ge-0/0/0.0 interface is bound to this zone.
		untrust	• IKE is the only allowed system service. • The ge-0/0/3.0 interface is bound to this zone.
		vpn	The st0.0 interface is bound to this zone.
Spoke	Security zones	trust	• All system services are allowed. • The ge-0/0/3.0 interface is bound to this zone.
		untrust	• IKE is the only allowed system service. • The ge-0/0/0.0 interface is bound to this zone.
		vpn	The st0.0 interface is bound to this zone.
Hub	Address book entries	local-net	• This address is for the trust zone's address book. • The address for this address book entry is 192.168.10.0/24.
		sunnyvale-net	• This address book is for the vpn zone's address book. • The address for this address book entry is 192.168.168.0/24.
		westford-net	• This address is for the vpn zone's address book. • The address for this address book entry is 192.168.178.0/24.

Copyright © 2016, Juniper Networks, Inc.

Table 509: Interface, Security Zone, and Address Book Information *(continued)*

Hub or Spoke	Feature	Name	Configuration Parameters
Spoke	Address book entries	local-net	• This address is for the trust zone's address book. • The address for this address book entry is 192.168.168.178.0/24.
		corp-net	• This address is for the vpn zone's address book. • The address for this address book entry is 192.168.10.0/24.
		sunnyvale-net	• This address is for the vpn zone's address book. • The address for this address book entry is 192.168.168.0/24.

Table 510: IKE Phase 1 Configuration Parameters

Hub or Spoke	Feature	Name	Configuration Parameters
Hub	Proposal	ike-phase1-proposal	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: aes-128-cbc
	Policy	ike-phase1-policy	• Mode: main • Proposal reference: ike-phase1-proposal • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
	Gateway	gw-westford	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/3.0 • Gateway address: 10.3.3.2
		gw-sunnyvale	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/3.0 • Gateway address: 10.2.2.2

Copyright © 2016, Juniper Networks, Inc.

Table 510: IKE Phase 1 Configuration Parameters *(continued)*

Hub or Spoke	Feature	Name	Configuration Parameters
Spoke	Proposal	ike-phase1-proposal	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: aes-128-cbc
	Policy	ike-phase1-policy	• Mode: main • Proposal reference: ike-phase1-proposal • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
	Gateway	gw-corporate	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/0.0 • Gateway address: 10.1.1.2

Table 511: IPsec Phase 2 Configuration Parameters

Hub or Spoke	Feature	Name	Configuration Parameters
Hub	Proposal	ipsec-phase2-proposal	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: aes-128-cbc
	Policy	ipsec-phase2-policy	• Proposal reference: ipsec-phase2-proposal • PFS: Diffie-Hellman group2
	VPN	vpn-sunnyvale	• IKE gateway reference: gw-sunnyvale • IPsec policy reference: ipsec-phase2-policy • Bind to interface: st0.0
		vpn-westford	• IKE gateway reference: gw-westford • IPsec policy reference: ipsec-phase2-policy • Bind to interface: st0.0
Spoke	Proposal	ipsec-phase2-proposal	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: aes-128-cbc
	Policy	ipsec-phase2-policy	• Proposal reference: ipsec-phase2-proposal • PFS: Diffie-Hellman group2

Copyright © 2016, Juniper Networks, Inc.

Table 511: IPsec Phase 2 Configuration Parameters *(continued)*

Hub or Spoke	Feature	Name	Configuration Parameters
	VPN	vpn-corporate	• IKE gateway reference: gw-corporate • IPsec policy reference: ipsec-phase2-policy • Bind to interface: st0.0

Table 512: Security Policy Configuration Parameters

Hub or Spoke	Purpose	Name	Configuration Parameters

Hub	The security policy permits traffic from the trust zone to the vpn zone.	local-to-spokes	• Match criteria:
			  • source-address local-net
			  • destination-address sunnyvale-net
			  • destination-address westford-net
			  • application any

	The security policy permits traffic from the vpn zone to the trust zone.	spokes-to-local	Match criteria: • source-address sunnyvale-net • source-address westford-net • destination-address local-net • application any
	The security policy permits intrazone traffic.	spoke-to-spoke	Match criteria: • source-address any • destination-address any • application any

Spoke	The security policy permits traffic from the trust zone to the vpn zone.	to-corp	• Match criteria:
			  • source-address local-net
			  • destination-address corp-net
			  • destination-address sunnyvale-net
			  • application any

	The security policy permits traffic from the vpn zone to the trust zone.	from-corp	Match criteria: • source-address corp-net • source-address sunnyvale-net • destination-address local-net • application any

Copyright © 2016, Juniper Networks, Inc.

Table 512: Security Policy Configuration Parameters *(continued)*

Hub or Spoke	Purpose	Name	Configuration Parameters
	The security policy permits traffic from the untrust zone to the trust zone.	permit-any	Match criteria: • source-address any • source-destination any • application any • Permit action: source-nat interface By specifying **source-nat interface**, the SRX Series device translates the source IP address and port for outgoing traffic, using the IP address of the egress interface as the source IP address and a random high-number port for the source port.

Table 513: TCP-MSS Configuration Parameters

Purpose	Configuration Parameters
TCC-MSS is negotiated as part of the TCP three-way handshake and limits the maximum size of a TCP segment to better fit the MTU limits on a network. For VPN traffic, the IPsec encapsulation overhead, along with the IP and frame overhead, can cause the resulting ESP packet to exceed the MTU of the physical interface, which causes fragmentation. Fragmentation results in increased use of bandwidth and device resources. NOTE: The value of 1350 is a recommended starting point for most Ethernet-based networks with an MTU of 1500 or greater. You might need to experiment with different TCP-MSS values to obtain optimal performance. For example, you might need to change the value if any device in the path has a lower MTU, or if there is any additional overhead such as PPP or Frame Relay.	MSS value: 1350

Configuration

- Configuring Basic Network, Security Zone, and Address Book Information for the Hub on page 6101
- Configuring IKE for the Hub on page 6104
- Configuring IPsec for the Hub on page 6106
- Configuring Security Policies for the Hub on page 6108
- Configuring TCP-MSS for the Hub on page 6110
- Configuring Basic Network, Security Zone, and Address Book Information for the Westford Spoke on page 6111
- Configuring IKE for the Westford Spoke on page 6114
- Configuring IPsec for the Westford Spoke on page 6115
- Configuring Security Policies for the Westford Spoke on page 6117
- Configuring TCP-MSS for the Westford Spoke on page 6119
- Configuring the Sunnyvale Spoke on page 6119

Copyright © 2016, Juniper Networks, Inc.

Configuring Basic Network, Security Zone, and Address Book Information for the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
set interfaces st0 unit 0 family inet address 10.11.11.10/24
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
set routing-options static route 192.168.168.0/24 next-hop 10.11.11.11
set routing-options static route 192.168.178.0/24 next-hop 10.11.11.12
set security zones security-zone untrust interfaces ge-0/0/3.0
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone vpn interfaces st0.0
set security address-book book1 address local-net 192.168.10.0/24
set security address-book book1 attach zone trust
set security address-book book2 address sunnyvale-net 192.168.168.0/24
set security address-book book2 address westford-net 192.168.178.0/24
set security address-book book2 attach zone vpn
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure basic network, security zone, and address book information for the hub:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@hub# set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
   user@hub# set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
   user@hub# set interfaces st0 unit 0 family inet address 10.11.11.10/24
   ```

2. Configure static route information.

   ```
   [edit]
   user@hub# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
   user@hub# set routing-options static route 192.168.168.0/24 next-hop 10.11.11.11
   user@hub# set routing-options static route 192.168.178.0/24 next-hop 10.11.11.12
   ```

3. Configure the untrust security zone.

   ```
   [edit ]
   user@hub# set security zones security-zone untrust
   ```

4. Assign an interface to the untrust security zone.

   ```
   [edit security zones security-zone untrust]
   user@hub# set interfaces ge-0/0/3.0
   ```

5. Specify allowed system services for the untrust security zone.

   ```
   [edit security zones security-zone untrust]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@hub# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@hub# edit security zones security-zone trust
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@hub# set interfaces ge-0/0/0.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@hub# set host-inbound-traffic system-services all
```

9. Create an address book and attach a zone to it.

```
[edit security address-book book1]
user@hub# set address local-net 10.10.10.0/24
user@hub# set attach zone trust
```

10. Configure the vpn security zone.

```
[edit]
user@hub# edit security zones security-zone vpn
```

11. Assign an interface to the vpn security zone.

```
[edit security zones security-zone vpn]
user@hub# set interfaces st0.0
```

12. Create another address book and attach a zone to it.

```
[edit security address-book book2]
user@hub# set address sunnyvale-net 192.168.168.0/24
user@hub# set address westford-net 192.168.178.0/24
user@hub# set attach zone vpn
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, and **show security address-book** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@hub# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 192.168.10.1/24;
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.1.2/30
        }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
st0{
    unit 0 {
        family inet {
            address 10.11.11.10/24
        }
    }
}

[edit]
user@hub# show routing-options
static {
    route 0.0.0.0/0 next-hop 10.1.1.1;
    route 192.168.168.0/24 next-hop 10.11.11.11;
    route 192.168.178.0/24 next-hop 10.11.11.12;
}

[edit]
user@hub# show security zones
security-zone untrust {
    host-inbound-traffic {
        system-services {
            ike;
        }
    }
    interfaces {
        ge-0/0/3.0;
    }
}
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
    }
    interfaces {
        ge-0/0/0.0;
    }
}
security-zone vpn {
    host-inbound-traffic {
    }
    interfaces {
        st0.0;
    }
}
[edit]
user@hub# show security address-book
book1 {
    address local-net 10.10.10.0/24;
    attach {
        zone trust;
    }
}
    book2 {
        address sunnyvale-net 192.168.168.0/24;
        address westford-net 192.168.178.0/24;
```

```
    attach {
      zone vpn;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text $ABC123
set security ike gateway gw-westford external-interface ge-0/0/3.0
set security ike gateway gw-westford ike-policy ike-phase1-policy
set security ike gateway gw-westford address 10.3.3.2
set security ike gateway gw-sunnyvale external-interface ge-0/0/3.0
set security ike gateway gw-sunnyvale ike-policy ike-phase1-policy
set security ike gateway gw-sunnyvale address 10.2.2.2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE for the hub:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@hub# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@hub# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@hub# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@hub# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@hub# set encryption-algorithm aes-128-cbc
   ```


Copyright © 2016, Juniper Networks, Inc.

6. Create an IKE Phase 1 policy.

```
[edit security ike]
user@hub# set policy ike-phase1-policy
```

7. Set the IKE Phase 1 policy mode.

```
[edit security ike policy ike-phase1-policy]
user@hub# set mode main
```

8. Specify a reference to the IKE proposal.

```
[edit security ike policy ike-phase1-policy]
user@hub# set proposals ike-phase1-proposal
```

9. Define the IKE Phase 1 policy authentication method.

```
[edit security ike policy ike-phase1-policy]
user@hub# set pre-shared-key ascii-text $ABC123
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike]
user@hub# set gateway gw-westford external-interface ge-0/0/3.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike]
user@hub# set gateway gw-westford ike-policy ike-phase1-policy
```

12. Define the IKE Phase 1 gateway address.

```
[edit security ike]
user@hub# set gateway gw-westford address 10.3.3.2
```

13. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike]
user@hub# set gateway gw-sunnyvale external-interface ge-0/0/3.0
```

14. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway]
user@hub# set gateway gw-sunnyvale ike-policy ike-phase1-policy
```

15. Define the IKE Phase 1 gateway address.

```
[edit security ike gateway]
user@hub# set gateway gw-sunnyvale address 10.2.2.2
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@hub# show security ike
proposal ike-phase1-proposal {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm sha1;
   encryption-algorithm aes-128-cbc;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
policy ike-phase1-policy {
    mode main;
    proposals ike-phase1-proposal;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway gw-sunnyvale {
    ike-policy ike-phase1-policy;
    address 10.2.2.2;
    external-interface ge-0/0/3.0;
}
gateway gw-westford {
    ike-policy ike-phase1-policy;
    address 10.3.3.2;
    external-interface ge-0/0/3.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec for the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy ipsec-phase2-policy proposals ipsec-phase2-proposal
set security ipsec policy ipsec-phase2-policy perfect-forward-secrecy keys group2
set security ipsec vpn vpn-westford ike gateway gw-westford
set security ipsec vpn vpn-westford ike ipsec-policy ipsec-phase2-policy
set security ipsec vpn vpn-westford bind-interface st0.0
set security ipsec vpn vpn-sunnyvale ike gateway gw-sunnyvale
set security ipsec vpn vpn-sunnyvale ike ipsec-policy ipsec-phase2-policy
set security ipsec vpn vpn-sunnyvale bind-interface st0.0
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet next-hop-tunnel 10.11.11.11 ipsec-vpn vpn-sunnyvale
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec for the hub:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@hub# set security ipsec proposal ipsec-phase2-proposal
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@hub# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ipsec proposal ipsec-phase2-proposal]
user@hub# set authentication-algorithm hmac-sha1-96
```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

```
[edit security ipsec proposal ipsec-phase2-proposal]
user@hub# set encryption-algorithm aes-128-cbc
```

5. Create the IPsec Phase 2 policy.

```
[edit security ipsec]
user@hub# set policy ipsec-phase2-policy
```

6. Specify the IPsec Phase 2 proposal reference.

```
[edit security ipsec policy ipsec-phase2-policy]
user@hub# set proposals ipsec-phase2-proposal
```

7. Specify IPsec Phase 2 PFS to use Diffie-Hellman group 2.

```
[edit security ipsec policy ipsec-phase2-policy]
user@host# set perfect-forward-secrecy keys group2
```

8. Specify the IKE gateways.

```
[edit security ipsec]
user@hub# set vpn vpn-westford ike gateway gw-westford
user@hub# set vpn vpn-sunnyvale ike gateway gw-sunnyvale
```

9. Specify the IPsec Phase 2 policies.

```
[edit security ipsec]
user@hub# set vpn vpn-westford ike ipsec-policy ipsec-phase2-policy
user@hub# set vpn vpn-sunnyvale ike ipsec-policy ipsec-phase2-policy
```

10. Specify the interface to bind.

```
[edit security ipsec]
user@hub# set vpn vpn-westford bind-interface st0.0
user@hub# set vpn vpn-sunnyvale bind-interface st0.0
```

11. Configure the st0 interface as multipoint.

```
[edit]
user@hub# set interfaces st0 unit 0 multipoint
```

12. Add static NHTB table entries for the Sunnyvale and Westford offices.

```
[edit]
user@hub# set interfaces st0 unit 0 family inet next-hop-tunnel 10.11.11.11 ipsec-vpn
  vpn-sunnyvale
user@hub# set interfaces st0 unit 0 family inet next-hop-tunnel 10.11.11.12 ipsec-vpn
  vpn-westford
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@hub# show security ipsec
proposal ipsec-phase2-proposal {
  protocol esp;
```

Copyright © 2016, Juniper Networks, Inc.

```
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
}
policy ipsec-phase2-policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec-phase2-proposal;
}
vpn vpn-sunnyvale {
    bind-interface st0.0;
    ike {
        gateway gw-sunnyvale;
        ipsec-policy ipsec-phase2-policy;
    }
}
vpn vpn-westford {
    bind-interface st0.0;
    ike {
        gateway gw-westford;
        ipsec-policy ipsec-phase2-policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies for the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone vpn policy local-to-spokes match
    source-address local-net
set security policies from-zone trust to-zone vpn policy local-to-spokes match
    destination-address sunnyvale-net
set security policies from-zone trust to-zone vpn policy local-to-spokes match
    destination-address westford-net
set security policies from-zone trust to-zone vpn policy local-to-spokes match application
    any
set security policies from-zone trust to-zone vpn policy local-to-spokes then permit
set security policies from-zone vpn to-zone trust policy spokes-to-local match
    source-address sunnyvale-net
set security policies from-zone vpn to-zone trust policy spokes-to-local match
    source-address westford-net
set security policies from-zone vpn to-zone trust policy spokes-to-local match
    destination-address local-net
set security policies from-zone vpn to-zone trust policy spokes-to-local match application
    any
set security policies from-zone vpn to-zone trust policy spokes-to-local then permit
set security policies from-zone vpn to-zone vpn policy spoke-to-spoke match
    source-address any
set security policies from-zone vpn to-zone vpn policy spoke-to-spoke match
    destination-address any
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies from-zone vpn to-zone vpn policy spoke-to-spoke match application
  any
set security policies from-zone vpn to-zone vpn policy spoke-to-spoke then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies for the hub:

1. Create the security policy to permit traffic from the trust zone to the vpn zone.

```
[edit security policies from-zone trust to-zone vpn]
user@hub# set policy local-to-spokes match source-address local-net
user@hub# set policy local-to-spokes match destination-address sunnyvale-net
user@hub# set policy local-to-spokes match destination-address westford-net
user@hub# set policy local-to-spokes match application any
user@hub# set policy local-to-spokes then permit
```

2. Create the security policy to permit traffic from the vpn zone to the trust zone.

```
[edit security policies from-zone vpn to-zone trust]
user@hub# set policy spokes-to-local match source-address sunnyvale-net
user@hub# set policy spokes-to-local match source-address westford-net
user@hub# set policy spokes-to-local match destination-address local-net
user@hub# set policy spokes-to-local match application any
user@hub# set policy spokes-to-local then permit
```

3. Create the security policy to permit intrazone traffic.

```
[edit security policies from-zone vpn to-zone vpn]
user@hub# set policy spoke-to-spoke match source-address any
user@hub# set policy spoke-to-spoke match destination-address any
user@hub# set policy spoke-to-spoke match application any
user@hub# set policy spoke-to-spoke then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@hub# show security policies
from-zone trust to-zone vpn {
  policy local-to-spokes {
    match {
      source-address local-net;
      destination-address [ sunnyvale-net westford-net ];
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone vpn to-zone trust {
  policy spokes-to-local {
```

Copyright © 2016, Juniper Networks, Inc.

```
        match {
            source-address [ sunnyvale-net westford-net ];
            destination-address local-net;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone vpn to-zone vpn {
    policy spoke-to-spoke {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring TCP-MSS for the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security flow tcp-mss ipsec-vpn mss 1350
```

Step-by-Step Procedure

To configure TCP-MSS information for the hub:

1. Configure TCP-MSS information.

```
[edit]
user@hub# set security flow tcp-mss ipsec-vpn mss 1350
```

Results

From configuration mode, confirm your configuration by entering the **show security flow** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@hub# show security flow
tcp-mss {
    ipsec-vpn {
        mss 1350;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring Basic Network, Security Zone, and Address Book Information for the Westford Spoke

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 10.3.3.2/30
set interfaces ge-0/0/3 unit 0 family inet address 192.168.178.1/24
set interfaces st0 unit 0 family inet address 10.11.11.12/24
set routing-options static route 0.0.0.0/0 next-hop 10.3.3.1
set routing-options static route 10.10.10.0/24 next-hop 10.11.11.10
set routing-options static route 192.168.168.0/24 next-hop 10.11.11.10
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone trust interfaces ge-0/0/3.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone vpn interfaces st0.0
set security address-book book1 address local-net 192.168.178.0/24
set security address-book book1 attach zone trust
set security address-book book2 address corp-net 10.10.10.0/24
set security address-book book2 address sunnyvale-net 192.168.168.0/24
set security address-book book2 attach zone vpn
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure basic network, security zone, and address book information for the Westford spoke:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@spoke# set interfaces ge-0/0/0 unit 0 family inet address 10.3.3.2/30
   user@spoke# set interfaces ge-0/0/3 unit 0 family inet address 192.168.178.1/24
   user@spoke# set interfaces st0 unit 0 family inet address 10.11.11.12/24
   ```

2. Configure static route information.

   ```
   [edit]
   user@spoke# set routing-options static route 0.0.0.0/0 next-hop 10.3.3.1
   user@spoke# set routing-options static route 10.10.10.0/24 next-hop 10.11.11.10
   user@spoke# set routing-options static route 192.168.168.0/24 next-hop 10.11.11.10
   ```

3. Configure the untrust security zone.

   ```
   [edit]
   user@spoke# set security zones security-zone untrust
   ```

4. Assign an interface to the security zone.

   ```
   [edit security zones security-zone untrust]
   user@spoke# set interfaces ge-0/0/0.0
   ```

5. Specify allowed system services for the untrust security zone.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security zones security-zone untrust]
user@spoke# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@spoke# edit security zones security-zone trust
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@spoke# set interfaces ge-0/0/3.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@spoke# set host-inbound-traffic system-services all
```

9. Configure the vpn security zone.

```
[edit]
user@spoke# edit security zones security-zone vpn
```

10. Assign an interface to the vpn security zone.

```
[edit security zones security-zone vpn]
user@spoke# set interfaces st0.0
```

11. Create an address book and attach a zone to it.

```
[edit security address-book book1]
user@spoke# set address local-net 192.168.178.0/24
user@spoke# set attach zone trust
```

12. Create another address book and attach a zone to it.

```
[edit security address-book book2]
user@spoke# set address corp-net 10.10.10.0/24
user@spoke# set address sunnyvale-net 192.168.168.0/24
user@spoke# set attach zone vpn
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, and **show security address-book** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@spoke# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.3.3.2/30;
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.178.1/24;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  st0 {
    unit 0 {
      family inet {
        address 10.11.11.10/24;
      }
    }
  }

[edit]
user@spoke# show routing-options
static {
  route 0.0.0.0/0 next-hop 10.3.3.1;
  route 192.168.168.0/24 next-hop 10.11.11.10;
  route 10.10.10.0/24 next-hop 10.11.11.10;
}

[edit]
user@spoke# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      ike;
    }
  }
  interfaces {
    ge-0/0/0.0;
  }
}
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
  }
  interfaces {
    ge-0/0/3.0;
  }
}
security-zone vpn {
  interfaces {
    st0.0;
  }
}
[edit]
user@spoke# show security address-book
book1 {
  address corp-net 10.10.10.0/24;
  attach {
    zone trust;
  }
}
  book2 {
    address local-net 192.168.178.0/24;
    address sunnyvale-net 192.168.168.0/24;
    attach {
```

```
        zone vpn;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Westford Spoke

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text $ABC123
set security ike gateway gw-corporate external-interface ge-0/0/0.0
set security ike gateway gw-corporate ike-policy ike-phase1-policy
set security ike gateway gw-corporate address 10.1.1.2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE for the Westford spoke:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@spoke# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@spoke# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@spoke# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@spoke# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@spoke# set encryption-algorithm aes-128-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike]
   user@spoke# set policy ike-phase1-policy
   ```


Copyright © 2016, Juniper Networks, Inc.

7. Set the IKE Phase 1 policy mode.

```
[edit security ike policy ike-phase1-policy]
user@spoke# set mode main
```

8. Specify a reference to the IKE proposal.

```
[edit security ike policy ike-phase1-policy]
user@spoke# set proposals ike-phase1-proposal
```

9. Define the IKE Phase 1 policy authentication method.

```
[edit security ike policy ike-phase1-policy]
user@spoke# set pre-shared-key ascii-text $ABC123
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike]
user@spoke# set gateway gw-corporate external-interface ge-0/0/0.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike]
user@spoke# set gateway gw-corporate ike-policy ike-phase1-policy
```

12. Define the IKE Phase 1 gateway address.

```
[edit security ike]
user@spoke# set gateway gw-corporate address 10.1.1.2
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@spoke# show security ike
proposal ike-phase1-proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
policy ike-phase1-policy {
    mode main;
    proposals ike-phase1-proposal;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway gw-corporate {
    ike-policy ike-phase1-policy;
    address 10.1.1.2;
    external-interface ge-0/0/0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec for the Westford Spoke

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy ipsec-phase2-policy proposals ipsec-phase2-proposal
set security ipsec policy ipsec-phase2-policy perfect-forward-secrecy keys group2
set security ipsec vpn vpn-corporate ike gateway gw-corporate
set security ipsec vpn vpn-corporate ike ipsec-policy ipsec-phase2-policy
set security ipsec vpn vpn-corporate bind-interface st0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec for the Westford spoke:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@spoke# set security ipsec proposal ipsec-phase2-proposal
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@spoke# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@spoke# set authentication-algorithm hmac-sha1-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@spoke# set encryption-algorithm aes-128-cbc
   ```

5. Create the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@spoke# set policy ipsec-phase2-policy
   ```

6. Specify the IPsec Phase 2 proposal reference.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@spoke# set proposals ipsec-phase2-proposal
   ```

7. Specify IPsec Phase 2 PFS to use Diffie-Hellman group 2.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@host# set perfect-forward-secrecy keys group2
   ```

8. Specify the IKE gateway.

   ```
   [edit security ipsec]
   user@spoke# set vpn vpn-corporate ike gateway gw-corporate
   ```

9. Specify the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@spoke# set vpn vpn-corporate ike ipsec-policy ipsec-phase2-policy
   ```


Copyright © 2016, Juniper Networks, Inc.

10. Specify the interface to bind.

```
[edit security ipsec]
user@spoke# set vpn vpn-corporate bind-interface st0.0
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@spoke# show security ipsec
proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
}
policy ipsec-phase2-policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec-phase2-proposal;
}
vpn vpn-corporate {
    bind-interface st0.0;
    ike {
        gateway gw-corporate;
        ipsec-policy ipsec-phase2-policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies for the Westford Spoke

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone vpn policy to-corporate match source-address
    local-net
set security policies from-zone trust to-zone vpn policy to-corporate match
    destination-address corp-net
set security policies from-zone trust to-zone vpn policy to-corporate match
    destination-address sunnyvale-net
set security policies from-zone trust to-zone vpn policy to-corporate application any
set security policies from-zone trust to-zone vpn policy to-corporate then permit
set security policies from-zone vpn to-zone trust policy from-corporate match
    source-address corp-net
set security policies from-zone vpn to-zone trust policy from-corporate match
    source-address sunnyvale-net
set security policies from-zone vpn to-zone trust policy from-corporate match
    destination-address local-net
set security policies from-zone vpn to-zone trust policy from-corporate application any
set security policies from-zone vpn to-zone trust policy from-corporate then permit
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies for the Westford spoke:

1. Create the security policy to permit traffic from the trust zone to the vpn zone.

```
[edit security policies from-zone trust to-zone vpn]
user@spoke# set policy to-corp match source-address local-net
user@spoke# set policy to-corp match destination-address corp-net
user@spoke# set policy to-corp match destination-address sunnyvale-net
user@spoke# set policy to-corp match application any
user@spoke# set policy to-corp then permit
```

2. Create the security policy to permit traffic from the vpn zone to the trust zone.

```
[edit security policies from-zone vpn to-zone trust]
user@spoke# set policy spokes-to-local match source-address corp-net
user@spoke# set policy spokes-to-local match source-address sunnyvale-net
user@spoke# set policy spokes-to-local match destination-address local-net
user@spoke# set policy spokes-to-local match application any
user@spoke# set policy spokes-to-local then permit
```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@spoke# show security policies
from-zone trust to-zone vpn {
  policy to-corp {
    match {
      source-address local-net;
      destination-address [ sunnyvale-net westford-net ];
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone vpn to-zone trust {
  policy spokes-to-local {
    match {
      source-address [ sunnyvale-net westford-net ];
      destination-address local-net;
      application any;
    }
    then {
      permit;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring TCP-MSS for the Westford Spoke

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security flow tcp-mss ipsec-vpn mss 1350
```

Step-by-Step Procedure

To configure TCP-MSS for the Westford spoke:

1. Configure TCP-MSS information.

```
[edit]
user@spoke# set security flow tcp-mss ipsec-vpn mss 1350
```

Results

From configuration mode, confirm your configuration by entering the **show security flow** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@spoke# show security flow
tcp-mss {
  ipsec-vpn {
    mss 1350;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Sunnyvale Spoke

CLI Quick Configuration

This example uses an SSG Series device for the Sunnyvale spoke. For reference, the configuration for the SSG Series device is provided. For information about configuring SSG Series devices, see the *Concepts and Examples ScreenOS Reference Guide*, which is located at http://www.juniper.net/techpubs.

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set zone name "VPN"
set interface ethernet0/6 zone "Trust"
set interface "tunnel.1" zone "VPN"
set interface ethernet0/6 ip 192.168.168.1/24
set interface ethernet0/6 route
set interface ethernet0/0 ip 10.2.2.2/30
set interface ethernet0/0 route
set interface tunnel.1 ip 10.11.11.11/24
set flow tcp-mss 1350
set address "Trust" "sunnyvale-net" 192.168.168.0 255.255.255.0
set address "VPN" "corp-net" 10.10.10.0 255.255.255.0
```

Copyright © 2016, Juniper Networks, Inc.

```
set address "VPN" "westford-net" 192.168.178.0 255.255.255.0
set ike gateway "corp-ike" address 10.1.1.2 Main outgoing-interface ethernet0/0 preshare
   "395psksecr3t" sec-level standard
set vpn corp-vpn monitor optimized rekey
set vpn "corp-vpn" bind interface tunnel.1
set vpn "corp-vpn" gateway "corp-ike" replay tunnel idletime 0 sec-level standard
set policy id 1 from "Trust" to "Untrust" "ANY" "ANY" "ANY" nat src permit
set policy id 2 from "Trust" to "VPN" "sunnyvale-net" "corp-net" "ANY" permit
set policy id 2
exit
set dst-address "westford-net"
exit
set policy id 3 from "VPN" to "Trust" "corp-net" "sunnyvale-net" "ANY" permit
set policy id 3
set src-address "westford-net"
exit
set route 10.10.10.0/24 interface tunnel.1
set route 192.168.178.0/24 interface tunnel.1
set route 0.0.0.0/0 interface ethernet0/0 gateway 10.2.2.1
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the IKE Phase 1 Status on page 6120
- Verifying the IPsec Phase 2 Status on page 6122
- Verifying Next-Hop Tunnel Bindings on page 6123
- Verifying Static Routes for Remote Peer Local LANs on page 6124
- Reviewing Statistics and Errors for an IPsec Security Association on page 6124
- Testing Traffic Flow Across the VPN on page 6125

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action

NOTE: Before starting the verification process, you need to send traffic from a host in the 192.168.10/24 network to a host in the 192.168.168/24 and 192.168.178/24 networks to bring the tunnels up. For route-based VPNs, you can send traffic initiated from the SRX Series device through the tunnel. We recommend that when testing IPsec tunnels, you send test traffic from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate a ping from 192.168.10.10 to 192.168.168.10.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@hub> show security ike security-associations
```


Copyright © 2016, Juniper Networks, Inc.

```
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode
6       10.3.3.2         UP     94906ae2263bbd8e  1c35e4c3fc54d6d3  Main
7       10.2.2.2         UP     7e7a1c0367dfe73c  f284221c656a5fbc  Main

user@hub> show security ike security-associations index 6 detail
IKE peer 10.3.3.2, Index 6,
  Role: Responder, State: UP
  Initiator cookie: 94906ae2263bbd8e,, Responder cookie: 1c35e4c3fc54d6d3
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.1.1.2:500, Remote: 10.3.3.2:500
  Lifetime: Expires in 3571 seconds
  Algorithms:
   Authentication        : sha1
   Encryption            : aes-cbc (128 bits)
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input bytes    :                1128
   Output bytes   :                 988
   Input packets  :                   6
   Output packets :                   5
  Flags: Caller notification sent
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 1
    Negotiation type: Quick mode, Role: Responder, Message ID: 1350777248
    Local: 10.1.1.2:500, Remote: 10.3.3.2:500
    Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Flags: Caller notification sent, Waiting for done
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote Address—Verify that the remote IP address is correct.
- State
  - UP—The Phase 1 SA has been established.
  - DOWN—There was a problem establishing the Phase 1 SA.
- Mode—Verify that the correct mode is being used.

Verify that the following information is correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

Copyright © 2016, Juniper Networks, Inc.

The **show security ike security-associations index 1 detail** command lists additional information about the security association with an index number of 1:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Initiator and responder role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying the IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@hub> show security ipsec security-associations
  total configured sa: 4
  ID    Gateway          Port  Algorithm          SPI      Life:sec/kb  Mon vsys

  <16384 10.2.2.2         500   ESP:aes-128/sha1   b2fc36f8 3364/ unlim   -   0
  >16384 10.2.2.2         500   ESP:aes-128/sha1   5d73929e 3364/ unlim   -   0
  ID    Gateway          Port  Algorithm          SPI      Life:sec/kb  Mon vsys

  <16385 10.3.3.2         500   ESP:3des/sha1      70f789c6 28756/unlim   -   0
  >16385 10.3.3.2         500   ESP:3des/sha1      80f4126d 28756/unlim   -   0

user@hub> show security ipsec security-associations index 16385 detail
  Virtual-system: Root
  Local Gateway: 10.1.1.2, Remote Gateway: 10.3.3.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/24)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 1895270854, AUX-SPI: 0
    Hard lifetime: Expires in 28729 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 28136 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)

    Anti-replay service: enabled, Replay window size: 32

    Direction: outbound, SPI: 2163479149, AUX-SPI: 0
    Hard lifetime: Expires in 28729 seconds
    Lifesize Remaining: Unlimited
```


Copyright © 2016, Juniper Networks, Inc.

```
    Soft lifetime: Expires in 28136 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)

    Anti-replay service: enabled, Replay window size: 32
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The ID number is 16385. Use this value with the **show security ipsec security-associations index** command to get more information about this particular SA.
- There is one IPsec SA pair using port 500, which indicates that no NAT-traversal is implemented. (NAT-traversal uses port 4500 or another random high-number port.)
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 28756/ unlim value indicates that the Phase 2 lifetime expires in 28756 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

The output from the **show security ipsec security-associations index 16385 detail** command lists the following information:

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, are correct for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be specified. For some third-party vendors, the proxy ID must be manually entered to match.
- Another common reason for Phase 2 failure is not specifying the ST interface binding. If IPsec cannot complete, check the kmd log or set traceoptions.

Verifying Next-Hop Tunnel Bindings

Purpose After Phase 2 is complete for all peers, verify the next-hop tunnel bindings.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@hub> show security ipsec next-hop-tunnels
Next-hop gateway  interface    IPSec VPN name              Flag
10.11.11.11       st0.0        sunnyvale-vpn               Static
10.11.11.12       st0.0        westford-vpn                Auto
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The next-hop gateways are the IP addresses for the st0 interfaces of all remote spoke peers. The next hop should be associated with the correct IPsec VPN name. If no NHTB entry exists, there is no way for the hub device to differentiate which IPsec VPN is associated with which next hop.

The Flag field has one of the following values:

- Static— NHTB was manually configured in the st0.0 interface configurations, which is required if the peer is not an SRX Series device.
- Auto— NHTB was not configured, but the entry was automatically populated into the NHTB table during Phase 2 negotiations between two SRX Series devices

There is no NHTB table for any of the spoke sites in this example. From the spoke perspective, the st0 interface is still a point-to-point link with only one IPsec VPN binding.

Verifying Static Routes for Remote Peer Local LANs

Purpose Verify that the static route references the spoke peer's st0 IP address.

Action From operational mode, enter the **show route** command.

```
user@hub> show route 192.168.168.10
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.168.0/24   *[Static/5] 00:08:33
                    > to 10.11.11.11 via st0.0

user@hub> show route 192.168.178.10
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.178.0/24   *[Static/5] 00:04:04
                    > to 10.11.11.12 via st0.0
```

The next hop is the remote peer's st0 IP address, and both routes point to st0.0 as the outgoing interface.

Reviewing Statistics and Errors for an IPsec Security Association

Purpose Review ESP and authentication header counters and errors for an IPsec security association.

Action From operational mode, enter the **show security ipsec statistics index** command.

```
user@hub> show security ipsec statistics index 16385
ESP Statistics:
  Encrypted bytes:              920
  Decrypted bytes:             6208
  Encrypted packets:              5
  Decrypted packets:             87
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
```


Copyright © 2016, Juniper Networks, Inc.

```
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

You can also use the **show security ipsec statistics** command to review statistics and errors for all SAs.

To clear all IPsec statistics, use the **clear security ipsec statistics** command.

Meaning If you see packet loss issues across a VPN, you can run the **show security ipsec statistics** or **show security ipsec statistics detail** command several times to confirm that the encrypted and decrypted packet counters are incrementing. You should also check whether the other error counters are incrementing.

Testing Traffic Flow Across the VPN

Purpose Verify the traffic flow across the VPN.

Action You can use the **ping** command from the SRX Series device to test traffic flow to a remote host PC. Make sure that you specify the source interface so that the route lookup is correct and the appropriate security zones are referenced during policy lookup.

From operational mode, enter the **ping** command.

```
user@hub> ping 192.168.168.10 interface ge-0/0/0 count 5
PING 192.168.168.10 (192.168.168.10): 56 data bytes
64 bytes from 192.168.168.10: icmp_seq=0 ttl=127 time=8.287 ms
64 bytes from 192.168.168.10: icmp_seq=1 ttl=127 time=4.119 ms
64 bytes from 192.168.168.10: icmp_seq=2 ttl=127 time=5.399 ms
64 bytes from 192.168.168.10: icmp_seq=3 ttl=127 time=4.361 ms
64 bytes from 192.168.168.10: icmp_seq=4 ttl=127 time=5.137 ms

--- 192.168.168.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.119/5.461/8.287/1.490 ms
```

You can also use the **ping** command from the SSG Series device.

```
user@hub> ping 192.168.10.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.10.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/4/5 ms

ssg-> ping 192.168.178.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.178.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=8/8/10 ms
```

Copyright © 2016, Juniper Networks, Inc.

Meaning If the **ping** command fails from the SRX Series or SSG Series device, there might be a problem with the routing, security policies, end host, or encryption and decryption of ESP packets.

Related Documentation

- Understanding Hub-and-Spoke VPNs on page 6093
- Example: Configuring a Route-Based VPN on page 6074
- Example: Configuring a Policy-Based VPN on page 6222

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 256

Configuring VPNs for IKEv2

- Understanding Internet Key Exchange Version 2 on page 6127
- Understanding IKEv2 Configuration Payload on page 6128
- Example: Configuring a Route-Based VPN for IKEv2 on page 6130
- Understanding Pico Cell Provisioning on page 6146
- Example: Configuring the SRX Series for Pico Cell Provisioning with IKEv2 Configuration Payload on page 6149

Understanding Internet Key Exchange Version 2

Internet Key Exchange version 2 (IKEv2) is the next generation standard for secure key exchange between peer VPN devices, as defined in RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*.

A VPN peer is configured as either IKEv1 or IKEv2. When a peer is configured as IKEv2, it cannot fall back to IKEv1 if its remote peer initiates IKEv1 negotiation. By default, Juniper Networks security devices are IKEv1 peers. As of Junos OS Release 11.3R1, you can configure a Juniper Networks security device as an IKEv2 peer.

Use the **version v2-only** configuration statement at the [**edit security ike gateway *gw-name***] hierarchy level to configure IKEv2. The IKE version is displayed in the output of the **show security ike security-associations** and **show security ipsec security-associations** CLI operational commands.

The advantages of using IKEv2 over IKEv1 are as follows:

- Replaces eight initial exchanges with a single four-message exchange.
- Reduces the latency for the IPsec SA setup and increases connection establishment speed.
- Increases robustness against DOS attacks.
- Improves reliability through the use of sequence numbers, acknowledgements, and error correction.
- Improves reliability, as all messages are requests or responses. The initiator is responsible for retransmitting if it does not receive a response.

Copyright © 2016, Juniper Networks, Inc.

IKEv2 includes support for:

- Route-based VPNs.

  NOTE: **IKEv2 does not support policy-based VPNs.**

- Site-to-site VPNs.
- Dead peer detection.
- Chassis cluster.
- Certificate-based authentication.
- Child SAs. An IKEv2 child SA is known as a Phase 2 SA in IKEv1. In IKEv2, a child SA cannot exist without the underlying IKE SA. If a child SA is required, it is rekeyed. However, if child SAs are currently active, the corresponding IKE SA is rekeyed.
- AutoVPN.
- Dynamic endpoint VPN.

IKEv2 does not support the following features:

- Policy-based VPN.
- Dialup tunnels.
- VPN monitoring.
- EAP.
- Multiple child SAs for the same traffic selectors for each QoS value.
- IP Payload Compression Protocol (IPComp).
- Traffic selectors.

Related Documentation

- Understanding IKEv2 Configuration Payload on page 6128
- Example: Configuring a Route-Based VPN for IKEv2 on page 6130

Understanding IKEv2 Configuration Payload

Configuration payload is an Internet Key Exchange version 2 (IKEv2) feature used to propagate provisioning information from a responder (or server) to an initiator (or client). IKEv2 configuration payload is supported with route-based VPNs only.

RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*, defines 15 different configuration attributes that can be returned to the initiator by the responder. Table 514 describes the IKEv2 configuration attributes supported on SRX Series devices.

Copyright © 2016, Juniper Networks, Inc.

Table 514: IKEv2 Configuration Attributes

Attribute Type	Value	Description	Length
INTERNAL_IP4_ADDRESS	1	Specifies an address on the internal network. Multiple internal addresses can be requested. The responder can send up to the number of addresses requested.	0 or 4 octets
INTERNAL_IP4_NETMASK	2	Specifies the internal network's netmask value. Only one netmask value is allowed in the request and response messages (for example, 255.255.255.0), and it must be used only with an INTERNAL_IP4_ADDRESS attribute.	0 or 4 octets
INTERNAL_IP4_DNS	3	Specifies an address of a DNS server within the network. Multiple DNS servers can be requested. The responder can respond with zero or more DNS server attributes.	0 or 4 octets
INTERNAL_IP4_NBNS	4	Specifies an address of a NetBIOS name server (NBNS), for example, a WINS server, within the network. Multiple NBNS servers can be requested. The responder can respond with zero or more NBNS server attributes.	0 or 4 octets
INTERNAL_IP4_DHCP	6	Instructs the host to send any internal DHCP request to the address contained within the attribute. Multiple DHCP servers can be requested. The responder can respond with zero or more DHCP server attributes.	0 or 4 octets

For the IKE responder to provide the initiator with provisioning information, it must acquire the information from a specified source such as a RADIUS server. Provisioning information can also be returned from a DHCP server through a RADIUS server. On the RADIUS server, the user information should not include an authentication password. The RADIUS server profile is bound to the IKE gateway using the **xauth *access-profile profile-name*** configuration at the [**edit security ike gateway *gateway-name***] hierarchy level.

In a route-based VPN, secure tunnel (st0) interfaces operate in either point-to-multipoint or point-to-point mode. Dynamic address assignment through the IKEv2 configuration payload is supported for point-to-multipoint interfaces only. For point-to-multipoint interfaces, the interfaces must be numbered and the addresses in the configuration payload INTERNAL_IP4_ADDRESS attribute type must be within the subnetwork range of the associated point-to-multipoint interface.

Related Documentation

- Understanding Internet Key Exchange Version 2 on page 6127
- Understanding Pico Cell Provisioning on page 6146
- Example: Configuring the SRX Series for Pico Cell Provisioning with IKEv2 Configuration Payload on page 6149
- Example: Configuring NAT-T with Dynamic Endpoint VPN on page 6298

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring a Route-Based VPN for IKEv2

This example shows how to configure a route-based IPsec VPN to allow data to be securely transferred between a branch office and a corporate office.

- Requirements on page 6130
- Overview on page 6130
- Configuration on page 6132
- Verification on page 6142

Requirements

This example uses the following hardware:

- SRX240 device
- SSG140 device

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

In this example, you configure a route-based VPN for a branch office in Chicago, Illinois, because you want to conserve tunnel resources but still get granular restrictions on VPN traffic. Users in the Chicago office will use the VPN to connect to their corporate headquarters in Sunnyvale, California.

In this example, you configure interfaces, an IPv4 default route, security zones, and address books. Then you configure IKE Phase 1, IPsec Phase 2, a security policy, and TCP-MSS parameters. See Table 515 through Table 519 for specific configuration parameters used in this example.

Table 515: Interface, Static Route, Security Zone, and Address Book Information

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/0.0	192.168.10.1/24
	ge-0/0/3.0	10.1.1.2/30
	st0.0 (tunnel interface)	10.11.11.10/24
Static routes	0.0.0.0/0 (default route)	The next hop is 10.1.1.1.
	192.168.168.0/24	The next hop is st0.0.

Copyright © 2016, Juniper Networks, Inc.

Table 515: Interface, Static Route, Security Zone, and Address Book Information *(continued)*

Feature	Name	Configuration Parameters
Security zones	trust	• All system services are allowed. • The ge-0/0/0.0 interface is bound to this zone.
	untrust	• IKE is the only allowed system service. • The ge-0/0/3.0 interface is bound to this zone.
	vpn-chicago	The st0.0 interface is bound to this zone.
Address book entries	sunnyvale	• This address is for the trust zone's address book. • The address for this address book entry is 192.168.10.0/24.
	chicago	• This address is for the untrust zone's address book. • The address for this address book entry is 192.168.168.0/24.

Table 516: IKE Phase 1 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ike-phase1-proposal	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: aes-128-cbc
Policy	ike-phase1-policy	• Mode: main • Proposal reference: ike-phase1-proposal • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gw-chicago	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/3.0 • Gateway address: 10.2.2.2

Table 517: IPsec Phase 2 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ipsec-phase2-proposal	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: aes-128-cbc
Policy	ipsec-phase2-policy	• Proposal reference: ipsec-phase2-proposal • PFS: Diffie-Hellman group2

Copyright © 2016, Juniper Networks, Inc.

Table 517: IPsec Phase 2 Configuration Parameters *(continued)*

Feature	Name	Configuration Parameters
VPN	ipsec-vpn-chicago	• IKE gateway reference: gw-chicago • IPsec policy reference: ipsec-phase2-policy • Bind to interface: st0.0

Table 518: Security Policy Configuration Parameters

Purpose	Name	Configuration Parameters

The security policy permits traffic from the trust zone to the vpn-chicago zone.	vpn-tr-chi	• Match criteria:
		  • source-address sunnyvale
		  • destination-address chicago
		  • application any
		• Action: permit

The security policy permits traffic from the vpn-chicago zone to the trust zone.	vpn-chi-tr	• Match criteria:
		  • source-address chicago
		  • destination-address sunnyvale
		  • application any
		• Action: permit

Table 519: TCP-MSS Configuration Parameters

Purpose	Configuration Parameters
TCP-MSS is negotiated as part of the TCP three-way handshake and limits the maximum size of a TCP segment to better fit the MTU limits on a network. For VPN traffic, the IPsec encapsulation overhead, along with the IP and frame overhead, can cause the resulting ESP packet to exceed the MTU of the physical interface, which causes fragmentation. Fragmentation increases bandwidth and device resources. NOTE: We recommend a value of 1350 as the starting point for most Ethernet-based networks with an MTU of 1500 or greater. You might need to experiment with different TCP-MSS values to obtain optimal performance. For example, you might need to change the value if any device in the path has a lower MTU, or if there is any additional overhead such as PPP or Frame Relay.	MSS value: 1350

Configuration

- Configuring Interface, Static Route, Security Zone, and Address Book Information on page 6133
- Configuring IKE on page 6135

Copyright © 2016, Juniper Networks, Inc.

- Configuring IPsec on page 6137
- Configuring Security Policies on page 6139
- Configuring TCP-MSS on page 6140
- Configuring the SSG Series Device on page 6141

Configuring Interface, Static Route, Security Zone, and Address Book Information

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
set interfaces st0 unit 0 family inet address 10.11.11.10/24
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
set routing-options static route 192.168.168.0/24 next-hop st0.0
set security zones security-zone untrust interfaces ge-0/0/3.0
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust address-book address sunnyvale 192.168.10.0/24
set security zones security-zone vpn-chicago interfaces st0.0
set security zones security-zone vpn-chicago address-book address chicago
  192.168.168.0/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interface, static route, security zone, and address book information:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
   user@host# set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
   user@host# set interfaces st0 unit 0 family inet address 10.11.11.10/24
   ```

2. Configure static route information.

   ```
   [edit]
   user@host# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
   user@host# set routing-options static route 192.168.168.0/24 next-hop st0.0
   ```

3. Configure the untrust security zone.

   ```
   [edit ]
   user@host# edit security zones security-zone untrust
   ```

4. Assign an interface to the security zone.

   ```
   [edit security zones security-zone untrust]
   user@host# set interfaces ge-0/0/3.0
   ```

5. Specify allowed system services for the security zone.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@host# edit security zones security-zone trust
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/0.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

9. Configure the address book entry for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set address-book address sunnyvale 192.168.10.0/24
```

10. Configure the vpn-chicago security zone.

```
[edit]
user@host# edit security zones security-zone vpn-chicago
```

11. Assign an interface to the security zone.

```
[edit security zones security-zone vpn-chicago]
user@host# set interfaces st0.0
```

12. Configure the address book entry for the vpn-chicago zone.

```
[edit security zones security-zone vpn-chicago]
user@host# set address-book address chicago 192.168.168.0/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.10.1/24;
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 10.1.1.2/30
    }
  }
}
st0{
```


Copyright © 2016, Juniper Networks, Inc.

```
    unit 0 {
      family inet {
        address 10.11.11.10/24
      }
    }
}

[edit]
user@host# show routing-options
static {
    route 0.0.0.0/0 next-hop 10.1.1.1;
    route 192.168.168.0/24 next-hop st0.0;
}

[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
        system-services {
            ike;
        }
    }
    interfaces {
        ge-0/0/3.0;
    }
}
security-zone trust {
    address-book {
        address sunnyvale 192.168.10.0/24;
    }
    host-inbound-traffic {
        system-services {
            all;
        }
    }
    interfaces {
        ge-0/0/0.0;
    }
}
security-zone vpn-chicago {
    host-inbound-traffic {
        address-book {
            address chicago 192.168.168.0/24;
        }
    }
    interfaces {
        st0.0;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text "$ABC123"
set security ike gateway gw-chicago external-interface ge-0/0/3.0
set security ike gateway gw-chicago ike-policy ike-phase1-policy
set security ike gateway gw-chicago address 10.2.2.2
set security ike gateway gw-chicago version v2-only
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set encryption-algorithm aes-128-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike]
   user@host# set policy ike-phase1-policy
   ```

7. Specify a reference to the IKE proposal.

   ```
   [edit security ike policy ike-phase1-policy]
   user@host# set proposals ike-phase1-proposal
   ```

8. Define the IKE Phase 1 policy authentication method.

   ```
   [edit security ike policy ike-phase1-policy]
   user@host# set pre-shared-key ascii-text "$ABC123"
   ```

9. Create an IKE Phase 1 gateway and define its external interface.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike]
user@host# set gateway gw-chicago external-interface ge-0/0/3.0
```

10. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw-chicago]
user@host# set ike-policy ike-phase1-policy
```

11. Define the IKE Phase 1 gateway address.

```
[edit security ike gateway gw-chicago]
user@host# set address 10.2.2.2
```

12. Define the IKE Phase 1 gateway version.

```
[edit security ike gateway gw-chicago]
user@host# set version v2-only
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security ike
proposal ike-phase1-proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
policy ike-phase1-policy {
    proposals ike-phase1-proposal;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway gw-chicago {
    ike-policy ike-phase1-policy;
    address 10.2.2.2;
    external-interface ge-0/0/3.0;
    version v2-only;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy ipsec-phase2-policy proposals ipsec-phase2-proposal
set security ipsec policy ipsec-phase2-policy perfect-forward-secrecy keys group2
set security ipsec vpn ipsec-vpn-chicago ike gateway gw-chicago
set security ipsec vpn ipsec-vpn-chicago ike ipsec-policy ipsec-phase2-policy
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec vpn ipsec-vpn-chicago bind-interface st0.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec-phase2-proposal
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set authentication-algorithm hmac-sha1-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec-phase2-proposal]
   user@host# set encryption-algorithm aes-128-cbc
   ```

5. Create the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@host# set policy ipsec-phase2-policy
   ```

6. Specify the IPsec Phase 2 proposal reference.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@host# set proposals ipsec-phase2-proposal
   ```

7. Specify IPsec Phase 2 PFS to use Diffie-Hellman group 2.

   ```
   [edit security ipsec policy ipsec-phase2-policy]
   user@host# set perfect-forward-secrecy keys group2
   ```

8. Specify the IKE gateway.

   ```
   [edit security ipsec]
   user@host# set vpn ipsec-vpn-chicago ike gateway gw-chicago
   ```

9. Specify the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@host# set vpn ipsec-vpn-chicago ike ipsec-policy ipsec-phase2-policy
   ```

10. Specify the interface to bind.

    ```
    [edit security ipsec]
    user@host# set vpn ipsec-vpn-chicago bind-interface st0.0
    ```

Results

From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security ipsec
proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
}
policy ipsec-phase2-policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec-phase2-proposal;
}
vpn ipsec-vpn-chicago {
    bind-interface st0.0;
    ike {
        gateway gw-chicago;
        ipsec-policy ipsec-phase2-policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    source-address sunnyvale
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    destination-address chicago
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi match
    application any
set security policies from-zone trust to-zone vpn-chicago policy vpn-tr-chi then permit
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
    source-address chicago
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
    destination-address sunnyvale
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr match
    application any
set security policies from-zone vpn-chicago to-zone trust policy vpn-chi-tr then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies:

1. Create the security policy to permit traffic from the trust zone to the vpn-chicago zone.

   ```
   [edit security policies from-zone trust to-zone vpn-chicago]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy vpn-tr-chi match source-address sunnyvale
user@host# set policy vpn-tr-chi match destination-address chicago
user@host# set policy vpn-tr-chi match application any
user@host# set policy vpn-tr-chi then permit
```

2. Create the security policy to permit traffic from the vpn-chicago zone to the trust zone.

```
[edit security policies from-zone vpn-chicago to-zone trust]
user@host# set policy vpn-chi-tr match source-address sunnyvale
user@host# set policy vpn-chi-tr match destination-address chicago
user@host# set policy vpn-chi-tr match application any
user@host# set policy vpn-chi-tr then permit
```

Results From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
from-zone trust to-zone vpn-chicago {
    policy vpn-tr-vpn {
        match {
            source-address sunnyvale;
            destination-address chicago;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone vpn-chicago to-zone trust {
    policy vpn-tr-vpn {
        match {
            source-address chicago;
            destination-address sunnyvale;
            application any;
        }
        then {
            permit;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring TCP-MSS

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security flow tcp-mss ipsec-vpn mss 1350
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure TCP-MSS information:

1. Configure TCP-MSS information.

```
[edit]
user@host# set security flow tcp-mss ipsec-vpn mss 1350
```

Results

From configuration mode, confirm your configuration by entering the **show security flow** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security flow
tcp-mss {
  ipsec-vpn {
    mss 1350;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the SSG Series Device

CLI Quick Configuration

For reference, the configuration for the SSG Series device is provided. For information about configuring SSG Series devices, see the *Concepts & Examples ScreenOS Reference Guide*, which is located at http://www.juniper.net/techpubs.

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set zone name vpn-chicago
set interface ethernet0/6 zone Trust
set interface ethernet0/0 zone Untrust
set interface tunnel.1 zone vpn-chicago
set interface ethernet0/6 ip 192.168.168.1/24
set interface ethernet0/6 route
set interface ethernet0/0 ip 10.2.2.2/30
set interface ethernet0/0 route
set interface tunnel.1 ip 10.11.11.11/24
set flow tcp-mss 1350
set address Trust “192.168.168-net” 192.168.168.0 255.255.255.0
set address vpn-chicago "192.168.10-net" 192.168.10.0 255.255.255.0
set ike gateway corp-ike address 10.1.1.2 IKEv2 outgoing-interface ethernet0/0 preshare
  395psksecr3t sec-level standard
set vpn corp-vpn gateway corp-ike replay tunnel idletime 0 sec-level standard
set vpn corp-vpn monitor optimized rekey
set vpn corp-vpn bind interface tunnel.1
set policy from Trust to Untrust “ANY” “ANY” “ANY” nat src permit
```

Copyright © 2016, Juniper Networks, Inc.

```
set policy from Trust to vpn-chicago “192.168.168-net” “192.168.10-net” “ANY” permit
set policy from vpn-chicago to Trust “192.168.10-net” “192.168.168-net” “ANY” permit
set route 192.168.10.0/24 interface tunnel.1
set route 0.0.0.0/0 interface ethernet0/0 gateway 10.2.2.1
```

Verification

Confirm that the configuration is working properly.

- Verifying the IKE Phase 1 Status on page 6142
- Verifying the IPsec Phase 2 Status on page 6143
- Reviewing Statistics and Errors for an IPsec Security Association on page 6145
- Testing Traffic Flow Across the VPN on page 6145

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action

NOTE: Before starting the verification process, you need to send traffic from a host in the 192.168.10/24 network to a host in the 192.168.168/24 network. For route-based VPNs, traffic can be initiated by the SRX Series device through the tunnel. We recommend that when testing IPsec tunnels, test traffic be sent from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate a ping from 192.168.10.10 to 192.168.168.10.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode
1       10.2.2.2         UP     744a594d957dd513  1e1307db82f58387  IKEv2

user@host> show security ike security-associations index 1 detail
IKE peer 10.2.2.2, Index 1,
  Role: Responder, State: UP
  Initiator cookie: 744a594d957dd513, Responder cookie: 1e1307db82f58387
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local: 10.1.1.2:500, Remote: 10.2.2.2:500
  Lifetime: Expires in 28570 seconds
  Algorithms:
   Authentication        : sha1
   Encryption            : aes-cbc (128 bits)
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input bytes    :                852
   Output bytes   :                940
   Input packets  :                  5
   Output packets :                  5
  Flags: Caller notification sent
```


Copyright © 2016, Juniper Networks, Inc.

```
IPSec security associations: 1 created, 0 deleted
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote Address—Verify that the remote IP address is correct.
- State
  - UP—The Phase 1 SA has been established.
  - DOWN—There was a problem establishing the Phase 1 SA.
- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets).
- IKE policy parameters.
- Preshared key information.
- Phase 1 proposal parameters (must match on both peers).

The **show security ike security-associations index 1 detail** command lists additional information about the SA with an index number of 1:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Initiator and responder information
- Number of IPsec SAs created

Verifying the IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
  total configured sa: 2
  ID    Gateway     Port  Algorithm       SPI      Life:sec/kb  Mon vsys
  <16384 10.2.2.2    500   ESP:aes-128/sha1   76d64d1d 3363/ unlim   -   0
  >16384 10.2.2.2    500   ESP:aes-128/sha1   a1024ee2 3363/ unlim   -   0

user@host> show security ipsec security-associations index 16384 detail
  Virtual-system: Root
  Local Gateway: 10.1.1.2, Remote Gateway: 10.2.2.2
  Local Identity: ipv4_subnet(any:0,[0..7]=192.168.10.0/24)
  Remote Identity: ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
  Version: IKEv2

    DF-bit: clear

    Direction: inbound, SPI: 1993755933, AUX-SPI: 0
    Hard lifetime: Expires in 3352 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 2775 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)

    Anti-replay service: enabled, Replay window size: 32

    Direction: outbound, SPI: 2701283042, AUX-SPI: 0
    Hard lifetime: Expires in 3352 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 2775 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc
 (128 bits)
    Anti-replay service: enabled, Replay window size: 32
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The ID number is 16384. Use this value with the **show security ipsec security-associations index** command to get more information about this particular SA.
- There is one IPsec SA pair using port 500.
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3363/ unlim value indicates that the Phase 2 lifetime expires in 3363 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, because Phase 2 is not dependent on Phase 1 after the VPN is up.
- The vsys is the root system, and it is always listed as 0.
- The IKEv2 allows connections from a version 2 peer and will initiate a version 2 negotiation.

The output from the **show security ipsec security-associations index 16384 detail** command lists the following information:

 Copyright © 2016, Juniper Networks, Inc.

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, are correct for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be specified. For some third-party vendors, the proxy ID must be manually entered to match.

- Another common reason for Phase 2 failure is not specifying the ST interface binding. If IPsec cannot complete, check the kmd log or set traceoptions.

Reviewing Statistics and Errors for an IPsec Security Association

Purpose Review ESP and authentication header counters and errors for an IPsec SA.

Action From operational mode, enter the **show security ipsec statistics index** *index_number* command, using the index number of the VPN for which you want to see statistics.

```
user@host> show security ipsec statistics index 16384
ESP Statistics:
  Encrypted bytes:              920
  Decrypted bytes:             6208
  Encrypted packets:              5
  Decrypted packets:             87
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

You can also use the **show security ipsec statistics** command to review statistics and errors for all SAs.

To clear all IPsec statistics, use the **clear security ipsec statistics** command.

Meaning If you see packet loss issues across a VPN, you can run the **show security ipsec statistics** or **show security ipsec statistics detail** command several times to confirm that the encrypted and decrypted packet counters are incrementing. You should also check that the other error counters are incrementing.

Testing Traffic Flow Across the VPN

Purpose Verify the traffic flow across the VPN.

Action You can use the **ping** command from the SRX Series device to test traffic flow to a remote host PC. Make sure that you specify the source interface so that the route lookup is correct and the appropriate security zones are referenced during policy lookup.

Copyright © 2016, Juniper Networks, Inc.

From operational mode, enter the **ping** command.

```
ssg-> ping 192.168.168.10 interface ge-0/0/0 count 5
PING 192.168.168.10 (192.168.168.10): 56 data bytes
64 bytes from 192.168.168.10: icmp_seq=0 ttl=127 time=8.287 ms
64 bytes from 192.168.168.10: icmp_seq=1 ttl=127 time=4.119 ms
64 bytes from 192.168.168.10: icmp_seq=2 ttl=127 time=5.399 ms
64 bytes from 192.168.168.10: icmp_seq=3 ttl=127 time=4.361 ms
64 bytes from 192.168.168.10: icmp_seq=4 ttl=127 time=5.137 ms

--- 192.168.168.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.119/5.461/8.287/1.490 ms
```

You can also use the **ping** command from the SSG Series device.

```
user@host> ping 192.168.10.10 from ethernet0/6
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.10.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/4/5 ms
```

Meaning If the **ping** command fails from the SRX Series or SSG Series device, there might be a problem with the routing, security policies, end host, or encryption and decryption of ESP packets.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Hub-and-Spoke VPN on page 6094
- Example: Configuring a Policy-Based VPN on page 6222
- Understanding Internet Key Exchange Version 2 on page 6127

Understanding Pico Cell Provisioning

IKEv2 configuration payload can be used to propagate provisioning information from an IKE responder, such as an SRX Series device, to multiple initiators, such as LTE pico cell base stations in a cellular network. The pico cells ship from the factory with a standard configuration that allows them to connect to the SRX Series device, but the pico cell provisioning information is stored on one or more provisioning servers within a protected network. The pico cells receive full provisioning information after establishing secure connections with the provisioning servers.

Copyright © 2016, Juniper Networks, Inc.

The workflow required to bootstrap and provision a pico cell and introduce it to service includes four distinct stages:

1. Initial addresses acquisition—The pico cell ships from the factory with the following information:

   - Configuration for the secure gateway tunnel to the SRX Series device
   - Digital certificate issued by the manufacturer
   - Fully qualified domain name (FQDN) of the provisioning servers that lie within the protected network

   The pico cell boots up and acquires an address to be used for IKE negotiation from a DHCP server. A tunnel is then built to the secure gateway on the SRX Series device using this address. An address for Operation, Administration, and Management (OAM) traffic is also assigned by the DHCP server for use on the protected network.

2. Pico cell provisioning—Using its assigned OAM traffic address, the pico cell requests its provisioning information—typically operator certificate, license, software, and configuration information—from servers within the protected network.

3. Reboot—The pico cell reboots and uses the acquired provisioning information to make it specific to the service provider's network and operation model.

4. Service provision—When the pico cell enters service, it uses a single certificate that contains distinguished name (DN) and subject alternative name values with a FQDN to build two tunnels to the secure gateway on the SRX Series device: one for OAM traffic and the other for Third-Generation Partnership Project (3GPP) data traffic.

Figure 252 shows a typical workflow for a pico cell deployment.

Copyright © 2016, Juniper Networks, Inc.

Figure 252: Typical Pico Cell Deployment Workflow

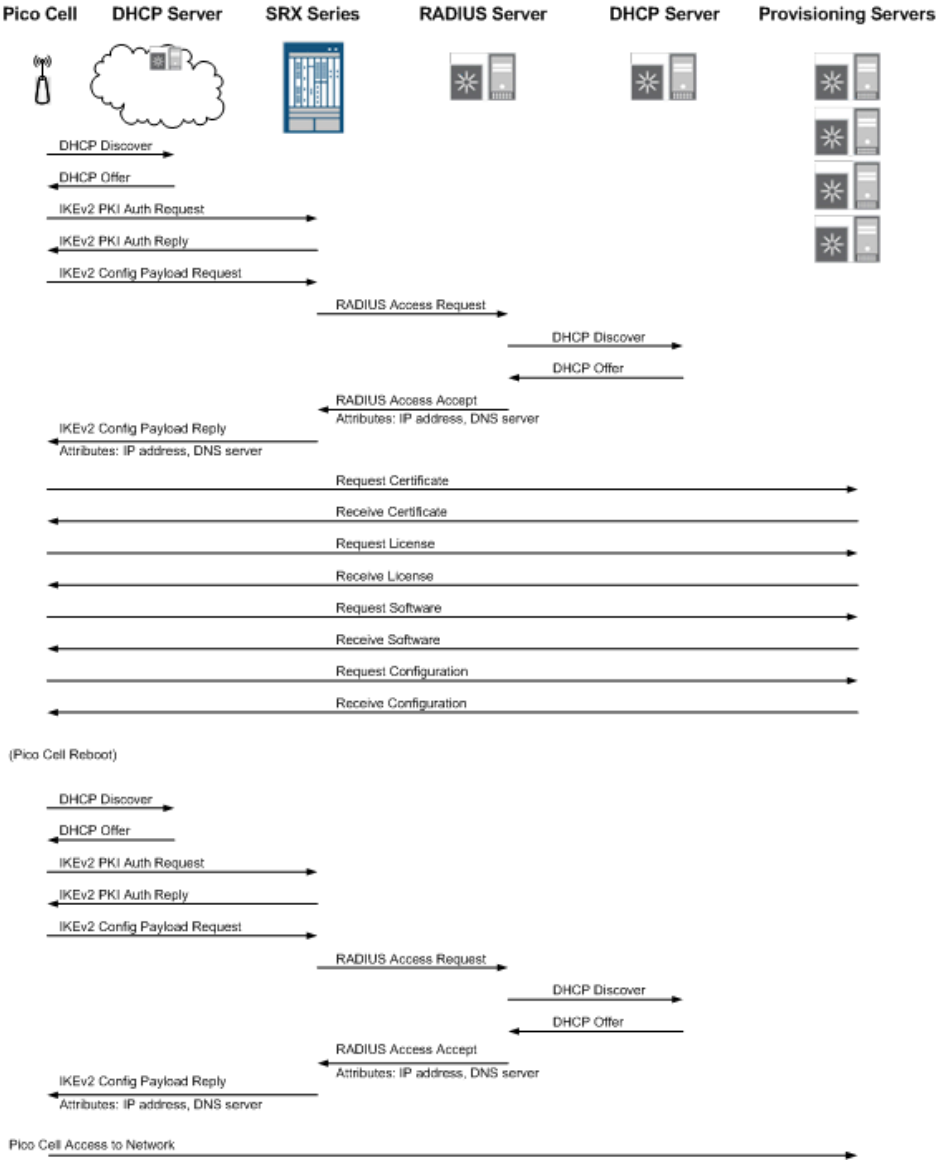

NOTE: The IKEv2 configuration payload feature is supported only for point-to-multipoint secure tunnel (st0) interfaces. Point-to-multipoint interfaces must be numbered, and the addresses provided in the configuration payload must be within the subnetwork range of the associated point-to-multipoint interface.

Related Documentation

- Understanding IKEv2 Configuration Payload on page 6128
- Example: Configuring the SRX Series for Pico Cell Provisioning with IKEv2 Configuration Payload on page 6149

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring the SRX Series for Pico Cell Provisioning with IKEv2 Configuration Payload

In networks where many devices are being deployed, managing the network needs to be simple. The IKEv2 configuration payload feature supports the provisioning of these devices without touching either the device configuration or the SRX Series configuration. This example shows how to configure an SRX Series to support pico cell provisioning using the IKEv2 configuration payload feature.

- Requirements on page 6149
- Overview on page 6149
- Configuration on page 6153
- Verification on page 6169

Requirements

This example uses the following hardware and software components:

- Two SRX Series devices configured in a chassis cluster
- One SRX Series device configured as an intermediate router
- Two pico cell clients
- One RADIUS server configured with pico cell client provisioning information
- Junos OS Release 12.1X46-D10 or later for IKEv2 configuration payload support

Overview

In this example, an SRX Series uses the IKEv2 configuration payload feature to propagate provisioning information to a series of pico cells. The pico cells ship from the factory with a standard configuration that allows them to connect to the SRX Series, but the pico cell provisioning information is stored on an external RADIUS server. The pico cells receive full provisioning information after establishing secure connections with provisioning servers in a protected network. The IKEv2 configuration payload feature is supported for IPv4 only.

Figure 253 shows a topology in which the SRX Series supports pico cell provisioning using the IKEv2 configuration payload feature.

Copyright © 2016, Juniper Networks, Inc.

Figure 253: SRX Series Support for Pico Cell Provisioning with IKEv2 Configuration Payload

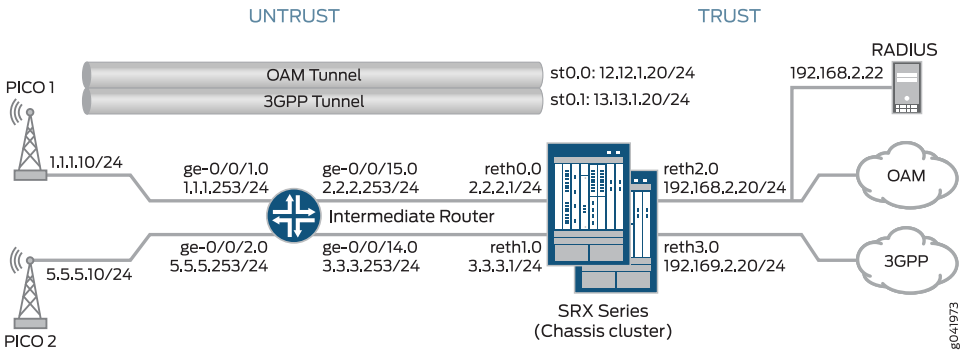

Each pico cell in this topology initiates two IPsec VPNs: one for management and one for data. In this example, management traffic uses the tunnel labeled OAM Tunnel, while the data traffic flows through the tunnel labeled 3GPP Tunnel. Each tunnel supports connections with OAM and 3GPP provisioning servers on separate, configurable networks, requiring separate routing instances and VPNs. This example provides the IKE Phase 1 and Phase 2 options for establishing the OAM and 3GPP VPNs.

In this example, the SRX Series acts as the IKEv2 configuration payload server, acquiring provisioning information from the RADIUS server and providing that information to the pico cell clients. The SRX Series returns the provisioning information for each authorized client in the IKEv2 configuration payload during tunnel negotiation. The SRX Series cannot be used as a client device.

Additionally, the SRX Series uses the IKEv2 configuration payload information to update the Traffic Selector initiator (TSi) and Traffic Selector responder (TSr) values exchanged with the client during tunnel negotiation. The configuration payload uses the TSi and TSr values that are configured on the SRX Series using the **proxy-identity** statement at the [**edit security ipsec vpn *vpn-name* ike**] hierarchy level. The TSi and TSr values define the network traffic for each VPN.

The intermediate router routes pico cell traffic to the appropriate interfaces on the SRX Series.

The following process describes the connection sequence:

1. The pico cell initiates an IPsec tunnel with the SRX Series using the factory configuration.
2. The SRX Series authenticates the client using the client certificate information and the root certificate of the CA that is enrolled in the SRX Series. After authentication, the SRX Series passes the IKE identity information from the client certificate to the RADIUS server in an authorization request.
3. After authorizing the client, the RADIUS server responds to the SRX Series with the client provisioning information:
   - IP address (TSi value)
   - IP subnet mask (optional; the default is 32 bit)

Copyright © 2016, Juniper Networks, Inc.

- DNS address (optional)

4. The SRX Series returns the provisioning information in the IKEv2 configuration payload for each client connection, and exchanges final TSi and TSr values with the pico cells. In this example, the SRX Series provides the following TSi and TSr information for each VPN:

VPN Connection	TSi/TSr Values Provided by SRX
Pico 1 OAM	TSi: 12.12.1.201/32, TSr: 192.168.2.0/24
Pico 1 3GPP	TSi: 13.13.1.201/32, TSr: 192.169.2.0/24, TSr: 13.13.0.0/16
Pico 2 OAM	TSi: 12.12.1.205/32, TSr: 192.168.2.0/24
Pico 2 3GPP	TSi: 13.13.1.205/32, TSr: 192.169.2.0/24, TSr: 13.13.0.0/16

NOTE: If the provisioning information supplied by the RADIUS server includes a subnet mask, the SRX Series returns a second TSr value for the client connection that includes the IP subnet. This enables intrapeer communication for devices on that subnet. In this example, intrapeer communication is enabled for the subnet associated with the 3GPP VPN (13.13.0.0/16).

NOTE: The IKEv2 configuration payload feature is supported only for point-to-multipoint secure tunnel (st0) interfaces. For point-to-multipoint interfaces, the interfaces must be numbered, and the addresses provided in the configuration payload must be within the subnetwork range of the associated point-to-multipoint interface.

Table 520 shows the Phase 1 and Phase 2 options configured on the SRX Series, including information for establishing both OAM and 3GPP tunnels.

Table 520: Phase 1 and Phase 2 Options for the SRX Series

Option	Value
IKE proposal:	
Proposal name	IKE_PROP
Authentication method	RSA digital certificates
Diffie-Hellman (DH) group	group5
Authentication algorithm	SHA-1
Encryption algorithm	AES 256 CBC

Copyright © 2016, Juniper Networks, Inc.

Table 520: Phase 1 and Phase 2 Options for the SRX Series *(continued)*

Option	Value
IKE policy:	
IKE Policy name	IKE_POL
Local certificate	SRX_1
IKE gateway (OAM):	
IKE policy	IKE_POL
Remote IP address	dynamic
IKE user type	group-ike-id
Local IKE ID	hostname srx_series.example.net
Remote IKE ID	hostname .pico_cell.net
External interface	reth0.0
Access profile	radius_pico
IKE version	v2-only
IKE gateway (3GPP):	
IKE policy	IKE_POL
Remote IP address	Dynamic
IKE user type	group-ike-id
Local IKE ID	distinguished-name wildcard OU=srx_series
Remote IKE ID	distinguished-name wildcard OU=pico_cell
External interface	reth1
Access profile	radius_pico
IKE version	v2-only
IPsec proposal:	
Proposal name	IPSEC_PROP
Protocol	ESP
Authentication algorithm	HMAC SHA-1 96

Copyright © 2016, Juniper Networks, Inc.

Table 520: Phase 1 and Phase 2 Options for the SRX Series *(continued)*

Option	Value
Encryption algorithm	AES 256 CBC
IPsec policy:	
Policy name	IPSEC_POL
Perfect Forward Secrecy (PFS) keys	group5
IPsec proposals	IPSEC_PROP
IPsec VPN (OAM):	
Bind interface	st0.0
IKE gateway	OAM_GW
Local proxy-identity	192.168.2.0/24
Remote proxy-identity	0.0.0.0/0
IPsec policy	IPSEC_POL
IPsec VPN (3GPP):	
Bind interface	st0.1
IKE gateway	3GPP_GW
Local proxy-identity	192.169.2.0/24
Remote proxy-identity	0.0.0.0/0
IPsec policy	IPSEC_POL

Certificates are stored on the pico cells and the SRX Series.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Configuration

- Configuring the SRX Series on page 6154
- Configuring the Intermediate Router on page 6163
- Configuring the Pico Cell (Sample Configuration) on page 6166
- Configuring the RADIUS Server (Sample Configuration) on page 6168

Copyright © 2016, Juniper Networks, Inc.

Configuring the SRX Series

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis cluster reth-count 5
set chassis cluster node 0
set chassis cluster node 1
set chassis cluster redundancy-group 0 node 0 priority 250
set chassis cluster redundancy-group 0 node 1 priority 150
set chassis cluster redundancy-group 1 node 0 priority 220
set chassis cluster redundancy-group 1 node 1 priority 149
set chassis cluster redundancy-group 1 interface-monitor ge-3/0/0 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/0/0 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-3/0/1 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/0/1 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-3/2/0 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/2/0 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-3/2/1 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/2/1 weight 255
set interfaces ge-3/0/0 gigether-options redundant-parent reth0
set interfaces ge-3/0/1 gigether-options redundant-parent reth1
set interfaces ge-3/2/0 gigether-options redundant-parent reth2
set interfaces ge-3/2/1 gigether-options redundant-parent reth3
set interfaces ge-8/0/0 gigether-options redundant-parent reth0
set interfaces ge-8/0/1 gigether-options redundant-parent reth1
set interfaces ge-8/2/0 gigether-options redundant-parent reth2
set interfaces ge-8/2/1 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 2.2.2.1./24
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 3.3.3.1/24
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth2 unit 0 family inet address 192.168.2.20/24
set interfaces reth3 redundant-ether-options redundancy-group 1
set interfaces reth3 unit 0 family inet address 192.169.2.20/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 12.12.1.20/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 13.13.1.20/24
set routing-options static route 1.1.0.0/16 next-hop 2.2.2.253
set routing-options static route 5.5.0.0/16 next-hop 2.2.2.253
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces reth0.0
set security zones security-zone untrust interfaces reth1.0
set security zones security-zone oam-trust host-inbound-traffic system-services all
set security zones security-zone oam-trust host-inbound-traffic protocols all
set security zones security-zone oam-trust interfaces reth2.0
set security zones security-zone oam-trust interfaces st0.0
set security zones security-zone 3gpp-trust host-inbound-traffic system-services all
set security zones security-zone 3gpp-trust host-inbound-traffic protocols all
set security zones security-zone 3gpp-trust interfaces reth3.0
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone 3gpp-trust interfaces st0.1
set access profile radius_pico authentication-order radius
set access profile radius_pico radius-server 192.168.2.22 secret "$ABC123"
set access profile radius_pico radius-server 192.168.2.22 routing-instance VR-OAM
set security ike proposal IKE_PROP authentication-method rsa-signatures
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL certificate local-certificate SRX_1
set security ike gateway OAM_GW ike-policy IKE_POL
set security ike gateway OAM_GW dynamic hostname .pico_cell.net
set security ike gateway OAM_GW dynamic ike-user-type group-ike-id
set security ike gateway OAM_GW local-identity hostname srx_series.example.net
set security ike gateway OAM_GW external-interface reth0.0
set security ike gateway OAM_GW xauth access-profile radius_pico
set security ike gateway OAM_GW version v2-only
set security ike gateway 3GPP_GW ike-policy IKE_POL
set security ike gateway 3GPP_GW dynamic distinguished-name wildcard OU=pico_cell
set security ike gateway 3GPP_GW dynamic ike-user-type group-ike-id
set security ike gateway 3GPP_GW local-identity distinguished-name wildcard
  OU=srx_series
set security ike gateway 3GPP_GW external-interface reth1.0
set security ike gateway 3GPP_GW xauth access-profile radius_pico
set security ike gateway 3GPP_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec proposal IPSEC_PROP lifetime-seconds 300
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn OAM_VPN bind-interface st0.0
set security ipsec vpn OAM_VPN ike gateway OAM_GW
set security ipsec vpn OAM_VPN ike proxy-identity local 192.168.2.0/24
set security ipsec vpn OAM_VPN ike proxy-identity remote 0.0.0.0/0
set security ipsec vpn OAM_VPN ike ipsec-policy IPSEC_POL
set security ipsec vpn 3GPP_VPN bind-interface st0.1
set security ipsec vpn 3GPP_VPN ike gateway 3GPP_GW
set security ipsec vpn 3GPP_VPN ike proxy-identity local 192.169.2.0/24
set security ipsec vpn 3GPP_VPN ike proxy-identity remote 0.0.0.0/0
set security ipsec vpn 3GPP_VPN ike ipsec-policy IPSEC_POL
set routing-instances VR-OAM instance-type virtual-router
set routing-instances VR-OAM interface reth2.0
set routing-instances VR-OAM interface st0.0
set routing-instances VR-3GPP instance-type virtual-router
set routing-instances VR-3GPP interface reth3.0
set routing-instances VR-3GPP interface st0.1
set security policies default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the SRX Series:

1. Configure the chassis cluster.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set node 0
user@host# set node 1
user@host#set redundancy-group 0 node 0 priority 250
user@host#set redundancy-group 0 node 1 priority 150
user@host#set redundancy-group 1 node 0 priority 220
user@host#set redundancy-group 1 node 1 priority 149
user@host# set redundancy-group 1 interface-monitor ge-3/0/0 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/0/0 weight 255
user@host# set redundancy-group 1 interface-monitor ge-3/0/1 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/0/1 weight 255
user@host# set redundancy-group 1 interface-monitor ge-3/2/0 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/2/0 weight 255
user@host# set redundancy-group 1 interface-monitor ge-3/2/1 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/2/1 weight 255
```

2. Configure interfaces.

```
[edit interfaces]
user@host# set ge-3/0/0 gigether-options redundant-parent reth0
user@host# set ge-3/0/1 gigether-options redundant-parent reth1
user@host# set ge-3/2/0 gigether-options redundant-parent reth2
user@host# set ge-3/2/1 gigether-options redundant-parent reth3
user@host# set ge-8/0/0 gigether-options redundant-parent reth0
user@host# set ge-8/0/1 gigether-options redundant-parent reth1
user@host# set ge-8/2/0 gigether-options redundant-parent reth2
user@host# set ge-8/2/1 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 2.2.2.1/24
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 3.3.3.1/24
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth2 unit 0 family inet address 192.168.2.20/24
user@host# set reth3 redundant-ether-options redundancy-group 1
user@host# set reth3 unit 0 family inet address 192.169.2.20/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 12.12.1.20/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 13.13.1.20/24
```

3. Configure routing options.

```
[edit routing-options]
user@host# set static route 1.1.0.0/16 next-hop 2.2.2.253
user@host# set static route 5.5.0.0/16 next-hop 2.2.2.253
```

4. Specify security zones.

```
[edit security zones security-zone untrust]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set host-inbound-traffic protocols all
user@host# set host-inbound-traffic system-services all
user@host# set interfaces reth0.0
user@host# set interfaces reth1.0

[edit security zones security-zone oam-trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces reth2.0
user@host# set interfaces st0.0

[edit security zones security-zone 3gpp-trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces reth3.0
user@host# set interfaces st0.1
```

5. Create the RADIUS profile.

```
[edit access profile radius_pico]
user@host# set authentication-order radius
user@host# set radius-server 192.168.2.22 secret "$ABC123"
user@host# set radius-server 192.168.2.22 routing-instance VR-OAM
```

6. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
user@host# set proposals IKE_PROP
user@host# set certificate local-certificate SRX_1

[edit security ike gateway OAM_GW]
user@host# set ike-policy IKE_POL
user@host# set dynamic hostname .pico_cell.net
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity hostname srx.example.net
user@host# set external-interface reth0.0
user@host# set xauth access-profile radius_pico
user@host# set version v2-only

[edit security ike gateway 3GPP_GW]
user@host# set ike-policy IKE_POL
user@host# set dynamic distinguished-name wildcard OU=pico_cell
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name wildcard OU=srx_series
user@host# set external-interface reth1.0
user@host# set xauth access-profile radius_pico
user@host# set version v2-only
```

7. Specify Phase 2 options.

Copyright © 2016, Juniper Networks, Inc.

```
[edit set security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 300

[edit security ipsec policy IPSEC_POL]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals IPSEC_PROP

[edit security ipsec vpn OAM_VPN]
user@host# set bind-interface st0.0
user@host# set ike gateway OAM_GW
user@host# set ike proxy-identity local 192.168.2.0/24
user@host# set ike proxy-identity remote 0.0.0.0/0
user@host# set ike ipsec-policy IPSEC_POL

[edit security ipsec vpn 3GPP_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway 3GPP_GW
user@host# set ike proxy-identity local 192.169.2.0/24
user@host# set ike proxy-identity remote 0.0.0.0/0
user@host# set ike ipsec-policy IPSEC_POL
```

8. Specify the routing instances.

```
[edit routing-instances VR-OAM]
user@host# set instance-type virtual router
user@host# set interface reth2.0
user@host# set interface st0.0

[edit routing-instances VR-3GPP]
user@host# set instance-type virtual router
user@host# set interface reth3.0
user@host# set interface st0.1
```

9. Specify security policies to permit site-to-site traffic.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show chassis cluster**, **show interfaces**, **show security zones**, **show access profile radius_pico**, **show security ike**, **show security ipsec**, **show routing-instances**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis cluster
reth-count 5
node 0
node 1
redundancy-group 0{
   node 0 priority 250;
   node 1 priority 150;
```


Copyright © 2016, Juniper Networks, Inc.

```
    redundancy-group 1 {
    node 0 priority 220;
    node 1 priority 149;
    interface-monitor {
        ge-3/0/0 weight 255;
        ge-8/0/0 weight 255;
        ge-3/0/1 weight 255;
        ge-8/0/1 weight 255;
        ge-3/2/0 weight 255;
        ge-8/2/0 weight 255;
        ge-3/2/1 weight 255;
        ge-8/2/1 weight 255;
    }
}
[edit]
user@host# show interfaces
ge-3/0/0 {
    gigether-options {
        redundant-parent reth0;
    }
}
ge-3/0/1 {
    gigether-options {
        redundant-parent reth1;
    }
}
ge-3/2/0 {
    gigether-options {
        redundant-parent reth2;
    }
}
ge-3/2/1 {
    gigether-options {
        redundant-parent reth3;
    }
}
ge-8/0/0 {
    gigether-options {
        redundant-parent reth0;
    }
}
ge-8/0/1 {
    gigether-options {
        redundant-parent reth1;
    }
}
ge-8/2/0 {
    gigether-options {
        redundant-parent reth2;
    }
}
ge-8/2/1 {
    gigether-options {
        redundant-parent reth3;
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
reth0 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 2.2.2.1/24;
    }
  }
}
reth1 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 3.3.3.1/24;
    }
  }
}
reth2 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 192.168.2.20/24;
    }
  }
}
reth3 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 192.169.2.20/24;
    }
  }
}
st0 {
  unit 0{
    multipoint;
    family inet {
      address 12.12.1.20/24;
    }
  }
  unit 1{
    multipoint;
    family inet {
      address 13.13.1.20/24;
    }
  }
}
[edit]
user@host# show routing-options
```


Copyright © 2016, Juniper Networks, Inc.

```
static {
    route 1.1.0.0/16 next-hop 2.2.2.253;
    route 5.5.0.0/16 next-hop 2.2.2.253;
}
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        reth1.0;
        reth0.0;
    }
}
security-zone oam-trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        reth2.0;
        st0.0;
    }
}
security-zone 3gpp-trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        reth3.0;
        st0.1;
    }
}
[edit]
user@host# show access profile radius_pico
authentication-order radius;
radius-server {
    192.168.2.22 {
        secret "$ABC123"; ## SECRET-DATA
        routing-instance VR-OAM;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show security ike
proposal IKE_PROP {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
policy IKE_POL {
    proposals IKE_PROP;
    certificate {
        local-certificate SRX_1;
    }
}
gateway OAM_GW {
    ike-policy IKE_POL;
    dynamic {
        hostname .pico_cell.net;
        ike-user-type group-ike-id;
    }
    local-identity hostname srx_series.example.net;
    external-interface reth0.0;
    xauth access-profile radius_pico;
    version v2-only;
}
gateway 3GPP_GW {
    ike-policy IKE_POL;
    dynamic {
        distinguished-name {
            wildcard OU=pico_cell;
        }
        ike-user-type group-ike-id;
    }
    local-identity distinguished-name;
    external-interface reth1.0;
    xauth access-profile radius_pico;
    version v2-only;
}
[edit]
user@host# show security ipsec
proposal IPSEC_PROP {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
    lifetime-seconds 300;
}
policy IPSEC_POL {
    perfect-forward-secrecy {
        keys group5;
    }
    proposals IPSEC_PROP;
}
vpn OAM_VPN {
    bind-interface st0.0;
    ike {
```


Copyright © 2016, Juniper Networks, Inc.

```
        gateway OAM_GW;
        proxy-identity {
          local 192.168.2.0/24;
          remote 0.0.0.0/0;
        }
        ipsec-policy IPSEC_POL;
      }
    }
    vpn 3GPP_VPN {
      bind-interface st0.1;
      ike {
        gateway 3GPP_GW;
        proxy-identity {
          local 192.169.2.0/24;
          remote 0.0.0.0/0;
        }
        ipsec-policy IPSEC_POL;
      }
    }
    [edit]
    user@host# show routing-instances
    VR-OAM {
      instance-type virtual-router;
      interface reth2.0;
      interface st0.0;
    }
    VR-3GPP {
      instance-type virtual-router;
      interface reth3.0;
      interface st0.1;
    }
    [edit]
    user@host# show security policies
    default-policy {
      permit-all;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Intermediate Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.253/24
set interfaces ge-0/0/2 unit 0 family inet address 5.5.5.253/24
set interfaces ge-0/0/14 unit 0 family inet address 3.3.3.253/24
set interfaces ge-0/0/15 unit 0 family inet address 2.2.2.253/24
set routing-options static route 192.169.2.0/24 next-hop 2.2.2.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/14.0
set security zones security-zone trust interfaces ge-0/0/15.0
```

Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces ge-0/0/2.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the intermediate router:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 1.1.1.253/24
user@host# set ge-0/0/2 unit 0 family inet address 5.5.5.253/24
user@host# set ge-0/0/14 unit 0 family inet address 3.3.3.253/24
user@host# set ge-0/0/15 unit 0 family inet address 2.2.2.253/24
```

2. Configure routing options.

```
[edit routing-options]
user@host# set static route 192.169.2.0/24 next-hop 2.2.2.1
```

3. Specify security zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic protocols all
user@host# set host-inbound-traffic system-services all
user@host# set interfaces ge-0/0/14.0
user@host# set interfaces ge-0/0/15.0
```

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces ge-0/0/2.0
```

4. Specify security policies.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 1.1.1.253/24;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 5.5.5.253/24;
    }
  }
}
ge-0/0/14 {
  unit 0 {
    family inet {
      address 3.3.3.253/24;
    }
  }
}
ge-0/0/15 {
  unit 0 {
    family inet {
      address 2.2.2.253/24;
    }
  }
}
[edit]
user@host# show routing-options
static {
  route 192.169.2.0/24 next-hop 2.2.2.1;
}
[edit]
user@host# show security zones
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
      all;
    }
  }
  interfaces {
    ge-0/0/14.0;
    ge-0/0/15.0;
  }
}
security-zone untrust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
      all;
    }
  }
  interfaces {
    ge-0/0/1.0;
    ge-0/0/2.0;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
}
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Pico Cell (Sample Configuration)

Step-by-Step Procedure

The pico cell information in this example is provided for reference. Detailed pico cell configuration information is beyond the scope of this document. The pico cell factory configuration must include the following information:

- Local certificate (X.509v3) and IKE identity information
- Traffic Selector (TSi, TSr) values set to any/any (0.0.0.0/0)
- SRX Series IKE identity information and public IP address
- Phase 1 and Phase 2 proposals that match the SRX Series configuration

The pico cells in this example use strongSwan open source software for IPsec-based VPN connections. This information is used by the SRX Series for pico cell provisioning using the IKEv2 configuration payload feature. In networks where many devices are being deployed, the pico cell configuration can be identical except for the certificate (leftcert) and identity (leftid) information. The following sample configurations illustrate factory settings.

1. Review the Pico 1 configuration:

```
conn %default
        ikelifetime=8h
        keylife=1h
        rekeymargin=1m
        keyingtries=1
        keyexchange=ikev2
        authby=pubkey
        mobike=no

conn oam
        left=%any
        leftsourceip=%config
        leftcert=/usr/local/etc/ipsec.d/certs/<cert_name>
        leftid=pico1.pico_cell.net
        leftfirewall=yes
        reauth=yes
        right=2.2.2.1/24
        rightid=srx_series.example.net
        rightsubnet=0.0.0.0/0 #peer net for proxy id
        ike=aes256-sha-modp1536!
        esp=aes256-sha-modp1536!
        auto=add

conn 3gpp
        left=%any
```


Copyright © 2016, Juniper Networks, Inc.

```
        leftsourceip=%config
        leftcert=/usr/local/etc/ipsec.d/certs/<cert_name>
        leftid="C=US, ST=CA, L=Sunnyvale, O=org, OU=pico_cell, CN=pico1"
        leftfirewall=yes
        reauth=yes
        right=3.3.3.1/24
        rightid="OU=srx_series"
        rightsubnet=0.0.0.0/0 #peer net for proxy id
        ike=aes256-sha-modp1536!
        esp=aes256-sha-modp1536!
        auto=add
```

2. Review the Pico 2 configuration:

```
conn %default
        ikelifetime=8h
        keylife=1h
        rekeymargin=1m
        keyingtries=1
        keyexchange=ikev2
        authby=pubkey
        mobike=no

conn oam
        left=%any
        leftsourceip=%config
        leftcert=/usr/local/etc/ipsec.d/certs/<cert_name>
        leftid=pico2.pico_cell.net
        leftfirewall=yes
        #reauth=no
        right=2.2.2.1/24
        rightid=srx_series.example.net
        rightsubnet=0.0.0.0/0 #peer net for proxy id
        ike=aes256-sha-modp1536!
        esp=aes256-sha-modp1536!
        auto=add

conn 3gpp
        left=%any
        leftsourceip=%config
        leftcert=/usr/local/etc/ipsec.d/certs/<cert_name>
        leftid="C=US, ST=CA, L=Sunnyvale, O=org, OU=pico_cell, CN=pico2"
        leftfirewall=yes
        #reauth=no
        right=3.3.3.1/24
        rightid="OU=srx_series"
        rightsubnet=0.0.0.0/0 #peer net for proxy id
        ike=aes256-sha-modp1536!
        esp=aes256-sha-modp1536!
        auto=add
```

Copyright © 2016, Juniper Networks, Inc.

Configuring the RADIUS Server (Sample Configuration)

Step-by-Step Procedure

The RADIUS server information in this example is provided for reference. Complete RADIUS server configuration information is beyond the scope of this document. The following information is returned to the SRX Series by the RADIUS server:

- Framed-IP-Address
- Framed-IP-Netmask (optional)
- Primary-DNS and Secondary-DNS (optional)

In this example, the RADIUS server has separate provisioning information for the OAM and 3GPP connections. The User-Name is taken from the client certificate information provided in the SRX Series authorization request.

NOTE: If the RADIUS server acquires client provisioning information from a DHCP server, the client identity information relayed to the DHCP server by the RADIUS server must be consistent with the client IKE identity information relayed to the RADIUS server by the SRX Series device. This ensures the continuity of the client identity across the various protocols.

1. Review the RADIUS configuration for the Pico 1 OAM VPN. The RADIUS server has the following information:

```
DEFAULT  User-Name =~  "CN\=pico1\,\ C\=US\,\ ST\=CA\,\ L\=Sunnyvale$",
Cleartext-Password := "$ABC123"
        Service-Type = Framed-User,
        Framed-IP-Address = 12.12.1.201,
        Framed-IP-Netmask = 255.255.255.255,
        Primary-Dns = 192.168.2.104,
        Secondary-Dns = 192.168.2.106,
```

   In this case, the RADIUS server provides the default subnet mask (255.255.255.255), which blocks intrapeer traffic.

2. Review the RADIUS configuration for the Pico 1 3GPP VPN. The RADIUS server has the following information:

```
DEFAULT  User-Name =~  "C\=US\,\ ST\=CA\,\ L\=Sunnyvale\,\ O\=org\,\
OU=pico_cell\,\ CN\=pico1\$", Cleartext-Password := "$ABC123"
        Service-Type = Framed-User,
        Framed-IP-Address = 13.13.1.201.10,
        Framed-IP-Netmask = 255.255.0.0,
        Primary-Dns = 192.168.2.104,
        Secondary-Dns = 192.168.2.106,
```

   In this case, the RADIUS server provides a subnet mask value (255.255.0.0), which enables intrapeer traffic.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The clear-text password is hard-coded and is not configurable. Additionally, this example creates two tunnels from the same client certificate by using different parts of the certificate for User-Name (IKE identity) information.

Verification

Confirm that the configuration is working properly.

- Verifying the IKE Phase 1 Status for the SRX Series on page 6169
- Verifying IPsec Security Associations for the SRX Series on page 6171

Verifying the IKE Phase 1 Status for the SRX Series

Purpose Verify the IKE Phase 1 status.

Action From operational mode on node 0, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations detail** command.

```
user@host# show security ike security-associations
node0:
--------------------------------------------------------------------------
Index      State  Initiator cookie  Responder cookie  Mode  Remote Address
553329718  UP     99919a471d1a5278  3be7c5a49172e6c2  IKEv2 1.1.1.1
1643848758 UP     9e31d4323195a195  4d142438106d4273  IKEv2 1.1.1.1

user@host# show security ike security-associations index 553329718 detail
node0:
--------------------------------------------------------------------------
IKE peer 1.1.1.1, Index 553329718, Gateway Name: OAM_GW
  Location: FPC 2, PIC 0, KMD-Instance 1
  Role: Responder, State: UP
  Initiator cookie: 99919a471d1a5278, Responder cookie: 3be7c5a49172e6c2
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 2.2.2.1:500, Remote: 1.1.1.1:500
  Lifetime: Expires in 28738 seconds
  Peer ike-id: C=US, ST=CA, L=Sunnyvale, O=org, OU=pico_cell, CN=pico1
  Xauth assigned IP: 12.12.1.201
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 2104
   Output bytes  :                  425
  Input  packets:                     2
   Output packets:                    1
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 1
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs with pico cells devices. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. This example shows only the IKE Phase 1 SA for the OAM VPN; however, a separate IKE Phase 1 SA will be displayed showing the IKE Phase 1 parameters for the 3GPP VPN.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA: you can use the **show security ike security-associations index detail** command to get more information about the SA.
- Remote address—Verify that the local IP address is correct and that port 500 is being used for peer-to-peer communication.
- Role responder state:
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.
- Peer (remote) IKE ID—Verify the certificate information is correct.
- Local identity and remote identity—Verify these addresses are correct.
- Mode—Verify that the correct mode is being used.

Verify that the following items are correct in your configuration:

- External interfaces (the interface must be the one that sends IKE packets)
- IKE policy parameters
- Phase 1 proposal parameters (must match between peers)

The **show security ike security-associations** command lists the following additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: Troubleshooting is best performed on the peer using the responder role.

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Copyright © 2016, Juniper Networks, Inc.

Verifying IPsec Security Associations for the SRX Series

Purpose Verify the IPsec status.

Action From operational mode on node 0, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations detail** command.

```
user@host# show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 2
  ID         Algorithm             SPI     Life:sec/kb Mon lsys Port Gateway
  <214171651 ESP:aes-cbc-256/sha1 cc2869e2 3529/        -  root 500  1.1.1.1
  >214171651 ESP:aes-cbc-256/sha1 c0a54936 3529/        -  root 500  1.1.1.1
  <205520899 ESP:aes-cbc-256/sha1 84e49026 3521/        -  root 500  1.1.1.1
  >205520899 ESP:aes-cbc-256/sha1 c4ed1849 3521/        -  root 500  1.1.1.1

user@host# show security ipsec security-associations detail
node0:
--------------------------------------------------------------------------
Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x604a29
Last Tunnel Down Reason: SA not initiated
  ID: 214171651 Virtual-system: root, VPN Name: 3GPP_VPN
  Local Gateway: 3.3.3.1, Remote Gateway: 1.1.1.1
  Local Identity: list(any:0,ipv4_subnet(any:0-65535,[0..7]=192.169.2.0/24),
ipv4_subnet(any:0-65535,[0..7]=13.13.0.0/16))
  Remote Identity: ipv4(any:0,[0..3]=13.13.1.201)
    DF-bit: clear
    Bind-interface: st0.1

  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Last Tunnel Down Reason: SA not initiated
    Location: FPC 6, PIC 0, KMD-Instance 2
    Direction: inbound, SPI: cc2869e2, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3523 seconds
    Lifesize Remaining:
    Soft lifetime: Expires in 2965 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

    Location: FPC 6, PIC 0, KMD-Instance 2
    Direction: outbound, SPI: c0a54936, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3523 seconds
    Lifesize Remaining:
    Soft lifetime: Expires in 2965 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

  ID: 205520899 Virtual-system: root, VPN Name: OAM_VPN
  Local Gateway: 2.2.2.1, Remote Gateway: 1.1.1.1
  Local Identity: ipv4_subnet(any:0-65535,[0..7]=192.168.2.0/24)
  Remote Identity: ipv4(any:0,[0..3]=12.12.1.201)
  Version: IKEv2
    DF-bit: clear
```

Copyright © 2016, Juniper Networks, Inc.

```
  Bind-interface: st0.0

Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
Last Tunnel Down Reason: SA not initiated
  Location: FPC 2, PIC 0, KMD-Instance 1
  Direction: inbound, SPI: 84e49026, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 3515 seconds
  Lifesize Remaining:
  Soft lifetime: Expires in 2933 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64

  Location: FPC 2, PIC 0, KMD-Instance 1
  Direction: outbound, SPI: c4ed1849, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 3515 seconds
  Lifesize Remaining:
  Soft lifetime: Expires in 2933 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning This examples shows the active IKE Phase 2 SAs for Pico 1. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IPsec policy parameters in your configuration. For each Phase 2 SA (OAM and 3GPP), information is provided in both the inbound and outboard direction. The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has an IP address of 1.1.1.1.
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3529/ value indicates that the Phase 2 lifetime expires in 3529 seconds, and that no lifesize has been specified, which indicates that it is unlimited. The Phase 2 lifetime can differ from the Phase 1 lifetime, because Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

The above output from the **show security ipsec security-associations index *index_id* detail** command lists the following information:

Copyright © 2016, Juniper Networks, Inc.

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, are correct for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be specified. For some third-party vendors, the proxy ID must be manually entered to match.

- Authentication and encryption algorithms used.
- Phase 2 proposal parameters (must match between peers).
- Secure tunnel (st0.0 and st0.1) bindings to the OAM and 3GPP gateways.

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding Internet Key Exchange Version 2 on page 6127
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 257

Configuring Secure Tunnel Interface in a Virtual Router

- Understanding Virtual Router Support for Route-Based VPNs on page 6175
- Understanding Virtual Router Limitations on page 6176
- Example: Configuring an st0 Interface in a Virtual Router on page 6176

Understanding Virtual Router Support for Route-Based VPNs

For route-based VPNs, you configure a unit of the secure tunnel (st0) interface and bind it to the IPsec VPN tunnel. You can configure a unit of the st0 interface in different virtual router instances. The following functions are supported for nondefault virtual router instances:

- Manual key management
- Transit traffic
- Self-traffic
- VPN monitoring
- Hub-and-spoke VPNs
- Encapsulating Security Payload (ESP) protocol
- Authentication Header (AH) protocol
- Aggressive mode or main mode
- st0 anchored on the loopback (lo0) interface
- Maximum number of virtual routers (VRs) supported on an SRX Series device
- Applications such as Application Layer Gateway (ALG), Intrusion Detection and Prevention (IDP), and Unified Threat Management (UTM)
- Dead peer detection (DPD)
- Chassis cluster active/backup
- Open Shortest Path First (OSPF) over st0

Copyright © 2016, Juniper Networks, Inc.

- Routing Information Protocol (RIP) over st0
- Policy-based VPN inside VR

Related Documentation

- Understanding Virtual Router Limitations on page 6176
- IPsec VPN Overview on page 6041

Understanding Virtual Router Limitations

The following features are not supported for virtual router (VR):

- Public key infrastructure (PKI) inside VR
- Chassis cluster active/active with VPN inside VR

When you configure VPN on SRX Series devices, overlapping of IP addresses across virtual routers is supported with the following limitations:

- An IKE external interface address cannot overlap with any other virtual router.
- An internal or trust interface address can overlap across any other virtual router.
- An st0 interface address cannot overlap in route-based VPN in point-to-multipoint tunnels such as NHTB.
- An st0 interface address can overlap in route-based VPN in point-to-point tunnels.

Related Documentation

- Understanding Virtual Router Support for Route-Based VPNs on page 6175
- IPsec VPN Overview on page 6041

Example: Configuring an st0 Interface in a Virtual Router

This example shows how to configure an st0 interface in a virtual router.

- Requirements on page 6176
- Overview on page 6176
- Configuration on page 6177
- Verification on page 6180

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones. See "Security Zones and Interfaces Overview" on page 1049.

Overview

In this example, you perform the following operations:

- Configure the interfaces.
- Configure IKE Phase 1 proposals.

Copyright © 2016, Juniper Networks, Inc.

- Configure IKE policies, and reference the proposals.
- Configure an IKE gateway, and reference the policy.
- Configure Phase 2 proposals.
- Configure policies, and reference the proposals.
- Configure AutoKey IKE, and reference the policy and gateway.
- Configure the security policy.
- Configure the routing instance.
- Configure the VPN bind to tunnel interface.
- Configure the routing options.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.2/30
set interfaces ge-0/0/1 unit 0 family inet address 10.2.2.2/30
set interfaces st0 unit 0 family inet address 10.3.3.2/30
set security ike proposal first_ikeprop authentication-method pre-shared-keys
set security ike proposal first_ikeprop dh-group group2
set security ike proposal first_ikeprop authentication-algorithm md5
set security ike proposal first_ikeprop encryption-algorithm 3des-cbc
set security ike policy first_ikepol mode main
set security ike policy first_ikepol proposals first_ikeprop
set security ike policy first_ikepol pre-shared-key ascii-text "$ABC123"
set security ike gateway first ike-policy first_ikepol
set security ike gateway first address 10.4.4.2
set security ike gateway first external-interface ge-0/0/0.0
set security ipsec proposal first_ipsecprop protocol esp
set security ipsec proposal first_ipsecprop authentication-algorithm hmac-md5-96
set security ipsec proposal first_ipsecprop encryption-algorithm 3des-cbc
set security ipsec policy first_ipsecpol perfect-forward-secrecy keys group1
set security ipsec policy first_ipsecpol proposals first_ipsecprop
set security ipsec vpn first_vpn bind-interface st0.0
set security ipsec vpn first_vpn ike gateway first
set security ipsec vpn first_vpn ike ipsec-policy first_ipsecpol
set security ipsec vpn first_vpn establish-tunnels immediately
set security policies default-policy permit-all
set routing-instances VR1 instance-type virtual-router
set routing-instances VR1 interface ge-0/0/1.0
set routing-instances VR1 interface st0.0
set routing-instances VR1 routing-options static route 10.6.6.0/24 next-hop st0.0
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an st0 in a VR:

1. Configure the interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.1.2/30
user@host# set interfaces ge-0/0/1 unit 0 family inet address 10.2.2.2/30
user@host# set interfaces st0 unit 0 family inet address 10.3.3.2/30
```

2. Configure Phase 1 of the IPsec tunnel.

```
[edit security ike]
user@host# set proposal first_ikeprop authentication-method pre-shared-keys
user@host# set proposal first_ikeprop dh-group group2
user@host# set proposal first_ikeprop authentication-algorithm md5
user@host# set proposal first_ikeprop encryption-algorithm 3des-cbc
```

3. Configure the IKE policies, and reference the proposals.

```
[edit security ike]
user@host# set policy first_ikepol mode main
user@host# set policy first_ikepol proposals first_ikeprop
user@host# set policy first_ikepol pre-shared-key ascii-text "$ABC123"
```

4. Configure the IKE gateway, and reference the policy.

```
[edit security ike]
user@host# set gateway first ike-policy first_ikepol
user@host# set gateway first address 10.4.4.2
user@host# set gateway first external-interface ge-0/0/0.0
```

5. Configure Phase 2 of the IPsec tunnel.

```
[edit security ipsec]
user@host# set proposal first_ipsecprop protocol esp
user@host# set proposal first_ipsecprop authentication-algorithm hmac-md5-96
user@host# set proposal first_ipsecprop encryption-algorithm 3des-cbc
```

6. Configure the policies, and reference the proposals.

```
[edit security ipsec]
user@host# set policy first_ipsecpol perfect-forward-secrecy keys group1
user@host# set policy first_ipsecpol proposals first_ipsecprop
```

7. Configure AutoKey IKE, and reference the policy and gateway.

```
[edit security ipsec]
user@host# set vpn first_vpn ike gateway first
user@host# set vpn first_vpn ike ipsec-policy first_ipsecpol
user@host# set vpn first_vpn establish-tunnels immediately
```

8. Configure the VPN bind to tunnel interface.

```
[edit security ipsec]
user@host# set vpn first_vpn bind-interface st0.0
```

9. Configure the security policy.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies]
user@host# set default-policy permit-all
```

10. Configure the st0 in the routing instance.

```
[edit routing-instances]
user@host# set VR1 instance-type virtual-router
user@host# set VR1 interface ge-0/0/1.0
user@host# set VR1 interface st0.0
```

11. Configure the routing options.

```
[edit routing-instances VR1 routing-options]
user@host# set static route 10.6.6.0/24 next-hop st0.0
```

Results From configuration mode, confirm your configuration by entering the **show security** and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show security
  ike {
  proposal first_ikeprop {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm md5;
    encryption-algorithm 3des-cbc;
  }
  policy first_ikepol {
    mode main;
    proposals first_ikeprop;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
  }
  gateway first {
    ike-policy first_ikepol;
    address 10.4.4.2;
    external-interface ge-0/0/0.0;
  }
}
  ipsec {
    proposal first_ipsecprop {
      protocol esp;
      authentication-algorithm hmac-md5-96;
      encryption-algorithm 3des-cbc;
    }
    policy first_ipsecpol {
      perfect-forward-secrecy {
        keys group1;
      }
      proposals first_ipsecprop;
    }
    vpn first_vpn {
      bind-interface st0.0;
      ike {
        gateway first;
        ipsec-policy first_ipsecpol;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
        establish-tunnels immediately;
      }
    }
  policies {
    default-policy {
      permit-all;
    }
  }
  user@host# show routing-instances
      VR1 {
        instance-type virtual-router;
        interface ge-0/0/1.0;
        interface st0.0;
        routing-options {
        static {
        route 10.6.6.0/24 next-hop st0.0;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying an st0 interface in the Virtual Router on page 6180

Verifying an st0 interface in the Virtual Router

Purpose Verify the st0 interface in the virtual router.

Action From operational mode, enter the **show interfaces st0.0 detail** command. The number listed for routing table corresponds to the order that the routing tables in the **show route all** command.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 258

Configuring Dual Stack Tunnels over an External Interface

- Understanding VPN Tunnel Modes on page 6181
- Understanding Dual-Stack Tunnels over an External Interface on page 6183
- Example: Configuring Dual-Stack Tunnels over an External Interface on page 6184

Understanding VPN Tunnel Modes

In VPN tunnel mode, IPsec encapsulates the original IP datagram—including the original IP header—within a second IP datagram. The outer IP header contains the IP address of the gateway, while the inner header contains the ultimate source and destination IP addresses. The outer and inner IP headers can have a protocol field of IPv4 or IPv6. SRX Series devices support four tunnel modes for route-based site-to-site VPNs.

IPv4-in-IPv4 tunnels encapsulate IPv4 packets inside IPv4 packets, as shown in Figure 254. The protocol fields for both the outer and the inner headers are IPv4.

Figure 254: IPv4-in-IPv4 Tunnel

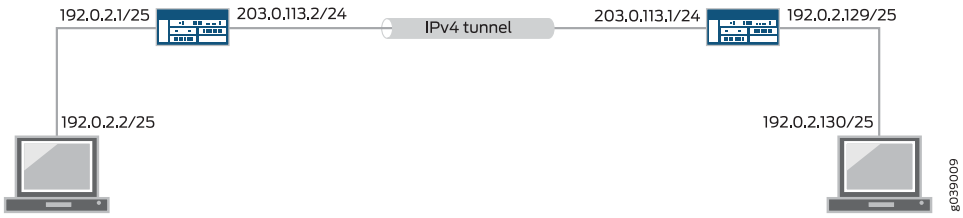

IPv6-in-IPv6 tunnels encapsulate IPv6 packets inside IPv6 packets, as shown in Figure 255. The protocol fields for both the outer and inner headers are IPv6.

Copyright © 2016, Juniper Networks, Inc.

Figure 255: IPv6-in-IPv6 Tunnel

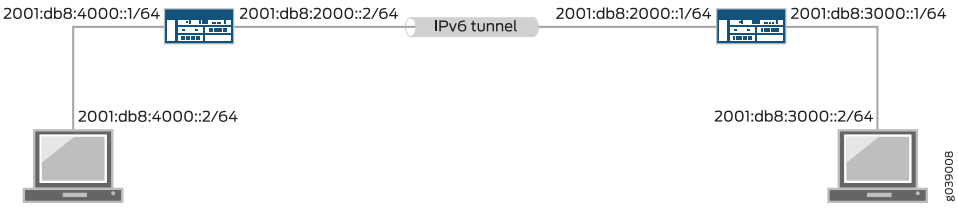

IPv6-in-IPv4 tunnels encapsulate IPv6 packets inside IPv4 packets, as shown in Figure 256. The protocol field for the outer header is IPv4 and the protocol field for the inner header is IPv6.

Figure 256: IPv6-in-IPv4 Tunnel

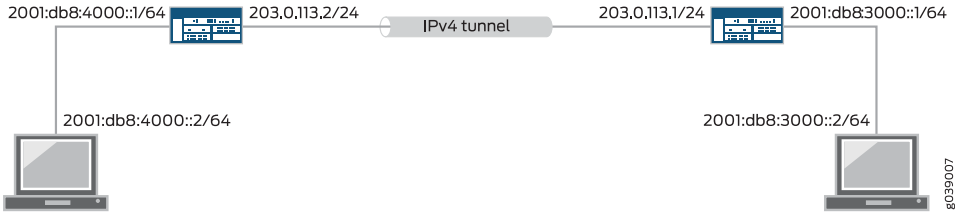

IPv4-in-IPv6 tunnels encapsulate IPv4 packets inside IPv6 packets, as shown in Figure 257. The protocol field for the outer header is IPv6 and the protocol field for the inner header is IPv4.

Figure 257: IPv4-in-IPv6 Tunnel

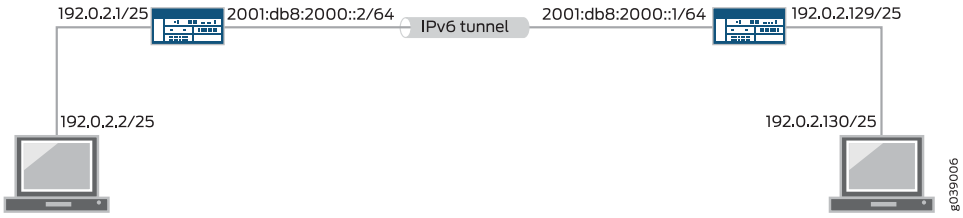

A single IPsec VPN tunnel can carry both IPv4 and IPv6 traffic. For example, an IPv4 tunnel can operate in both IPv4-in-IPv4 and IPv6-in-IPv4 tunnel modes at the same time. To allow both IPv4 and IPv6 traffic over a single IPsec VPN tunnel, the st0 interface bound to that tunnel must be configured with both **family inet** and **family inet6**.

A physical interface configured with both IPv4 and IPv6 addresses can be used as the external interface for parallel IPv4 and IPv6 tunnels to a peer in a route-based site-to-site VPN. This feature is known as dual-stack tunnels and requires separate st0 interfaces for each tunnel.

Copyright © 2016, Juniper Networks, Inc.

NOTE: For policy-based VPNs, IPv6-in-IPv6 is the only tunnel mode supported and it is only supported on branch SRX Series devices.

Related Documentation

- VPN Feature Support for IPv6 Addresses on page 6331
- Understanding Dual-Stack Tunnels over an External Interface on page 6183
- Understanding IPv6 IKE and IPsec Packet Processing on page 6335

Understanding Dual-Stack Tunnels over an External Interface

Dual-stack tunnels—parallel IPv4 and IPv6 tunnels over a single physical interface to a peer—are supported for route-based site-to-site VPNs. A physical interface configured with both IPv4 and IPv6 addresses can be used as the external interface to IPv4 and IPv6 gateways on the same peer or on different peers at the same time. In Figure 258, the physical interfaces reth0.0 and ge-0/0/0.1 support parallel IPv4 and IPv6 tunnels between two devices.

Figure 258: Dual-Stack Tunnels

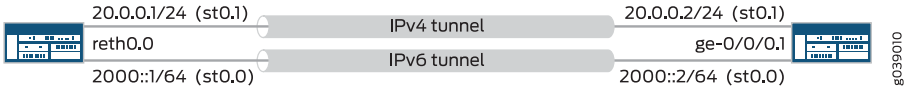

NOTE: In Figure 258, separate secure tunnel (st0) interfaces must be configured for each IPsec VPN tunnel. Parallel IPv4 and IPv6 tunnels that are bound to the same st0 interface are not supported.

A single IPsec VPN tunnel can carry both IPv4 and IPv6 traffic. For example, an IPv4 tunnel can operate in both IPv4-in-IPv4 and IPv6-in-IPv4 tunnel modes at the same time. To allow both IPv4 and IPv6 traffic over a single IPsec VPN tunnel, the st0 interface bound to that tunnel must be configured with both **family inet** and **family inet6**.

If multiple addresses in the same address family are configured on the same external interface to a VPN peer, we recommend that you configure **local-address** at the [**edit security ike gateway *gateway-name***] hierarchy level.

If **local-address** is configured, the specified IPv4 or IPv6 address is used as the local gateway address. If only one IPv4 and one IPv6 address is configured on a physical external interface, **local-address** configuration is not required.

Copyright © 2016, Juniper Networks, Inc.

NOTE: **local-address** must be an IP address that is configured on an interface on the SRX Series device. We recommend that **local-address** belong to the external interface of the IKE gateway. If **local-address** does not belong to the external interface of the IKE gateway, the interface must be in the same zone as the external interface of the IKE gateway and an intra-zone security policy must be configured to permit traffic.

local-address and the remote IKE gateway address must be in the same address family, either IPv4 or IPv6.

If **local-address** is not configured, the local gateway address is based on the remote gateway address. If the remote gateway address is an IPv4 address, the local gateway address is the primary IPv4 address of the external physical interface. If the remote gateway address is an IPv6 address, the local gateway address is the primary IPv6 address of the external physical interface.

Related Documentation

- Example: Configuring Dual-Stack Tunnels over an External Interface on page 6184
- Understanding VPN Tunnel Modes on page 6181
- VPN Feature Support for IPv6 Addresses on page 6331

Example: Configuring Dual-Stack Tunnels over an External Interface

This example shows how to configure parallel IPv4 and IPv6 tunnels over a single external physical interface to a peer for route-based site-to-site VPNs.

- Requirements on page 6184
- Overview on page 6184
- Configuration on page 6187
- Verification on page 6191

Requirements

Before you begin, read "Understanding Dual-Stack Tunnels over an External Interface" on page 6183.

NOTE: The configuration shown in this example is only supported with route-based site-to-site VPNs.

Overview

In this example, a redundant Ethernet interface on the local device supports parallel IPv4 and IPv6 tunnels to a peer device:

- The IPv4 tunnel carries IPv6 traffic; it operates in IPv6-in-IPv4 tunnel mode. The secure tunnel interface st0.0 bound to the IPv4 tunnel is configured with family inet6 only.

Copyright © 2016, Juniper Networks, Inc.

- The IPv6 tunnel carries both IPv4 and IPv6 traffic; it operates in both IPv4-in-IPv6 and IPv6-in-IPv6 tunnel modes. The secure tunnel interface st0.1 bound to the IPv6 tunnel is configured with both family inet and family inet6.

Table 521 shows the Phase 1 options used in this example. The Phase 1 option configuration includes two IKE gateway configurations, one to the IPv6 peer and the other to the IPv4 peer.

Table 521: Phase 1 Options for Dual-Stack Tunnel Configuration

Option	Value
IKE proposal	ike_proposal
Authentication method	Preshared keys
Authentication algorithm	MD5
Encryption algorithm	3DES CBC
Lifetime	3600 seconds
IKE policy	ike_policy
Mode	Aggressive
IKE proposal	ike_proposal
Preshared key	ASCII text
IPv6 IKE gateway	ike_gw_v6
IKE policy	ike_policy
Gateway address	2000::2
External interface	reth1.0
IKE version	IKEv2
IPv4 IKE gateway	ike_gw_v4
IKE policy	ike_policy
Gateway address	20.0.0.2
External interface	reth1.0

Table 522 shows the Phase 2 options used in this example. The Phase 2 option configuration includes two VPN configurations, one for the IPv6 tunnel and the other for the IPv4 tunnel.

Copyright © 2016, Juniper Networks, Inc.

Table 522: Phase 2 Options for Dual-Stack Tunnel Configuration

Option	Value
IPsec proposal	ipsec_proposal
Protocol	ESP
Authentication algorithm	HMAC SHA-1 96
Encryption algorithm	3DES CBC
IPsec policy	ipsec_policy
Proposal	ipsec_proposal
IPv6 VPN	test_s2s_v6
Bind interface	st0.1
IKE gateway	ike_gw_v6
IKE IPsec policy	ipsec_policy
Establish tunnels	Immediately
IPv4 VPN	test_s2s_v4
Bind interface	st0.0
IKE gateway	ike_gw_4
IKE IPsec policy	ipsec_policy

The following static routes are configured in the IPv6 routing table:

- Route IPv6 traffic to 3000::1/128 through st0.0.
- Route IPv6 traffic to 3000::2/128 through st0.1.

A static route is configured in the default (IPv4) routing table to route IPv4 traffic to 30.0.0.0/24 through st0.1.

NOTE: **Flow-based processing of IPv6 traffic must be enabled with the mode flow-based configuration option at the [edit security forwarding-options family inet6] hierarchy level.**

Copyright © 2016, Juniper Networks, Inc.

Topology

In Figure 259, the SRX Series device A supports IPv4 and IPv6 tunnels to device B. IPv6 traffic to 3000::1/128 is routed through the IPv4 tunnel, while IPv6 traffic to 3000::2/128 and IPv4 traffic to 30.0.0.0/24 are routed through the IPv6 tunnel.

Figure 259: Dual-Stack Tunnel Example

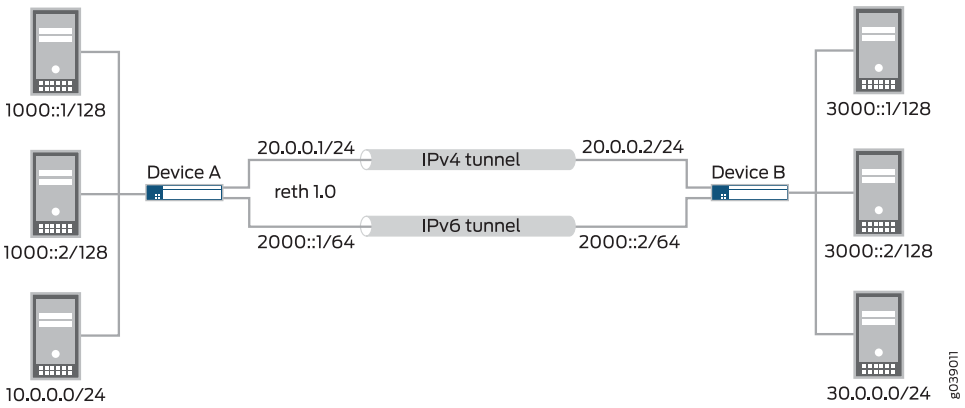

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 gigether-options redundant-parent reth1
set interfaces ge-8/0/1 gigether-options redundant-parent reth1
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 20.0.0.1/24
set interfaces reth1 unit 0 family inet6 address 2000::1/64
set interfaces st0 unit 0 family inet6
set interfaces st0 unit 1 family inet
set interfaces st0 unit 1 family inet6
set security ike proposal ike_proposal authentication-method pre-shared-keys
set security ike proposal ike_proposal authentication-algorithm md5
set security ike proposal ike_proposal encryption-algorithm 3des-cbc
set security ike proposal ike_proposal lifetime-seconds 3600
set security ike policy ike_policy mode aggressive
set security ike policy ike_policy proposals ike_proposal
set security ike policy ike_policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike_gw_v6 ike-policy ike_policy
set security ike gateway ike_gw_v6 address 2000::2
set security ike gateway ike_gw_v6 external-interface reth1.0
set security ike gateway ike_gw_v6 version v2-only
set security ike gateway ike_gw_v4 ike-policy ike_policy
set security ike gateway ike_gw_v4 address 20.0.0.2
set security ike gateway ike_gw_v4 external-interface reth1.0
set security ipsec proposal ipsec_proposal protocol esp
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec proposal ipsec_proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_proposal encryption-algorithm 3des-cbc
set security ipsec policy ipsec_policy proposals ipsec_proposal
set security ipsec vpn test_s2s_v6 bind-interface st0.1
set security ipsec vpn test_s2s_v6 ike gateway ike_gw_v6
set security ipsec vpn test_s2s_v6 ike ipsec-policy ipsec_policy
set security ipsec vpn test_s2s_v6 establish-tunnels immediately
set security ipsec vpn test_s2s_v4 bind-interface st0.0
set security ipsec vpn test_s2s_v4 ike gateway ike_gw_v4
set security ipsec vpn test_s2s_v4 ike ipsec-policy ipsec_policy
set routing-options rib inet6.0 static route 3000::1/128 next-hop st0.0
set routing-options rib inet6.0 static route 3000::2/128 next-hop st0.1
set routing-options static route 30.0.0.0/24 next-hop st0.1
set security forwarding-options family inet6 mode flow-based
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure dual-stack tunnels:

1. Configure the external interface.

```
[edit interfaces]
user@host# set ge-0/0/1 gigether-options redundant-parent reth1
user@host# set ge-8/0/1 gigether-options redundant-parent reth1
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 20.0.0.1/24
user@host# set reth1 unit 0 family inet6 address 2000::1/64
```

2. Configure the secure tunnel interfaces.

```
[edit interfaces]
user@host# set st0 unit 0 family inet6
user@host# set st0 unit 1 family inet
user@host# set st0 unit 1 family inet6
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike_proposal]
user@host# set authentication-method pre-shared-keys
user@host# set authentication-algorithm md5
user@host# set encryption-algorithm 3des-cbc
user@host# set lifetime-seconds 3600

[edit security ike policy ike_policy]
user@host# set mode aggressive
user@host# set proposals ike_proposal
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security ike gateway ike_gw_v6]
user@host# set ike-policy ike_policy
user@host# set address 2000::2
user@host# set external-interface reth1.0
user@host# set version v2-only
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike gateway ike_gw_v4]
user@host# set ike-policy ike_policy
user@host# set address 20.0.0.2
user@host# set external-interface reth1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm 3des-cbc

[edit security ipsec policy ipsec_policy]
user@host# set proposals ipsec_proposal

[edit security ipsec vpn test_s2s_v6 ]
user@host# set bind-interface st0.1
user@host# set ike gateway ike_gw_v6
user@host# set ike ipsec-policy ipsec_policy
user@host# set establish-tunnels immediately

[edit security ipsec vpn test_s2s_v4]
user@host# set bind-interface st0.0
user@host# set ike gateway ike_gw_v4
user@host# set ike ipsec-policy ipsec_policy
```

5. Configure static routes.

```
[edit routing-options rib inet6.0]
user@host# set static route 3000::1/128 next-hop st0.0
user@host# set static route 3000::2/128 next-hop st0.1

[edit routing-options]
user@host# set static route 30.0.0.0/24 next-hop st0.1
```

6. Enable IPv6 flow-based forwarding.

```
[edit security forwarding-options]
user@host# set family inet6 mode flow-based
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show routing-options**, and **show security forwarding-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces
  ge-0/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-8/0/1 {
    gigether-options {
      redundant-parent reth1;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  reth1 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 20.0.0.1/24;
      }
      family inet6 {
        address 2000::1/64;
      }
    }
  }
  st0 {
    unit 0 {
      family inet;
      family inet6;
    }
    unit 1 {
      family inet6;
    }
  }
  [edit]
  user@host# show security ike
  proposal ike_proposal {
    authentication-method pre-shared-keys;
    authentication-algorithm md5;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 3600;
  }
  policy ike_policy {
    mode aggressive;
    proposals ike_proposal;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
  }
  gateway ike_gw_v6 {
    ike-policy ike_policy;
    address 2000::2;
    external-interface reth1.0;
    version v2-only;
  }
  gateway ike_gw_4 {
    ike-policy ike_policy;
    address 20.0.0.2;
    external-interface reth1.0;
  }
  [edit]
  user@host# show security ipsec
  proposal ipsec_proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
  }
  policy ipsec_policy {
```


Copyright © 2016, Juniper Networks, Inc.

```
    proposals ipsec_proposal;
  }
  vpn test_s2s_v6 {
    bind-interface st0.1;
    ike {
      gateway ike_gw_v6;
      ipsec-policy ipsec_policy;
    }
    establish-tunnels immediately;
  }
  vpn test_s2s_v4 {
    bind-interface st0.0;
    ike {
      gateway ike_gw_4;
      ipsec-policy ipsec_policy;
    }
  }
  [edit]
  user@host# show routing-options
  rib inet6.0 {
    static {
      route 3000::1/128 next-hop st0.0;
      route 3000::2/128 next-hop st0.1;
    }
  }
  static {
    route 30.0.0.0/24 next-hop st0.1;
  }
  [edit]
user@host# show security forwarding-options
  family {
    inet6 {
      mode flow-based;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying IKE Phase 1 Status on page 6191
- Verifying IPsec Phase 2 Status on page 6192
- Verifying Routes on page 6192

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address
```

```
1081812113 UP  51d9e6df8a929624  7bc15bb40781a902  IKEv2          2000::2

1887118424 UP  d80b55b949b54f0a  b75ecc815529ae8f  Aggressive     20.0.0.2
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the peer devices.

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command.

```
user@host> show security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <131074 ESP:3des/sha1 8828bd36 3571/  unlim     -   root 500   20.0.0.2

  >131074 ESP:3des/sha1 c968afd8 3571/  unlim     -   root 500   20.0.0.2

  <131073 ESP:3des/sha1 8e9e695a 3551/  unlim     -   root 500   2000::2

  >131073 ESP:3des/sha1 b3a254d1 3551/  unlim     -   root 500   2000::2
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the peer devices.

Verifying Routes

Purpose Verify active routes.

Action From operational mode, enter the **show route** command.

```
user@host> show route
inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.5.0.0/16         *[Static/5] 3d 01:43:23
                     > to 10.157.64.1 via fxp0.0
10.10.0.0/16        *[Static/5] 3d 01:43:23
                     > to 10.157.64.1 via fxp0.0
10.150.0.0/16       *[Static/5] 3d 01:43:23
                     > to 10.157.64.1 via fxp0.0
10.150.48.0/21      *[Static/5] 3d 01:43:23
                     > to 10.157.64.1 via fxp0.0
10.155.0.0/16       *[Static/5] 3d 01:43:23
                     > to 10.157.64.1 via fxp0.0
10.157.64.0/19      *[Direct/0] 3d 01:43:23
                     > via fxp0.0
10.157.72.36/32     *[Local/0] 3d 01:43:23
```


Copyright © 2016, Juniper Networks, Inc.

```
                    Local via fxp0.0
10.204.0.0/16      *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
10.206.0.0/16      *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
10.209.0.0/16      *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
20.0.0.0/24        *[Direct/0] 03:45:41
                    > via reth1.0
20.0.0.1/32        *[Local/0] 03:45:41
                      Local via reth1.0
30.0.0.0/24        *[Static/5] 00:07:49
                    > via st0.1
50.0.0.0/24        *[Direct/0] 03:45:42
                    > via reth0.0
50.0.0.1/32        *[Local/0] 03:45:42
                      Local via reth0.0
172.16.0.0/12      *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
192.168.0.0/16     *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
192.168.102.0/23   *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
207.17.136.0/24    *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0
207.17.136.192/32  *[Static/5] 3d 01:43:23
                    > to 10.157.64.1 via fxp0.0

inet6.0: 10 destinations, 14 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2000::/64          *[Direct/0] 03:45:41
                    > via reth1.0
2000::1/128        *[Local/0] 03:45:41
                      Local via reth1.0
3000::1/128        *[Static/5] 00:03:45
                    > via st0.0
3000::2/128        *[Static/5] 00:03:45
                    > via st0.1
5000::/64          *[Direct/0] 03:45:42
                    > via reth0.0
5000::1/128        *[Local/0] 03:45:42
                      Local via reth0.0
fe80::/64          *[Direct/0] 03:45:42
                    > via reth0.0
                    [Direct/0] 03:45:41
                    > via reth1.0
                    [Direct/0] 03:45:41
                    > via st0.0
                    [Direct/0] 03:45:13
                    > via st0.1
fe80::210:dbff:feff:1000/128
                   *[Local/0] 03:45:42
                      Local via reth0.0
fe80::210:dbff:feff:1001/128
                   *[Local/0] 03:45:41
                      Local via reth1.0
```

Meaning The **show route** command lists active entries in the routing tables.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Dual-Stack Tunnels over an External Interface on page 6183
- Understanding VPN Tunnel Modes on page 6181

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 259

Configuring Traffic Selectors in Route-Based VPNs

- Understanding Traffic Selectors in Route-Based VPNs on page 6195
- Example: Configuring Traffic Selectors in a Route-Based VPN on page 6198
- Understanding Auto Route Insertion on page 6213
- Understanding Traffic Selectors and Overlapping IP Addresses on page 6213

Understanding Traffic Selectors in Route-Based VPNs

A traffic selector (also known as a *proxy ID* in IKEv1) is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses. With this feature, you can define a traffic selector within a specific route-based VPN, which can result in multiple Phase 2 IPsec security associations (SAs). Only traffic that conforms to a traffic selector is permitted through an SA.

- Traffic Selector Configuration on page 6195
- Traffic Selector Flexible Matches on page 6196
- Multiple Tunnels for Traffic Selector Configuration on page 6197
- Limitations on page 6198

Traffic Selector Configuration

To configure a traffic selector, use the **traffic-selector** configuration statement at the [**edit security ipsec vpn** *vpn-name*] hierarchy level. The traffic selector is defined with the mandatory **local-ip** *ip-address* and **remote-ip** *ip-address* statements. The CLI operational command **show security ipsec security-association detail** displays traffic selector information for SAs. The **show security ipsec security-association traffic-selector** ***traffic-selector-name*** CLI command displays information for a specified traffic selector.

For a given traffic selector, a single address or subnetwork is specified for the local and remote addresses. Traffic selectors can be configured with IPv4 or IPv6 addresses. Address books cannot be used to specify local or remote addresses.

Multiple traffic selectors can be configured for the same VPN. A maximum of 200 traffic selectors can be configured for each VPN. Traffic selectors can be used with IPv4-in-IPv4, IPv4-in-IPv6, IPv6-in-IPv6, or IPv6-in-IPv4 tunnel modes.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Traffic selectors on AutoVPN hubs can only be configured with IPv4 addresses. IPv4-in-IPv6, IPv6-in-IPv6, or IPv6-in-IPv4 tunnel modes are not supported for AutoVPN with traffic selectors.

When traffic selectors are configured, static routes are automatically added during configuration processing or when traffic selectors are negotiated; this process is known as auto route insertion (ARI). These routes might conflict with those that are populated through routing protocols. We recommend that you do not configure routing protocols on st0 interfaces that are bound to VPNs where traffic selectors are configured.

When a traffic selector is deleted, all corresponding IPsec SAs, routes, and tunnel sessions are cleared. This might affect traffic passing through these tunnels.

When a traffic selector is modified, deleted, or added, traffic selectors that follow it in the configuration are affected. The tunnels, SAs, and routes are cleared and reinstalled. Traffic selectors that precede the new or modified traffic selector in the configuration are unaffected.

For example, three traffic selectors are configured for the same VPN in the following order:

1. ts-red
2. ts-blue
3. ts-green

Changes in the traffic selector configuration have the following results:

Action	Result
Modify local or remote IP address in ts-blue.	Tunnels, SAs, and routes for ts-blue and ts-green are cleared. Tunnel, SA, and route for ts-red are not affected.
Delete ts-blue.	Tunnel, SA, and route for ts-green are cleared. Tunnel, SA, and route for ts-red are not affected.
Insert ts-white after ts-blue.	Tunnels, SAs, and routes for ts-white and ts-green are cleared. Tunnels, SAs, and routes for ts-red and ts-blue are not affected.

Traffic Selector Flexible Matches

During IKE negotiation, the responder can accept from the initiator a proposed traffic selector that is a subset of the traffic selector configured on the responder. There can be a wide subnetwork configured in a traffic selector on the hub and a narrow portion of

Copyright © 2016, Juniper Networks, Inc.

the subnetwork configured in a traffic selector on each spoke. For example, consider the following traffic selectors configured on an AutoVPN hub and spoke:

	Local IP	Remote IP
Spoke:	10.30.1.0/24	10.40.1.0/24
Hub:	10.40.0.0/16	10.30.0.0/16

The initiator (the spoke) negotiates with the responder (the hub) using the /24 netmask. Because the traffic selector IP addresses configured on the spoke are subsets of the traffic selector configured on the hub, the hub accepts the negotiation. The /24 subnetworks are used to match traffic for the negotiated tunnel and the 10.30.1.0/24 static route to the spoke is installed through ARI in the hub's routing table.

Traffic selector flexible matches are supported for both IKEv1 and IKEv2.

Multiple Tunnels for Traffic Selector Configuration

A single traffic selector configuration can result in multiple tunnels. In Figure 260, subnetworks 10.1.0.0/16 and 10.2.0.0/16 are behind device A while subnetwork 192.168.0.0/16 is behind device B.

Figure 260: Multiple Tunnels for Traffic Selector Configuration

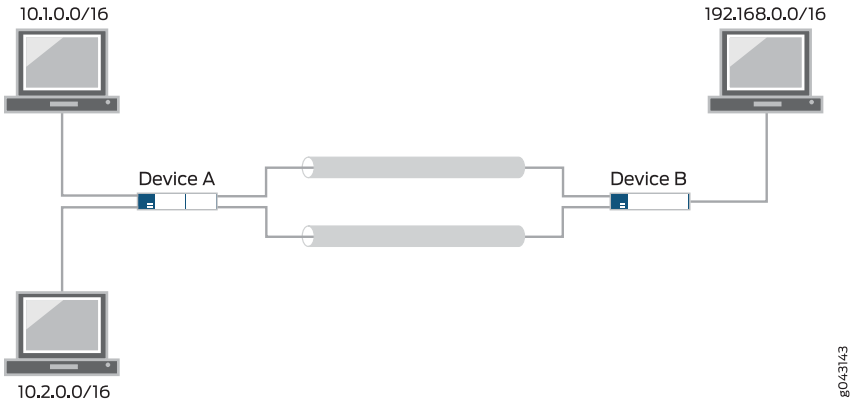

Two traffic selectors are configured on device A and one traffic selector is configured on device B, as follows:

	Local IP address	Remote IP address
Device A traffic selectors		
TS1:	10.1.0.0/16	192.168.0.0/16
TS2:	10.2.0.0/16	192.168.0.0/16

Copyright © 2016, Juniper Networks, Inc.

	Device B traffic selector	
TS:	192.168.0.0/16	10.0.0.0/8

Device A initiates two tunnels with device B using its proposed traffic selectors. Both of device A's proposed traffic selectors match the configured traffic selector on device B. Device B creates two different tunnels for its single traffic selector configuration.

Limitations

Traffic selectors cannot be configured with the following features:

- Policy-based VPNs
- Shared IKE IDs
- VPN monitoring
- Different address families configured for the local and remote IP addresses
- A remote address of 0.0.0.0/0 (IPv4) or 0::0 (IPv6)
- IKEv2 site-to-site VPN
- Dynamic routing protocols configured on st0 interfaces

Related Documentation

- Understanding Auto Route Insertion on page 6213
- Understanding AutoVPN with Traffic Selectors on page 6607
- Understanding VPN Tunnel Modes on page 6181
- Understanding Traffic Selectors and Overlapping IP Addresses on page 6213

Example: Configuring Traffic Selectors in a Route-Based VPN

This example shows how to configure traffic selectors for a route-based VPN.

- Requirements on page 6198
- Overview on page 6198
- Configuration on page 6201
- Verification on page 6210

Requirements

Before you begin, read "Understanding Traffic Selectors in Route-Based VPNs" on page 6195.

Overview

This example configures traffic selectors to allow traffic to flow between subnetworks on SRX_A and subnetworks on SRX_B.

Table 523 shows the traffic selectors used in this example. Traffic selectors are configured with other Phase 2 options (shown in Table 525).

Copyright © 2016, Juniper Networks, Inc.

Table 523: Traffic Selector Configurations

SRX_A			SRX_B		
Traffic Selector Name	Local IP	Remote IP	Traffic Selector Name	Local IP	Remote IP
TS1-ipv6	2001:db8:10::0/64	2001:db8:20::0/64	TS1-ipv6	2001:db8:20::0/64	2001:db8:10::0/64
TS2-ipv4	192.168.10.0/24	192.168.0.0/16	TS2-ipv4	192.168.0.0/16	192.168.10.0/24
TS3-ipv4	192.168.10.0/24	192.168.20.0/24	TS3-ipv4	192.168.20.0/24	192.168.10.0/24

Table 524 shows the Phase 1 options used in this example. The Phase 1 option configuration on each device includes an IKE gateway configuration to the IPv6 peer.

Table 524: Phase 1 Options for Traffic Selector Configurations

Option	SRX_A	SRX_B
IKE proposal	phase1_psk_proposal	phase1_psk_proposal
Authentication method	preshared keys	preshared keys
DH group	group2	group2
Authentication algorithm	SHA 1	SHA 1
Encryption algorithm	3DES CBC	3DES CBC
Lifetime	180 seconds	180 seconds
IKE policy	ike_psk_policy	ike_psk_policy
Mode	main	main
Proposal	phase1_psk_proposal	phase1_psk_proposal
Preshared key	ASCII text	ASCII text
IKE gateway	ike-gateway-to-he-srx	ike-gateway-to-branch-srx
IKE policy	ike_psk_policy	ike_psk_policy
Gateway address	2001:db8:2000::2	2001:db8:2000::1
External interface	ge-0/0/1.0	ge-0/0/1.0
Local address	2001:db8:2000::1	2001:db8:2000::2
IKE version	v1-only (default)	v1-only

Copyright © 2016, Juniper Networks, Inc.

Table 525 shows the Phase 2 options used in this example. Traffic selectors shown in Table 523 are configured with the Phase 2 options.

Table 525: Phase 2 Options for Traffic Selector Configurations

Option	SRX_A	SRX_B
IPsec proposal	phase2-proposal	phase2-proposal
Protocol	ESP	ESP
Authentication algorithm	HMAC SHA-1 96	HMAC SHA-1 96
Encryption algorithm	DES CBC	DES CBC
IPsec policy	ipsec-policy	ipsec-policy
Proposal	phase2-proposal	phase2-proposal
VPN	ipsec-vpn-to-he-srx	ipsec-vpn-to-branch-srx
Bind interface	st0.1	st0.1
IKE gateway	ike-gateway-to-he-srx	ike-gateway-to-branch-srx
IPsec policy	ipsec-policy	ipsec-policy

NOTE: On both devices, flow-based processing of IPv6 traffic must be enabled with the **mode flow-based** configuration option at the **[edit security forwarding-options family inet6]** hierarchy level.

Topology

In Figure 261, an IPv6 VPN tunnel carries both IPv4 and IPv6 traffic between the SRX_A and SRX_B devices. That is, the tunnel operates in both IPv4-in-IPv6 and IPv6-in-IPv6 tunnel modes.

Figure 261: Traffic Selector Configuration Example

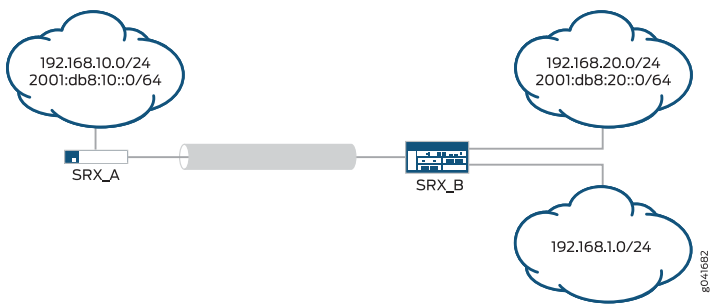

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Configuring SRX_A on page 6201
- Configuring SRX_B on page 6205

Configuring SRX_A

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet6 address 2001:db8:2000::1/64
set interfaces st0 unit 1 family inet
set interfaces st0 unit 1 family inet6
set interfaces ge-1/0/1 unit 0 family inet address 192.168.10.1/24
set interfaces ge-1/0/1 unit 0 family inet6 address 10::1/64
set security ike proposal phase1_psk_proposal authentication- method pre-shared-keys
set security ike proposal phase1_psk_proposal dh-group group2
set security ike proposal phase1_psk_proposal authentication- algorithm sha1
set security ike proposal phase1_psk_proposal encryption-algorithm 3des-cbc
set security ike proposal phase1_psk_proposal lifetime-seconds 180
set security ike policy ike_psk_policy mode main
set security ike policy ike_psk_policy proposals phase1_psk_proposal
set security ike policy ike_psk_policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike-gateway-to-he-srx ike-policy ike_psk_policy
set security ike gateway ike-gateway-to-he-srx address 2001:db8:2000::2
set security ike gateway ike-gateway-to-he-srx external-interface ge-0/0/1.0
set security ike gateway ike-gateway-to-he-srx local-address 2001:db8:2000::1
set security ipsec proposal phase2-proposal protocol esp
set security ipsec proposal phase2-proposal authentication- algorithm hmac-sha1-96
set security ipsec proposal phase2-proposal encryption-algorithm 3des-cbc
set security ipsec policy ipsec_policy proposals phase2-proposal
set security ipsec vpn ipsec-vpn-to-he-srx bind-interface st0.1
set security ipsec vpn ipsec-vpn-to-he-srx ike ipsec-policy ipsec_policy
set security ipsec vpn ipsec-vpn-to-he-srx ike gateway ike- gateway-to-he-srx
set security ipsec vpn ipsec-vpn-to-he-srx traffic-selector TS1- ipv6 local-ip
  2001:db8:10::0/64 remote-ip 2001:db8:20::0/64
set security ipsec vpn ipsec-vpn-to-he-srx traffic-selector TS2- ipv4 local-ip
  192.168.10.0/24 remote-ip 20.1.0.0/16
set security ipsec vpn ipsec-vpn-to-he-srx traffic-selector TS3- ipv4 local-ip
  192.168.10.0/24 remote-ip 192.168.20.0/24
set security forwarding-options family inet6 mode flow-based
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-1/0/1.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces st0.1
set security policies default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure traffic selectors:

1. Configure the external interface.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet6 address 2001:db8:2000::1/64
```

2. Configure the secure tunnel interface.

```
[edit interfaces]
user@host# set st0 unit 1 family inet
user@host# set st0 unit 1 family inet6
```

3. Configure the internal interface.

```
[edit interfaces]
user@host# set ge-1/0/1 unit 0 family inet address 192.168.10.1/24
user@host# set ge-1/0/1 unit 0 family inet6 address 10::1/64
```

4. Configure Phase 1 options.

```
[edit security ike proposal phase1_psk_proposal]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm 3des-cbc
user@host# set lifetime-seconds 180

[edit security ike policy ike_psk_policy]
user@host# set mode main
user@host# set proposals phase1_psk_proposal
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security ike gateway ike-gateway-to-he-srx]
user@host# set ike-policy ike_psk_policy
user@host# set address 2001:db8:2000::2
user@host# set external-interface ge-0/0/1.0
user@host# set local-address 2001:db8:2000::1
```

5. Configure Phase 2 options.

```
[edit security ipsec proposal phase2-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm 3des-cbc

[edit security ipsec policy ipsec_policy]
user@host# set proposals phase2-proposal

[edit security ipsec vpn ipsec-vpn-to-he-srx]
user@host# set bind-interface st0.1
user@host# set ike gateway ike-gateway-to-he-srx
user@host# set ike ipsec-policy ipsec_policy
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set traffic-selector TS1-ipv6 local-ip 2001:db8:10::0/64 remote-ip
  2001:db8:20::0/64
user@host# set traffic-selector TS2-ipv4 local-ip 192.168.10.0/24 remote-ip
  20.1.0.0/16
user@host# set traffic-selector TS3-ipv4 local-ip 192.168.10.0/24 remote-ip
  192.168.20.0/24
```

6. Enable IPv6 flow-based forwarding.

```
[edit security forwarding-options]
user@host# set family inet6 mode flow-based
```

7. Configure security zones and the security policy.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-1/0/1.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces ge-0/0/1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show security forwarding-options**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces
  ge-0/0/1 {
    unit 0 {
      family inet6 {
        address 2001:db8:2000::1/64;
      }
    }
  }
  ge-1/0/1 {
    unit 0 {
      family inet {
        address 192.168.10.1/24;
      }
      family inet6 {
        address 10::1/64;
      }
    }
  }
  st0 {
    unit 1 {
      family inet;
      family inet6;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
[edit]
user@host# show security ike
proposal phase1_psk_proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 180;
}
policy ike_psk_policy {
    mode main;
    proposals phase1_psk_proposal;
    pre-shared-key ascii-text
    "$ABC123"; ## SECRET-DATA
}
gateway ike-gateway-to-he-srx {
    ike-policy ike_psk_policy;
    address 2001:db8:2000::2;
    external-interface ge-0/0/1.0;
    local-address 2001:db8:2000::1;
}
[edit]
user@host# show security ipsec
proposal phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
}
policy ipsec_policy {
    proposals phase2-proposal;
}
vpn ipsec-vpn-to-he-srx {
    bind-interface st0.1;
    ike {
        ipsec-policy ipsec_policy;
        gateway ike-gateway-to-he-srx;
    }
    traffic-selector TS1-ipv6 {
        local-ip 2001:db8:10::0/64;
        remote-ip 2001:db8:20::0/64;
    }
    traffic-selector TS2-ipv4 {
        local-ip 192.168.10.0/24;
        remote-ip 20.1.0.0/16;
    }
    traffic-selector TS3-ipv4 {
        local-ip 192.168.10.0/24;
        remote-ip 192.168.20.0/24;
    }
}
[edit]
user@host# show security forwarding-options
family {
    inet6 {
```


Copyright © 2016, Juniper Networks, Inc.

```
        mode flow-based;
      }
    }
    [edit]
    user@host# show security zones
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        ge-1/0/1.0;
      }
    }
    security-zone untrust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        st0.1;
        ge-0/0/1.0;
      }
    }
    [edit]
  user@host# show security policies
    default-policy {
      permit-all;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring SRX_B

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet6 address 2001:db8:2000::2/64
set interfaces st0 unit 1 family inet
set interfaces st0 unit 1 family inet6
set interfaces ge-1/0/1 unit 0 family inet address 192.168.20.1/24
set interfaces ge-1/0/1 unit 0 family inet6 address 20::1/64
set interfaces ge-1/1/1 unit 0 family inet address 192.168.1.1/24
set security ike proposal phase1_psk_proposal authentication-method pre-shared-keys
set security ike proposal phase1_psk_proposal dh-group group2
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ike proposal phase1_psk_proposal authentication-algorithm sha1
set security ike proposal phase1_psk_proposal encryption-algorithm 3des-cbc
set security ike proposal phase1_psk_proposal lifetime-seconds 180
set security ike policy ike_psk_policy mode main
set security ike policy ike_psk_policy proposals phase1_psk_proposal
set security ike policy ike_psk_policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike-gateway-to-branch-srx ike-policy ike_psk_policy
set security ike gateway ike-gateway-to-branch-srx address 2001:db8:2000::1
set security ike gateway ike-gateway-to-branch-srx external-interface ge-0/0/1.0
set security ike gateway ike-gateway-to-branch-srx local-address 2001:db8:2000::2
set security ipsec proposal phase2-proposal protocol esp
set security ipsec proposal phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal phase2-proposal encryption-algorithm 3des-cbc
set security ipsec policy ipsec_policy proposals phase2-proposal
set security ipsec vpn ipsec-vpn-to-branch-srx bind-interface st0.1
set security ipsec vpn ipsec-vpn-to-branch-srx ike ipsec-policy ipsec_policy
set security ipsec vpn ipsec-vpn-to-branch-srx ike gateway ike-gateway-to-branch-srx
set security ipsec vpn ipsec-vpn-to-branch-srx traffic-selector TS1-ipv6 local-ip
  2001:db8:20::0/64 remote-ip 2001:db8:10::0/64
set security ipsec vpn ipsec-vpn-to-branch-srx traffic-selector TS2-ipv4 local-ip
  20.1.0.0/16 remote-ip 192.168.10.0/24
set security ipsec vpn ipsec-vpn-to-branch-srx traffic-selector TS3-ipv4 local-ip
  192.168.20.0/24 remote-ip 192.168.10.0/24
set security forwarding-options family inet6 mode flow-based
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-1/0/1.0
set security zones security-zone trust interfaces ge-1/1/1.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.1
set security zones security-zone untrust interfaces ge-0/0/1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure traffic selectors:

1. Configure the external interface.

   ```
   [edit interfaces]
   user@host# set ge-0/0/1 unit 0 family inet6 address 2001:db8:2000::2/64
   ```

2. Configure the secure tunnel interface.

   ```
   [edit interfaces]
   user@host# set st0 unit 1 family inet
   user@host# set st0 unit 1 family inet6
   ```

3. Configure the internal interfaces.

   ```
   [edit interfaces]
   user@host# set ge-1/0/1 unit 0 family inet address 192.168.20.1/24
   user@host# set ge-1/0/1 unit 0 family inet6 address 20::1/64
   user@host# set ge-1/1/1 unit 0 family inet address 192.168.1.1/24
   ```


Copyright © 2016, Juniper Networks, Inc.

4. Configure Phase 1 options.

```
[edit security ike proposal phase1_psk_proposal]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm 3des-cbc
user@host# set lifetime-seconds 180

[edit security ike policy ike_psk_policy]
user@host# set mode main
user@host# set proposals phase1_psk_proposal
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security ike gateway ike-gateway-to-branch-srx]
user@host# set ike-policy ike_psk_policy
user@host# set address 2001:db8:2000::1
user@host# set external-interface ge-0/0/1.0
user@host# set local-address 2001:db8:2000::2
```

5. Configure Phase 2 options.

```
[edit security ipsec proposal phase2-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm 3des-cbc

[edit security ipsec policy ipsec_policy]
user@host# set proposals phase2-proposal

[edit security ipsec vpn ipsec-vpn-to-branch-srx]
user@host# set bind-interface st0.1
user@host# set ike gateway ike-gateway-to-branch-srx
user@host# set ike ipsec-policy ipsec_policy
user@host# set traffic-selector TS1-ipv6 local-ip 2001:db8:20::0/64 remote-ip
  2001:db8:10::0/64
user@host# set traffic-selector TS2-ipv4 local-ip 20.1.0.0/16 remote-ip
  192.168.10.0/24
user@host# set traffic-selector TS3-ipv4 local-ip 192.168.20.0/24 remote-ip
  192.168.10.0/24
```

6. Enable IPv6 flow-based forwarding.

```
[edit security forwarding-options]
user@host# set family inet6 mode flow-based
```

7. Configure security zones and the security policy.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-1/0/1.0
user@host# set interfaces ge-1/1/1.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces ge-0/0/1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show security forwarding-options**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show interfaces
  ge-0/0/1 {
    unit 0 {
      family inet6 {
        address 2001:db8:2000::2/64;
      }
    }
  }
  ge-1/0/1 {
    unit 0 {
      family inet {
        address 192.168.20.1/24;
      }
      family inet6 {
        address 20::1/64;
      }
    }
  }
  ge-1/1/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  st0 {
    unit 1 {
      family inet;
      family inet6;
    }
  }
  [edit]
  user@host# show security ike
  proposal phase1_psk_proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 180;
  }
  policy ike_psk_policy {
```

 Copyright © 2016, Juniper Networks, Inc.

```
    mode main;
    proposals phase1_psk_proposal;
    pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway ike-gateway-to-branch-srx {
    ike-policy ike_psk_policy;
    address 2001:db8:2000::1;
    external-interface ge-0/0/1.0;
    local-address 2001:db8:2000::2;
}
[edit]
user@host# show security ipsec
proposal phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
}
policy ipsec_policy {
    proposals phase2-proposal;
}
vpn ipsec-vpn-to-branch-srx {
    bind-interface st0.1;
    ike {
        ipsec-policy ipsec_policy;
        gateway ike-gateway-to-branch-srx;
    }
    traffic-selector TS1-ipv6 {
        local-ip 2001:db8:20::0/64;
        remote-ip 2001:db8:10::0/64;
    }
    traffic-selector TS2-ipv4 {
        local-ip 20.1.0.0/16;
        remote-ip 192.168.10.0/24;
    }
    traffic-selector TS3-ipv4 {
        local-ip 192.168.20.0/24;
        remote-ip 192.168.10.0/24;
    }
}
[edit]
user@host# show security forwarding-options
family {
    inet6 {
        mode flow-based;
    }
}
[edit]
user@host# show security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    interfaces {
      ge-1/0/1.0;
      ge-1/1/1.0;
    }
  }
  security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      st0.1;
      ge-0/0/1.0;
    }
  }
  [edit]
user@host# show security policies
  default-policy {
    permit-all;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying IPsec Phase 2 Status on page 6210
- Verifying Routes on page 6212

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command.

```
user@host> show security ipsec security-associations
  Total active tunnels: 3
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <268173313 ESP:des/ sha1 3d75aeff 2984/ unlim -  root 500   2001:db8:2000::2
  >268173313 ESP:des/ sha1 a468fece 2984/ unlim -  root 500   2001:db8:2000::2
  <268173316 ESP:des/ sha1 417f3cea 3594/ unlim -  root 500   2001:db8:2000::2
  >268173316 ESP:des/ sha1 a4344027 3594/ unlim -  root 500   2001:db8:2000::2
  <268173317 ESP:des/ sha1 cc9fb573 3556/ unlim -  root 500   2001:db8:2000::2
  >268173317 ESP:des/ sha1 a4bde69b 3556/ unlim -  root 500   2001:db8:2000::2
```

From operational mode, enter the **show security ipsec security-associations detail** command.

```
user@host> show security ipsec security-associations detail
```


Copyright © 2016, Juniper Networks, Inc.

```
ID: 268173313 Virtual-system: root, VPN Name: ipsec-vpn-to-he-srx
Local Gateway: 2001:db8:2000::1, Remote Gateway: 2001:db8:2000::2
Traffic Selector Name: TS1-ipv6
Local Identity: ipv6(2001:db8:10::-2001:db8:10::ffff:ffff:ffff:ffff)
Remote Identity: ipv6(2001:db8:20::-2001:db8:20::ffff:ffff:ffff:ffff)
Version: IKEv1
  DF-bit: clear
  Bind-interface: st0.1

Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: c608b29
Tunnel Down Reason: SA not initiated
  Direction: inbound, SPI: 3d75aeff, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 2976 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2354 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64

  Direction: outbound, SPI: a468fece, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 2976 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2354 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64

ID: 268173316 Virtual-system: root, VPN Name: ipsec-vpn-to-he-srx
Local Gateway: 2001:db8:2000::1, Remote Gateway: 2001:db8:2000::2
Traffic Selector Name: TS2-ipv4
Local Identity: ipv4(192.168.10.0-192.168.10.255)
Remote Identity: ipv4(192.168.20.0-192.168.20.255)
Version: IKEv1
  DF-bit: clear
  Bind-interface: st0.1

Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: c608b29
Tunnel Down Reason: SA not initiated
  Direction: inbound, SPI: 417f3cea, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 3586 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2948 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64

  Direction: outbound, SPI: a4344027, AUX-SPI: 0
                            , VPN Monitoring: -
  Hard lifetime: Expires in 3586 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2948 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64

ID: 268173317 Virtual-system: root, VPN Name: ipsec-vpn-to-he-srx
Local Gateway: 2000::1, Remote Gateway: 2000::2
Traffic Selector Name: TS3-ipv4
```

Copyright © 2016, Juniper Networks, Inc.

```
Local Identity: ipv4(192.168.10.0-192.168.10.255)
Remote Identity: ipv4(192.168.20.0-192.168.20.255)
Version: IKEv1
  DF-bit: clear
  Bind-interface: st0.1

Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: c608b29
Tunnel Down Reason: SA not initiated
  Direction: inbound, SPI: cc9fb573, AUX-SPI: 0
                              , VPN Monitoring: -
  Hard lifetime: Expires in 3548 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2925 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64

  Direction: outbound, SPI: a4bde69b, AUX-SPI: 0
                              , VPN Monitoring: -
  Hard lifetime: Expires in 3548 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 2925 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: des-cbc
  Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the peer devices.

Verifying Routes

Purpose Verify active routes

Action From operational mode, enter the **show route** command.

```
user@host> show route
inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16        *[Static/5] 00:00:32
                    > via st0.1
192.168.20.0/24        *[Static/5] 00:00:32
                    > via st0.1
2001:db8:20::/64       *[Static/5] 00:00:34
                    > via st0.1
```

Meaning The **show route** command lists active entries in the routing tables. Routes to the remote IP address configured in each traffic selector should be present with the correct st0 interface.

Related Documentation

- Understanding Traffic Selectors in Route-Based VPNs on page 6195

Copyright © 2016, Juniper Networks, Inc.

Understanding Auto Route Insertion

Auto route insertion (ARI) automatically inserts a static route for the remote network and hosts protected by a remote tunnel endpoint. A route is created based on the remote IP address configured in the traffic-selector. In the case of traffic selectors, the configured remote address is inserted as a route in the routing instance associated with the st0 interface that is bound to the VPN.

NOTE: **Routing protocols and traffic selector configuration are mutually exclusive ways of steering traffic to a tunnel. ARI routes might conflict with routes that are populated through routing protocols. Therefore, you should not configure routing protocols on an st0 interface that is bound to a VPN on which traffic selectors are configured.**

ARI is also known as reverse route insertion (RRI). ARI routes are inserted in the routing table as follows:

- If the **establish-tunnels immediately** option is configured at the [**edit security ipsec vpn *vpn-name***] hierarchy level, ARI routes are added after Phase 1 and Phase 2 negotiations are complete. Because a route is not added until SAs are established, a failed negotiation does not result in traffic being routed to a st0 interface that is down. An alternate or backup tunnel is used instead.
- If the **establish-tunnels immediately** option is not configured at the [**edit security ipsec vpn *vpn-name***] hierarchy level, ARI routes are added at configuration commit.
- An ARI route is not added if the configured or negotiated remote address in a traffic selector is 0.0.0.0/0 or 0::0.

The preference for the static ARI route is 5. This value is necessary to avoid conflict with similar routes that might be added by a routing protocol process. There is no configuration of the metric for the static ARI route.

Related Documentation

- Understanding Traffic Selectors in Route-Based VPNs on page 6195
- Understanding AutoVPN with Traffic Selectors on page 6607

Understanding Traffic Selectors and Overlapping IP Addresses

This section discusses overlapping IP addresses in traffic selector configurations.

- Overlapping IP Addresses in Different VPNs Bound to the Same st0 Interface on page 6214
- Overlapping IP Addresses in the Same VPN Bound to the Same st0 Interface on page 6214
- Overlapping IP Addresses in Different VPNs Bound to Different st0 Interfaces on page 6215

Copyright © 2016, Juniper Networks, Inc.

Overlapping IP Addresses in Different VPNs Bound to the Same st0 Interface

This scenario is not supported with traffic selectors. Traffic selectors cannot be configured on different VPNs that are bound to the same point-to-multipoint st0 interface, as shown in the following example:

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
   bind-interface st0.1;
}
vpn vpn-2 {
   bind-interface st0.1;
}
```

Overlapping IP Addresses in the Same VPN Bound to the Same st0 Interface

When overlapping IP addresses are configured for multiple traffic selectors in the same VPN, the first configured traffic selector that matches the packet determines the tunnel used for packet encryption.

In the following example, four traffic selectors (ts-1, ts-2, ts-3, and ts-4) are configured for the VPN (vpn-1), which is bound to the point-to-point st0.1 interface:

```
[edit]
user@host# show security ipsec vpn vpn-1
vpn vpn-1 {
   bind-interface st0.1;
   traffic-selector ts-1 {
      local-ip 192.168.5.0/24;
      remote-ip 10.1.5.0/24;
   }
   traffic-selector ts-2 {
      local-ip 192.168.0.0/16;
      remote-ip 10.1.0.0/16;
   }
   traffic-selector ts-3 {
      local-ip 172.16.0.0/16;
      remote-ip 10.2.0.0/16;
   }
   traffic-selector ts-4 {
      local-ip 172.16.5.0/24;
      remote-ip 10.2.5.0/24;
   }
}
```

A packet with a source address 192.168.5.5 and a destination address 10.1.5.10 matches traffic selectors ts-1 and ts-2. However, traffic selector ts-1 is the first configured match and the tunnel associated with ts-1 is used for packet encryption.

A packet with a source address 172.16.5.5 and a destination address 10.2.5.10 matches the traffic selectors ts-3 and ts-4. However, traffic selector ts-3 is the first configured match and the tunnel associated with traffic selector ts-3 is used for packet encryption.

Copyright © 2016, Juniper Networks, Inc.

Overlapping IP Addresses in Different VPNs Bound to Different st0 Interfaces

When overlapping IP addresses are configured for multiple traffic selectors in different VPNs that are bound to different point-to-point st0 interfaces, an st0 interface is first selected by the longest prefix match for a given packet. Within the VPN that is bound to the selected st0 interface, the traffic selector is then selected based on the first configured match for the packet.

In the following example, a traffic selector is configured in each of two VPNs. The traffic selectors are configured with the same local subnetwork but different remote subnetworks.

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
    bind-interface st0.1;
    traffic-selector ts-1 {
        local-ip 192.168.1.0/24;
        remote-ip 10.1.1.0/24;
    }
}
vpn vpn-2 {
    bind-interface st0.2;
    traffic-selector ts-2 {
        local-ip 192.168.1.0/24;
        remote-ip 10.2.2.0/24;
    }
}
```

Different remote subnetworks are configured in each traffic selector, therefore two different routes are added to the routing table. Route lookup uses the st0 interface bound to the appropriate VPN.

In the following example, a traffic selector is configured in each of two VPNs. The traffic selectors are configured with different remote subnetworks. The same local subnetwork is configured for each traffic selector, but different netmask values are specified.

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
    bind-interface st0.1;
    traffic-selector ts-1 {
        local-ip 192.168.0.0/8;
        remote-ip 10.1.1.0/24;
    }
}
vpn vpn-2 {
    bind-interface st0.2;
    traffic-selector ts-2 {
        local-ip 192.168.0.0/16;
        remote-ip 10.2.2.0/24;
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

A different remote subnetwork is configured in each traffic selector, therefore two different routes are added to the routing table. Route lookup uses the st0 interface bound to the appropriate VPN.

In the following example, traffic selectors are configured in each of two VPNs. The traffic selectors are configured with different local and remote subnetworks.

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
  bind-interface st0.1;
  traffic-selector ts-1 {
    local-ip 192.168.1.0/24;
    remote-ip 10.1.1.0/24;
  }
}
vpn vpn-2 {
  bind-interface st0.2;
  traffic-selector ts-2 {
    local-ip 172.16.1.0/24;
    remote-ip 10.2.2.0/24;
  }
}
```

In this case, the traffic selectors do not overlap. The remote subnetworks configured in the traffic selectors are different, therefore two different routes are added to the routing table. Route lookup uses the st0 interface bound to the appropriate VPN.

In the following example, a traffic selector is configured in each of two VPNs. The traffic selectors are configured with the same local subnetwork. The same remote subnetwork is configured for each traffic selector, but different netmask values are specified.

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
  bind-interface st0.1;
  traffic-selector ts-1 {
    local-ip 192.168.1.0/24;
    remote-ip 10.1.1.0/24;
  }
}
vpn vpn-2 {
  bind-interface st0.2;
  traffic-selector ts-2 {
    local-ip 192.168.1.0/24;
    remote-ip 10.1.0.0/16;
  }
}
```

Note that the **remote-ip** configured for ts-1 is 10.1.1.0/24 while the **remote-ip** configured for ts-2 is 10.1.0.0/16. For a packet destined to 10.1.1.1, route lookup selects the st0.1 interface as it has the longer prefix match. The packet is encrypted based on the tunnel corresponding to the st0.1 interface.

Copyright © 2016, Juniper Networks, Inc.

In some cases, valid packets can be dropped due to traffic selector traffic enforcement. In the following example, traffic selectors are configured in each of two VPNs. The traffic selectors are configured with different local subnetworks. The same remote subnetwork is configured for each traffic selector, but different netmask values are specified.

```
[edit]
user@host# show security ipsec
vpn vpn-1 {
   bind-interface st0.1;
   traffic-selector ts-1 {
      local-ip 192.168.1.0/24;
      remote-ip 10.1.1.0/24;
   }
}
vpn vpn-2 {
   bind-interface st0.2;
   traffic-selector ts-2 {
      local-ip 172.16.1.0/16;
      remote-ip 10.1.0.0/16;
   }
}
```

Two routes to 10.1.1.0 (10.1.1.0/24 via interface st0.1 and 10.1.0.0/16 via interface st0.2) are added to the routing table. A packet sent from source 172.16.1.1 to destination 10.1.1.1 matches the routing table entry for 10.1.1.0/24 via interface st0.1. However, the packet does not match the traffic specified by traffic selector ts-1 and is dropped.

NOTE: If multiple traffic selectors are configured with the same remote subnetwork and netmask, equal cost routes are added to the routing table. This case is not supported with traffic selectors as the route chosen cannot be predicted.

Related Documentation

- Understanding Traffic Selectors in Route-Based VPNs on page 6195

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 81

Configuring Policy-Based IPsec VPNs

- Configuring Policy-Based VPNs on page 6221

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 260

Configuring Policy-Based VPNs

- Understanding Policy-Based IPsec VPNs on page 6221
- Example: Configuring a Policy-Based VPN on page 6222

Understanding Policy-Based IPsec VPNs

For policy-based IPsec VPNs, a security policy specifies as its action the VPN tunnel to be used for transit traffic that meets the policy's match criteria. A VPN is configured independent of a policy statement. The policy statement refers to the VPN by name to specify the traffic that is allowed access to the tunnel. For policy-based VPNs, each policy creates an individual IPsec security association (SA) with the remote peer, each of which counts as an individual VPN tunnel. For example, if a policy contains a group source address and a group destination address, whenever one of the users belonging to the address set attempts to communicate with any one of the hosts specified as the destination address, a new tunnel is negotiated and established. Because each tunnel requires its own negotiation process and separate pair of SAs, the use of policy-based IPsec VPNs can be more resource-intensive than route-based VPNs.

Examples of where policy-based VPNs can be used:

- You are implementing a dial-up VPN.
- Policy-based VPNs allow you to direct traffic based on firewall policies.

NOTE: We recommend that you use route-based VPN when you want to configure a VPN between multiple remote sites. Route-based VPNs can provide the same capabilities as policy-based VPNs.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Route-Based VPN on page 6074
- Example: Configuring a Hub-and-Spoke VPN on page 6094
- Example: Configuring a Policy-Based VPN on page 6222

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring a Policy-Based VPN

This example shows how to configure a policy-based IPsec VPN to allow data to be securely transferred between a branch office and the corporate office.

- Requirements on page 6222
- Overview on page 6222
- Configuration on page 6226
- Verification on page 6235

Requirements

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

In this example, you configure a policy-based VPN for a branch office in Chicago, Illinois, because you do not need to conserve tunnel resources or configure many security policies to filter traffic through the tunnel. Users in the Chicago office will use the VPN to connect to their corporate headquarters in Sunnyvale, California.

Figure 262 shows an example of a policy-based VPN topology. In this topology, the SRX Series device is located in Sunnyvale, and an SSG Series device (or it can be another third-party device) is located in Chicago.

Copyright © 2016, Juniper Networks, Inc.

Figure 262: Policy-Based VPN Topology

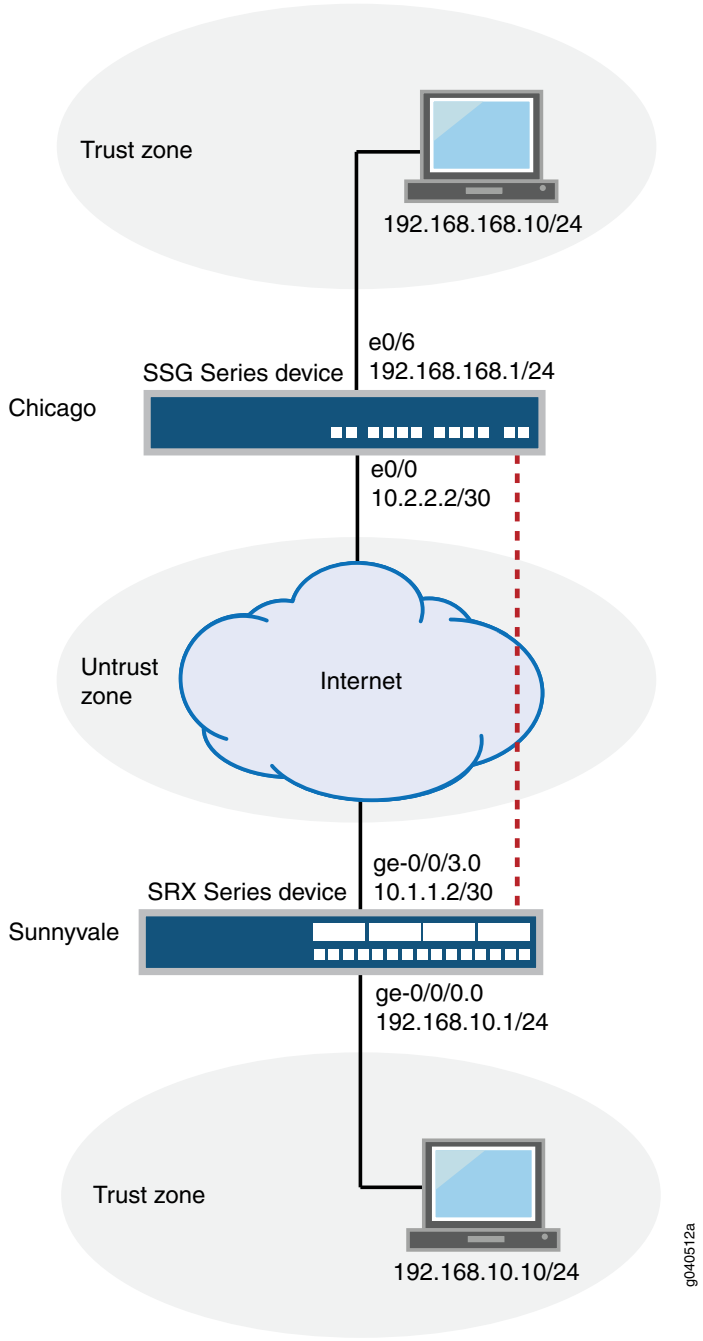

IKE IPsec tunnel negotiation occurs in two phases. In Phase 1, participants establish a secure channel in which to negotiate the IPsec security association (SA). In Phase 2, participants negotiate the IPsec SA for authenticating traffic that will flow through the

Copyright © 2016, Juniper Networks, Inc.

tunnel. Just as there are two phases to tunnel negotiation, there are two phases to tunnel configuration.

In this example, you configure interfaces, an IPv4 default route, security zones, and address books. Then you configure IKE Phase 1, IPsec Phase 2, security policy, and TCP-MSS parameters. See Table 526 through Table 530.

Table 526: Interface, Security Zone, and Address Book Information

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/0.0	192.168.10.1/24
	ge-0/0/3.0	10.1.1.2/30
Security zones	trust	• All system services are allowed. • The ge-0/0/0.0 interface is bound to this zone.
	untrust	• IKE is the only allowed system service. • The ge-0/0/3.0 interface is bound to this zone.
Address book entries	sunnyvale	• This address is an entry in the address book **book1**, which is attached to a zone called **trust**. • The address for this address book entry is 192.168.10.0/24.
	chicago	• This address is an entry in the address book **book2**, which is attached to a zone called **untrust**. • The address for this address book entry is 192.168.168.0/24.

Table 527: IKE Phase 1 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ike-phase1-proposal	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: aes-128-cbc
Policy	ike-phase1-policy	• Mode: main • Proposal reference: ike-phase1-proposal • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gw-chicago	• IKE policy reference: ike-phase1-policy • External interface: ge-0/0/3.0 • Gateway address: 10.2.2.2

Copyright © 2016, Juniper Networks, Inc.

Table 528: IPsec Phase 2 Configuration Parameters

Feature	Name	Configuration Parameters
Proposal	ipsec-phase2-proposal	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: aes-128-cbc
Policy	ipsec-phase2-policy	• Proposal reference: ipsec-phase2-proposal • PFS: Diffie-Hellman group2
VPN	ike-vpn-chicago	• IKE gateway reference: gw-chicago • IPsec policy reference: ipsec-phase2-policy

Table 529: Security Policy Configuration Parameters

Purpose	Name	Configuration Parameters

This security policy permits traffic from the trust zone to the untrust zone.	vpn-tr-untr	• Match criteria:
		  ◦ source-address sunnyvale
		  ◦ destination-address chicago
		  ◦ application any
		• Permit action: tunnel ipsec-vpn ike-vpn-chicago
		• Permit action: tunnel pair-policy vpn-untr-tr

This security policy permits traffic from the untrust zone to the trust zone.	vpn-untr-tr	• Match criteria:
		  ◦ source-address chicago
		  ◦ destination-address sunnyvale
		  ◦ application any
		• Permit action: tunnel ipsec-vpn ike-vpn-chicago
		• Permit action: tunnel pair-policy vpn-tr-untr

This security policy permits all traffic from the trust zone to the untrust zone.
 	
NOTE: You must put the vpn-tr-untr policy before the permit-any security policy. Junos OS performs a security policy lookup starting at the top of the list. If the permit-any policy comes before the vpn-tr-untr policy, all traffic from the trust zone will match the permit-any policy and be permitted. Thus, no traffic will ever match the vpn-tr-untr policy.
	permit-any	• Match criteria:
		  ◦ source-address any
		  ◦ source-destination any
		  ◦ application any
		• Action: permit

Copyright © 2016, Juniper Networks, Inc.

Table 530: TCP-MSS Configuration Parameters

Purpose	Configuration Parameters
TCP-MSS is negotiated as part of the TCP three-way handshake and limits the maximum size of a TCP segment to better fit the maximum transmission unit (MTU) limits on a network. This is especially important for VPN traffic, as the IPsec encapsulation overhead, along with the IP and frame overhead, can cause the resulting Encapsulating Security Payload (ESP) packet to exceed the MTU of the physical interface, thus causing fragmentation. Fragmentation results in increased use of bandwidth and device resources. NOTE: We recommend a value of 1350 as the starting point for most Ethernet-based networks with an MTU of 1500 or greater. You might need to experiment with different TCP-MSS values to obtain optimal performance. For example, you might need to change the value if any device in the path has a lower MTU, or if there is any additional overhead such as PPP or Frame Relay.	MSS value: 1350

Configuration

Configuring Basic Network, Security Zone, and Address Book Information

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
set security zones security-zone untrust interfaces ge-0/0/3.0
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone trust interfaces ge-0/0/0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security address-book book1 address sunnyvale 192.168.10.0/24
set security address-book book1 attach zone trust
set security address-book book2 address chicago 192.168.168.0/24
set security address-book book2 attach zone untrust
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure basic network, security zone, and address book information:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@host# set interfaces ge-0/0/0 unit 0 family inet address 192.168.10.1/24
   user@host# set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.2/30
   ```

2. Configure static route information.

   ```
   [edit]
   user@host# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
   ```

3. Configure the untrust security zone.

Copyright © 2016, Juniper Networks, Inc.

```
[edit ]
user@host# edit security zones security-zone untrust
```

4. Assign an interface to the security zone.

```
[edit security zones security-zone untrust]
user@host# set interfaces ge-0/0/3.0
```

5. Specify allowed system services for the security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@host# edit security zones security-zone trust
```

7. Assign an interface to the security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/0.0
```

8. Specify allowed system services for the security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

9. Create an address book and attach it to a zone.

```
[edit security address-book book1]
user@host# set address sunnyvale 192.168.10.0/24
user@host# set attach zone trust
```

10. Create another address book and attach it to a zone.

```
[edit security address-book book2]
user@host# set address chicago 192.168.168.0/24
user@host# set attach zone untrust
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, and **show security address-book** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.10.1/24;
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 10.1.1.2/30
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}

[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 10.1.1.1;
}

[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      ike;
    }
  }
  interfaces {
    ge-0/0/3.0;
  }
}
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
  }
  interfaces {
    ge-0/0/0.0;
  }
}
[edit]
user@host# show security address-book
book1 {
  address sunnyvale 192.168.10.0/24;
  attach {
    zone trust;
  }
}
  book2 {
    address chicago 192.168.168.0/24;
    attach {
      zone untrust;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text $ABC123
set security ike gateway gw-chicago external-interface ge-0/0/3.0
set security ike gateway gw-chicago ike-policy ike-phase1-policy
set security ike gateway gw-chicago address 10.2.2.2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike-phase1-proposal]
   user@host# set encryption-algorithm aes-128-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike]
   user@host# set policy ike-phase1-policy
   ```

7. Set the IKE Phase 1 policy mode.

   ```
   [edit security ike policy ike-phase1-policy]
   user@host# set mode main
   ```

8. Specify a reference to the IKE proposal.

   ```
   [edit security ike policy ike-phase1-policy]
   user@host# set proposals ike-phase1-proposal
   ```

9. Define the IKE Phase 1 policy authentication method.

   ```
   [edit security ike policy ike-phase1-policy]
   user@host# set pre-shared-key ascii-text $ABC123
   ```

10. Create an IKE Phase 1 gateway and define its external interface.

    ```
    [edit security ike]
    ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set gateway gw-chicago external-interface ge-0/0/3.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw-chicago]
user@host# set ike-policy ike-phase1-policy
```

12. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike gateway gw-chicago]
user@host# set gateway gw-chicago external-interface ge-0/0/3.0
```

13. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw-chicago]
user@host# set ike-policy ike-phase1-policy
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security ike
proposal ike-phase1-proposal {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm sha1;
   encryption-algorithm aes-128-cbc;
}
policy ike-phase1-policy {
   mode main;
   proposals ike-phase1-proposal;
   pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
}
gateway gw-chicago {
   ike-policy ike-phase1-policy;
   address 10.2.2.2;
   external-interface ge-0/0/3.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy ipsec-phase2-policy proposals ipsec-phase2-proposal
set security ipsec policy ipsec-phase2-policy perfect-forward-secrecy keys group2
set security ipsec vpn ike-vpn-chicago ike gateway gw-chicago
set security ipsec vpn ike-vpn-chicago ike ipsec-policy ipsec-phase2-policy
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

```
[edit]
user@host# set security ipsec proposal ipsec-phase2-proposal
```

2. Specify the IPsec Phase 2 proposal protocol.

```
[edit security ipsec proposal ipsec-phase2-proposal]
user@host# set protocol esp
```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

```
[edit security ipsec proposal ipsec-phase2-proposal]
user@host# set authentication-algorithm hmac-sha1-96
```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

```
[edit security ipsec proposal ipsec-phase2-proposal]
user@host# set encryption-algorithm aes-128-cbc
```

5. Create the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set policy ipsec-phase2-policy
```

6. Specify the IPsec Phase 2 proposal reference.

```
[edit security ipsec policy ipsec-phase2-policy]
user@host# set proposals ipsec-phase2-proposal
```

7. Specify IPsec Phase 2 PFS to use Diffie-Hellman group 2.

```
[edit security ipsec policy ipsec-phase2-policy]
user@host# set perfect-forward-secrecy keys group2
```

8. Specify the IKE gateway.

```
[edit security ipsec]
user@host# set vpn ike-vpn-chicago ike gateway gw-chicago
```

9. Specify the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set vpn ike-vpn-chicago ike ipsec-policy ipsec-phase2-policy
```

Results

From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security ipsec
proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
```

Copyright © 2016, Juniper Networks, Inc.

```
}
policy ipsec-phase2-policy {
  perfect-forward-secrecy {
    keys group2;
  }
  proposals ipsec-phase2-proposal;
}
vpn ike-vpn-chicago {
  ike {
    gateway gw-chicago;
    ipsec-policy ipsec-phase2-policy;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone untrust policy vpn-tr-untr match
  source-address sunnyvale
set security policies from-zone trust to-zone untrust policy vpn-tr-untr match
  destination-address chicago
set security policies from-zone trust to-zone untrust policy vpn-tr-untr match application
  any
set security policies from-zone trust to-zone untrust policy vpn-tr-untr then permit tunnel
  ipsec-vpn ike-vpn-chicago
set security policies from-zone trust to-zone untrust policy vpn-tr-untr then permit tunnel
  pair-policy vpn-untr-tr
set security policies from-zone untrust to-zone trust policy vpn-untr-tr match
  source-address chicago
set security policies from-zone untrust to-zone trust policy vpn-untr-tr match
  destination-address sunnyvale
set security policies from-zone untrust to-zone trust policy vpn-untr-tr match application
  any
set security policies from-zone untrust to-zone trust policy vpn-untr-tr then permit tunnel
  ipsec-vpn ike-vpn-chicago
set security policies from-zone untrust to-zone trust policy vpn-untr-tr then permit tunnel
  pair-policy vpn-tr-untr
set security policies from-zone trust to-zone untrust policy permit-any match
  source-address any
set security policies from-zone trust to-zone untrust policy permit-any match
  destination-address any
set security policies from-zone trust to-zone untrust policy permit-any match application
  any
set security policies from-zone trust to-zone untrust policy permit-any then permit
insert security policies from-zone trust to-zone untrust policy vpn-tr-untr before policy
  permit-any
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies:

1. Create the security policy to permit traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy vpn-tr-untr match source-address sunnyvale
   user@host# set policy vpn-tr-untr match destination-address chicago
   user@host# set policy vpn-tr-untr match application any
   user@host# set policy vpn-tr-untr then permit tunnel ipsec-vpn ike-vpn-chicago
   user@host# set policy vpn-tr-untr then permit tunnel pair-policy vpn-untr-tr
   ```

2. Create the security policy to permit traffic from the untrust zone to the trust zone.

   ```
   [edit security policies from-zone untrust to-zone trust]
   user@host# set policy vpn-untr-tr match source-address sunnyvale
   user@host# set policy vpn-untr-tr match destination-address chicago
   user@host# set policy vpn-untr-tr match application any
   user@host# set policy vpn-untr-tr then permit tunnel ipsec-vpn ike-vpn-chicago
   user@host# set policy vpn-untr-tr then permit tunnel pair-policy vpn-tr-untr
   ```

3. Create the security policy to permit traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy permit-any match source-address any
   user@host# set policy vpn-untr-tr match destination-address any
   user@host# set policy vpn-untr-tr match application any
   user@host# set policy vpn-untr-tr then permit
   ```

4. Reorder the security policies so that the vpn-tr-untr security policy is placed above the permit-any security policy.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# insert policy vpn-tr-untr before policy permit-any
   ```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy vpn-tr-untr {
    match {
      source-address sunnyvale;
      destination-address chicago;
      application any;
    }
    then {
      permit {
        tunnel {
          ipsec-vpn ike-vpn-chicago;
          pair-policy vpn-untr-tr;
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
    policy permit-any {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit
      }
    }
  }
  from-zone untrust to-zone trust {
    policy vpn-untr-tr {
      match {
        source-address chicago;
        destination-address sunnyvale;
        application any;
      }
      then {
        permit {
          tunnel {
            ipsec-vpn ike-vpn-chicago;
            pair-policy vpn-tr-untr;
          }
        }
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring TCP-MSS

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security flow tcp-mss ipsec-vpn mss 1350
```

Step-by-Step Procedure

To configure TCP-MSS information:

1. Configure TCP-MSS information.

```
[edit]
user@host# set security flow tcp-mss ipsec-vpn mss 1350
```

Results

From configuration mode, confirm your configuration by entering the **show security flow** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security flow
tcp-mss {
  ipsec-vpn {
    mss 1350;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the SSG Series Device

CLI Quick Configuration

For reference, the configuration for the SSG Series device is provided. For information about configuring SSG Series devices, see the *Concepts and Examples ScreenOS Reference Guide*, which is located at http://www.juniper.net/techpubs.

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interface ethernet0/6 zone Trust
set interface ethernet0/0 zone Untrust
set interface ethernet0/6 ip 192.168.168.1/24
set interface ethernet0/6 route
set interface ethernet0/0 ip 10.2.2.2/30
set interface ethernet0/0 route
set flow tcp-mss 1350
set address Trust “local-net” 192.168.168.0 255.255.255.0
set address Untrust "corp-net" 192.168.10.0 255.255.255.0
set ike gateway corp-ike address 10.1.1.2 Main outgoing-interface ethernet0/0 preshare
  395psksecr3t sec-level standard
set vpn corp-vpn gateway corp-ike replay tunnel idletime 0 sec-level standard
set policy id 11 from Trust to Untrust “local-net” “corp-net” “ANY” tunnel vpn “corp-vpn”
  pair-policy 10
set policy id 10 from Untrust to Trust “corp-net” “local-net” “ANY” tunnel vpn “corp-vpn”
  pair-policy 11
set policy id 1 from Trust to Untrust “ANY” “ANY” “ANY” nat src permit
set route 0.0.0.0/0 interface ethernet0/0 gateway 10.2.2.1
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the IKE Phase 1 Status on page 6235
- Verifying the IPsec Phase 2 Status on page 6237
- Reviewing Statistics and Errors for an IPsec Security Association on page 6238

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Copyright © 2016, Juniper Networks, Inc.

Action

NOTE: Before starting the verification process, you need to send traffic from a host in the 192.168.10/24 network to a host in the 192.168.168/24 network. For policy-based VPNs, a separate host must generate the traffic; traffic initiated from the SRX Series device will not match the VPN policy. We recommend that the test traffic be from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate ping from 192.168.10.10 to 192.168.168.10.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode
4       10.2.2.2         UP     5e1db3f9d50b0de6  e50865d9ebf134f8  Main

user@host> show security ike security-associations index 4 detail
IKE peer 10.2.2.2, Index 4,
  Role: Responder, State: UP
  Initiator cookie: 5e1db3f9d50b0de6, Responder cookie: e50865d9ebf134f8
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.1.1.2:500, Remote: 10.2.2.2:500
  Lifetime: Expires in 28770 seconds
  Algorithms:
   Authentication        : sha1
   Encryption            : aes-128-cbc
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input bytes    :                852
   Output bytes   :                856
   Input packets  :                  5
   Output packets :                  4
  Flags: Caller notification sent
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0
```

Meaning

The **show security ike security-associations** command lists all active IKE Phase 1 security associations (SAs). If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote Address—Verify that the remote IP address is correct.
- State
  - UP—The Phase 1 SA has been established.
  - DOWN—There was a problem establishing the Phase 1 SA.
- Mode—Verify that the correct mode is being used.

Copyright © 2016, Juniper Networks, Inc.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations index 1 detail** command lists additional information about the security association with an index number of 1:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Initiator and responder role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying the IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
  total configured sa: 2
  ID    Gateway         Port  Algorithm          SPI      Life:sec/kb  Mon vsys
  <2    10.2.2.2         500   ESP:aes-128/sha1   a63eb26f 3565/ unlim   -   0
  >2    10.2.2.2         500   ESP:aes-128/sha1   a1024ed9 3565/ unlim   -   0

user@host> show security ipsec security-associations index 2 detail
  Virtual-system: Root
  Local Gateway: 10.1.1.2, Remote Gateway: 10.2.2.2
  Local Identity: ipv4_subnet(any:0,[0..7]=192.168.10.0/24)
  Remote Identity: ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
    DF-bit: clear
    Policy-name: vpnpolicy-unt-tr

    Direction: inbound, SPI: 2789126767, AUX-SPI: 0
    Hard lifetime: Expires in 3558 seconds
    Lifesize Remaining: Unlimited
    Soft lifetime: Expires in 2986 seconds
    Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
```

Copyright © 2016, Juniper Networks, Inc.

```
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)

Anti-replay service: enabled, Replay window size: 32

Direction: outbound, SPI: 2701283033,, AUX-SPI: 0
Hard lifetime: Expires in 3558 seconds
Lifesize Remaining: Unlimited
Soft lifetime: Expires in 2986 seconds
Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc
Anti-replay service: enabled, Replay window size: 32
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The ID number is 2. Use this value with the **show security ipsec security-associations index** command to get more information about this particular SA.
- There is one IPsec SA pair using port 500, which indicates that no NAT-traversal is implemented. (NAT-traversal uses port 4500 or another random high-number port.)
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3565/ unlim value indicates that the Phase 2 lifetime expires in 3565 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U (up) or D (down) is listed.
- The virtual system (vsys) is the root system, and it always lists 0.

The output from the **show security ipsec security-associations index 16384 detail** command lists the following information:

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common reasons for a Phase 2 failure. For policy-based VPNs, the proxy ID is derived from the security policy. The local address and remote address are derived from the address book entries, and the service is derived from the application configured for the policy. If Phase 2 fails because of a proxy ID mismatch, you can use the policy to confirm which address book entries are configured. Verify that the addresses match the information being sent. Check the service to ensure that the ports match the information being sent.

Reviewing Statistics and Errors for an IPsec Security Association

Purpose Review ESP and authentication header counters and errors for an IPsec security association.

Action From operational mode, enter the **show security ipsec statistics index** *index_number* command, using the index number of the VPN for which you want to see statistics.

```
user@host> show security ipsec statistics index 2
```


Copyright © 2016, Juniper Networks, Inc.

```
ESP Statistics:
  Encrypted bytes:              920
  Decrypted bytes:             6208
  Encrypted packets:              5
  Decrypted packets:             87
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

You can also use the **show security ipsec statistics** command to review statistics and errors for all SAs.

To clear all IPsec statistics, use the **clear security ipsec statistics** command.

Meaning If you see packet loss issues across a VPN, you can run the **show security ipsec statistics** or **show security ipsec statistics detail** command several times to confirm that the encrypted and decrypted packet counters are incrementing. You should also check if the other error counters are incrementing.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Route-Based VPN on page 6074
- Example: Configuring a Hub-and-Spoke VPN on page 6094

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 82

Configuring VPNs with NAT-T

- Configuring Route-Based and Policy-Based VPNs with NAT-T on page 6243

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 261

Configuring Route-Based and Policy-Based VPNs with NAT-T

- Understanding NAT-T on page 6243
- Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device on page 6244
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272
- Example: Configuring NAT-T with Dynamic Endpoint VPN on page 6298

Understanding NAT-T

Network Address Translation-Traversal (NAT-T) is a method for getting around IP address translation issues encountered when data protected by IPsec passes through a NAT device for address translation. Any changes to the IP addressing, which is the function of NAT, causes IKE to discard packets. After detecting one or more NAT devices along the data path during Phase 1 exchanges, NAT-T adds a layer of User Datagram Protocol (UDP) encapsulation to IPsec packets so they are not discarded after address translation. NAT-T encapsulates both IKE and ESP traffic within UDP with port 4500 used as both the source and destination port. Because NAT devices age out stale UDP translations, keepalive messages are required between the peers.

There are two broad categories of NAT:

- Static NAT, where there is a one-to-one relationship between the private and public addresses. Static NAT works in both inbound and outbound directions.
- Dynamic NAT, where there is a many-to-one or many-to-many relationship between the private and public addresses. Dynamic NAT works in the outbound direction only.

The location of a NAT device can be such that:

- Only the IKEv1 or IKEv2 initiator is behind a NAT device. Multiple initiators can be behind separate NAT devices. Initiators can also connect to the responder through multiple NAT devices.
- Only the IKEv1 or IKEv2 responder is behind a NAT device.
- Both the IKEv1 or IKEv2 initiator and the responder are behind a NAT device.

Copyright © 2016, Juniper Networks, Inc.

Dynamic endpoint VPN covers the situation where the initiator's IKE external address is not fixed and is therefore not known by the responder. This can occur when the initiator's address is dynamically assigned by an ISP or when the initiator's connection crosses a dynamic NAT device that allocates addresses from a dynamic address pool.

Configuration examples for NAT-T are provided for the topology in which only the responder is behind a NAT device and the topology in which both the initiator and responder are behind a NAT device. Site-to-site IKE gateway configuration for NAT-T is supported on both the initiator and responder. A remote IKE ID is used to validate a peer's local IKE ID during Phase 1 of IKE tunnel negotiation. Both the initiator and responder require a **local-identity** and a **remote-identity** setting.

On all high-end SRX Series devices, the IPsec NAT-T tunnel scaling and sustaining issues are as follows:

- For a given private IP address, the NAT device should translate both 500 and 4500 private ports to the same public IP address.
- The total number of tunnels from a given public translated IP cannot exceed 1000 tunnels.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device on page 6244
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272
- Example: Configuring NAT-T with Dynamic Endpoint VPN on page 6298

Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device

This example shows how to configure a route-based VPN with a responder behind a NAT device to allow data to be securely transferred between a branch office and the corporate office.

- Requirements on page 6244
- Overview on page 6244
- Configuration on page 6250
- Verification on page 6265

Requirements

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

In this example, you configure a route-based VPN for a branch office in Chicago, Illinois, because you want to conserve tunnel resources but still get granular restrictions on VPN

Copyright © 2016, Juniper Networks, Inc.

traffic. Users in the Chicago office will use the VPN to connect to their corporate headquarters in Sunnyvale, California.

Figure 263 shows an example of a topology for route-based VPN with only the responder behind a NAT device.

Copyright © 2016, Juniper Networks, Inc.

Figure 263: Route-Based VPN Topology with Only the Responder Behind a NAT Device

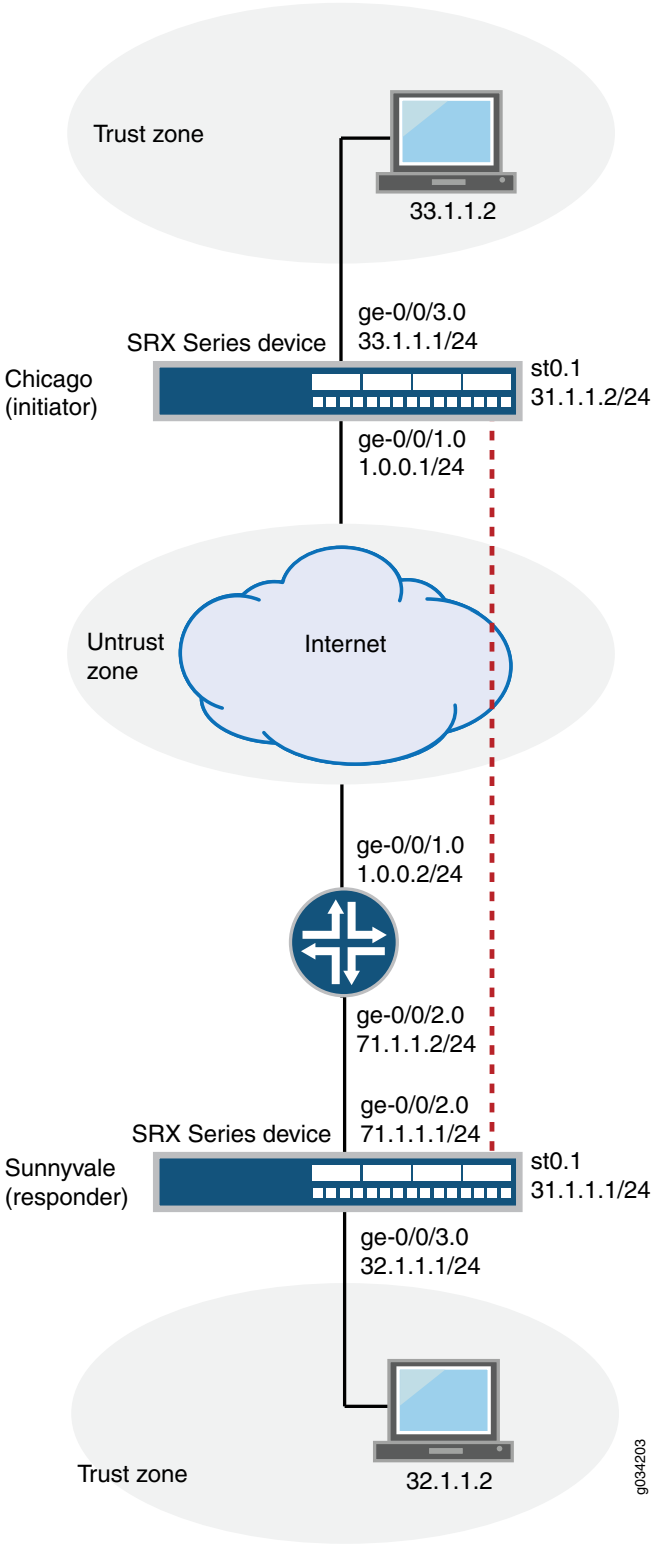

Copyright © 2016, Juniper Networks, Inc.

In this example, you configure interfaces, routing options, security zones, and security policies for both an initiator in Chicago and a responder in Sunnyvale. Then you configure IKE Phase 1 and IPsec Phase 2 parameters.

Packets sent from the initiator with a destination address 1.1.1.1/32 are translated to the destination address 71.1.1.1/32 on the NAT device.

See Table 531 through Table 533 for specific configuration parameters used for the initiator in the examples.

Table 531: Interface, Routing Options, Zones, and Security Policies for the Initiator

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/1	1.0.0.1/24
	ge-0/0/3	33.1.1.1/24
	st0.1 (tunnel interface)	31.1.1.2/24
Static routes	32.1.1.0/24	The next hop is st0.1.
	1.1.1.1/32	The next hop is 1.0.0.2.
Security zones	untrust	• Only IKE system service is allowed. • The ge-0/0/1.0 and the st0.1 interfaces are bound to this zone.
	trust	• All system services are allowed. • All protocols are allowed. • The ge-0/0/3.0 interface is bound to this zone.
Security policies	to-sunnyvale	Permit traffic from 33.1.1.1/24 in the trust zone to 32.1.1.1/24 in the untrust zone.
	from-sunnyvale	Permit traffic from 32.1.1.1/24 in the untrust zone to 33.1.1.1/24 in the trust zone.

Table 532: IKE Phase 1 Configuration Parameters for the Initiator

Feature	Name	Configuration Parameters
Proposal	ike_prop	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: 3des-cbc
Policy	ike_pol	• Mode: main • Proposal reference: ike_prop • IKE Phase 1 policy authentication method: pre-shared-key ascii-text

Copyright © 2016, Juniper Networks, Inc.

Table 532: IKE Phase 1 Configuration Parameters for the Initiator *(continued)*

Feature	Name	Configuration Parameters
Gateway	gw1	• IKE policy reference: ike_pol • External interface: ge-0/0/1.0 • Gateway address: 1.1.1.1 • Local peer (initiator): branch_natt1@example.net • Remote peer (responder): responder_natt1@example.net

Table 533: IPsec Phase 2 Configuration Parameters for the Initiator

Feature	Name	Configuration Parameters
Proposal	ipsec_prop	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: 3des-cbc
Policy	ipsec_pol	• Proposal reference: ipsec_prop • Perfect forward secrecy (PFS) keys: group2
VPN	vpn1	• IKE gateway reference: gw1 • IPsec policy reference: ipsec_pol • Bind to interface: st0.1 • Establish tunnels immediately

See Table 534 through Table 536 for specific configuration parameters used for the responder in the examples.

Table 534: Interface, Routing Options, Zones, and Security Policies for the Responder

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/2	71.1.1.1/24
	ge-0/0/3	32.1.1.1/24
	st0.1 (tunnel interface)	31.1.1.1/24
Static routes	0.0.0.0/0 (default route)	The next hop is 71.1.1.2.
	33.1.1.0/24	The next hop is st0.1.
Security zones	untrust	• Only IKE system service is allowed. • The ge-0/0/2.0 and the st0.1 interfaces are bound to this zone.
	trust	• All system services are allowed. All protocols are allowed. • The ge-0/0/3.0 interface is bound to this zone.

Copyright © 2016, Juniper Networks, Inc.

Table 534: Interface, Routing Options, Zones, and Security Policies for the Responder *(continued)*

Feature	Name	Configuration Parameters
Security policies	to-chicago	Permit traffic from 32.1.1.1/24 in the trust zone to 33.1.1.1/24 in the untrust zone.
	from-chicago	Permit traffic from 33.1.1.1/24 in the untrust zone to 32.1.1.1/24 in the trust zone.

Table 535: IKE Phase 1 Configuration Parameters for the Responder

Feature	Name	Configuration Parameters
Proposal	ike_prop	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: sha1 • Encryption algorithm: 3des-cbc
Policy	ike_pol	• Mode: main • Proposal reference: ike_prop • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gw1	• IKE policy reference: ike_pol • External interface: ge-0/0/2.0 • Gateway address: 1.0.0.1 • Local peer (responder): responder_natt1@example.net • Remote peer (initiator): branch_natt1@example.net

Table 536: IPsec Phase 2 Configuration Parameters for the Responder

Feature	Name	Configuration Parameters
Proposal	ipsec_prop	• Protocol: esp • Authentication algorithm: hmac-sha1-96 • Encryption algorithm: 3des-cbc
Policy	ipsec_pol	• Proposal reference: ipsec_prop • PFS keys: group2
VPN	vpn1	• IKE gateway reference: gw1 • IPsec policy reference: ipsec_pol • Bind to interface: st0.1 • Establish tunnels immediately

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Configuring Interface, Routing Options, Security Zones, and Security Policies for the Initiator on page 6250
- Configuring IKE for the Initiator on page 6254
- Configuring IPsec for the Initiator on page 6256
- Configuring Interfaces, Routing Options, Security Zones, and Security Policies for the Responder on page 6258
- Configuring IKE for the Responder on page 6261
- Configuring IPsec for the Responder on page 6264

Configuring Interface, Routing Options, Security Zones, and Security Policies for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.0.0.1/24
set interfaces ge-0/0/3 unit 0 family inet address 33.1.1.1/24
set interfaces st0 unit 1 family inet address 31.1.1.2/24
set routing-options static route 32.1.1.0/24 next-hop st0.1
set routing-options static route 1.1.1.1/32 next-hop 1.0.0.2
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone untrust interfaces st0.1
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security address-book book1 address Chicago-lan 33.1.1.1/24
set security address-book book1 attach zone trust
set security address-book book2 address Sunnyvale-lan 32.1.1.1/24
set security address-book book2 attach zone untrust
set security policies from-zone trust to-zone untrust policy to-sunnyvale match
    source-address Chicago-lan
set security policies from-zone trust to-zone untrust policy to-sunnyvale match
    destination-address Sunnyvale-lan
set security policies from-zone trust to-zone untrust policy to-sunnyvale match application
    any
set security policies from-zone trust to-zone untrust policy to-sunnyvale then permit
set security policies from-zone untrust to-zone trust policy from-sunnyvale match
    source-address Sunnyvale-lan
set security policies from-zone untrust to-zone trust policy from-sunnyvale match
    destination-address Chicago-lan
set security policies from-zone untrust to-zone trust policy from-sunnyvale match
    application any
set security policies from-zone untrust to-zone trust policy from-sunnyvale then permit
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interface, static route, security zone, and security policy information:

1. Configure Ethernet interface information.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family inet address 1.0.0.1/24
user@host# set interfaces ge-0/0/3 unit 0 family inet address 33.1.1.1/24
user@host# set interfaces st0 unit 1 family inet address 31.1.1.2/24
```

2. Configure static route information.

```
[edit]
user@host# set routing-options static route 32.1.1.0/24 next-hop st0.1
user@host# set routing-options static route 1.1.1.1/32 next-hop 1.0.0.2
```

3. Configure the untrust security zone.

```
[edit ]
user@host# set security zones security-zone untrust
```

4. Assign interfaces to the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces st0.1
```

5. Specify allowed system services for the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@host# set security zones security-zone trust host-inbound-traffic protocols
  all
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/3.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

9. Configure address books.

```
[edit security address-book]
user@host# set book1 address Chicago-lan 33.1.1.1/24
user@host# set book1 attach zone trust
user@host# set book2 address Sunnyvale-lan 32.1.1.1/24
user@host# set book2 attach zone untrust
```

10. Create security policies.

```
[edit security security-policies from-zone trust to-zone untrust]
user@host# set policy to-sunnyvale match source-address Chicago-lan
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy to-sunnyvale match destination-address Sunnyvale-lan
user@host# set policy to-sunnyvale match application any
user@host# set policy to-sunnyvale then permit

[edit security security-policies from-zone untrust to-zone trust]
user@host# set policy from-sunnyvale match source-address Sunnyvale-lan
user@host# set policy from-sunnyvale match destination-address Chicago-lan
user@host# set policy from-sunnyvale match application any
user@host# set policy from-sunnyvale then permit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, **show security address-book**, and **show security policies**commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 1.0.0.1/24;
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 33.1.1.1/24;
    }
  }
}
st0 {
  unit 1 {
    family inet {
      address 31.1.1.2/24
    }
  }
}
```

```
[edit]
user@host# show routing-options
static {
  route 32.1.1.0/24 next-hop st0.1;
  route 1.1.1.1/32 next-hop 1.0.0.2;
}
```

```
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      ike;
    }
  }
  interfaces {
```

 Copyright © 2016, Juniper Networks, Inc.

```
        st0.1;
        ge-0/0/1.0;
    }
}
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
        all;
        }
    }
    interfaces {
        ge-0/0/3.0;
        }
    [edit]
[edit]
user@host# show security address-book
    book1 {
        address Chicago-lan 33.1.1.1/24;
        attach {
            zone trust;
        }
    }
    book2 {
        address Sunnyvale-lan 32.1.1.1/24;
        attach {
            zone untrust;
        }
    }
[edit]
user@host# show security policies
    from-zone trust to-zone untrust {
        policy to-sunnyvale {
            match {
                source-address Chicago-lan;
                destination-address Sunnyvale-lan;
                application any;
            }
            then {
                permit;
            }
        }
    }
    from-zone untrust to-zone trust {
        policy from-sunnyvale {
            match {
                source-address Sunnyvale-lan;
                destination-address Chicago-lan;
                application any;
            }
            then {
                permit;
            }
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike_prop authentication-method pre-shared-keys
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_pol mode main
set security ike policy ike_pol proposals ike_prop
set security ike policy ike_pol pre-shared-key ascii-text "$ABC123"
set security ike gateway gw1 ike-policy ike_pol
set security ike gateway gw1 address 1.1.1.1
set security ike gateway gw1 local-identity user-at-hostname branch_natt1@example.net
set security ike gateway gw1 remote-identity user-at-hostname
  responder_natt1@example.net
set security ike gateway gw1 external-interface ge-0/0/1.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike_prop
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike_prop]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy ike_pol
```

7. Set the IKE Phase 1 policy mode.

```
[edit security ike policy ike_pol]
user@host# set mode main
```

8. Specify a reference to the IKE proposal.

```
[edit security ike policy ike_pol]
user@host# set proposals ike_prop
```

9. Define the IKE Phase 1 policy authentication method.

```
[edit security ike policy ike_pol]
user@host# set pre-shared-key ascii-text "$ABC123"
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike gateway gw1]
user@host# set external-interface ge-0/0/1.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw1]
user@host# set ike-policy ike_pol
```

12. Define the IKE Phase 1 gateway address.

```
[edit security ike gateway gw1]
user@host# set address 1.1.1.1
```

13. Set **local-identity** of the local peer.

```
[edit security ike gateway gw1]
user@host# set local-identity user-at-hostname branch_natt1@example.net
```

14. Set **remote-identity** of the responder. This is the IKE identifier.

```
[edit security ike gateway gw1]
user@host# set remote-identity user-at-hostname responder_natt1@example.net
```

15. Define the external interface.

```
[edit security ike gateway gw1]
user@host# set external-interface ge-0/0/1.0
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ike
proposal ike_prop {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
}
policy ike_pol {
    mode main;
    proposals ike_prop;
```

Copyright © 2016, Juniper Networks, Inc.

```
    pre-shared-key ascii-text "$ABC123";
}
gateway gw1 {
    ike-policy ike_poly;
    address 1.1.1.1;
    local-identity user-at-hostname branch_natt1@example.net;
    remote-identity user-at-hostname responder_natt1@example.net;
    external-interface ge-0/0/1.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec policy ipsec_pol perfect-forward-secrecy keys group2
set security ipsec policy ipsec_pol proposals ipsec_prop
set security ipsec vpn vpn1 bind-interface st0.1
set security ipsec vpn vpn1 ike gateway gw1
set security ipsec vpn vpn1 ike ipsec-policy ipsec_pol
set security ipsec vpn vpn1 establish-tunnels immediately
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec_prop
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set authentication-algorithm hmac-sha1-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

5. Create the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy ipsec_pol
```

6. Specify IPsec Phase 2 to use perfect forward secrecy (PFS).

```
[edit security ipsec policy ipsec_pol]
user@host# set perfect-forward-secrecy keys group2
```

7. Specify the IPsec Phase 2 proposal reference.

```
[edit security ipsec policy ipsec_pol]
user@host# set proposals ipsec_prop
```

8. Specify the IKE gateway.

```
[edit security ipsec]
user@host# set vpn vpn1 ike gateway gw1
```

9. Specify the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set vpn vpn1 ike ipsec-policy ipsec_pol
```

10. Specify the interface to bind.

```
[edit security ipsec]
user@host# set vpn vpn1 bind-interface st0.1
```

11. Specify that the tunnel be brought up immediately without waiting for a verification packet to be sent.

```
[edit security ipsec]
user@host# set vpn vpn1 establish-tunnels immediately
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ipsec
proposal ipsec_prop {
   protocol esp;
   authentication-algorithm hmac-sha1-96;
   encryption-algorithm 3des-cbc;
}
policy ipsec_pol {
   perfect-forward-secrecy {
      keys group2;
   }
   proposals ipsec_prop;
}
vpn vpn1 {
   bind-interface st0.1;
      ike {
         gateway gw1;
         ipsec-policy ipsec_pol;
         }
      establish-tunnels immediately;
   }
   proposals ipsec_prop;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Interfaces, Routing Options, Security Zones, and Security Policies for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/2 unit 0 family inet address 71.1.1.1/24
set interfaces ge-0/0/3 unit 0 family inet address 32.1.1.1/24
set interfaces st0 unit 1 family inet address 31.1.1.1/24
set routing-options static route 0.0.0.0/0 next-hop 71.1.1.2
set routing-options static route 33.1.1.0/24 next-hop st0.1
set security zones security-zone untrust host-inbound-traffic system-services ike
set security zones security-zone untrust interfaces ge-0/0/2.0
set security zones security-zone untrust interfaces st0.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security address-book book1 address Sunnyvale-lan 32.1.1.1/24
set security address-book book1 attach zone trust
set security address-book book2 address Chicago-lan 33.1.1.1/24
set security address-book book2 attach zone untrust
set security policies from-zone trust to-zone untrust policy to-chicago match
  source-address Sunnyvale-lan
set security policies from-zone trust to-zone untrust policy to-chicago match
  destination-address Chicago-lan
set security policies from-zone trust to-zone untrust policy to-chicago match application
  any
set security policies from-zone trust to-zone untrust policy to-chicago then permit
set security policies from-zone untrust to-zone trust policy from-chicago match
  source-address Chicago-lan
set security policies from-zone untrust to-zone trust policy from-chicago match
  destination-address Sunnyvale-lan
set security policies from-zone untrust to-zone trust policy from-chicago match application
  any
set security policies from-zone untrust to-zone trust policy from-chicago then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interface, static route, security zones, policies and gateways:

1. Configure Ethernet interface information.

```
[edit]
user@host# set interfaces ge-0/0/2 unit 0 family inet address 71.1.1.1/24
user@host# set interfaces ge-0/0/3 unit 0 family inet address 32.1.1.1/24
user@host# set interfaces st0 unit 1 family inet address 31.1.1.1/24
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure static route information.

```
[edit]
user@host# set routing-options static route 0.0.0.0/0 next-hop 71.1.1.2
user@host# set routing-options static route 33.1.1.0/24 next-hop st0.1
```

3. Configure the untrust security zone.

```
[edit ]
user@host# set security zones security-zone untrust
```

4. Assign interfaces to the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set security zones security-zone untrust interfaces ge-0/0/2.0
user@host# set security zones security-zone untrust interfaces st0.1
```

5. Specify allowed system services for the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services ike
```

6. Configure the trust security zone.

```
[edit]
user@host# set security zones security-zone trust host-inbound-traffic protocols
  all
```

7. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/3.0
```

8. Specify allowed system services for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

9. Configure address books.

```
[edit security address-book]
user@host# set book1 address Sunnyvale-lan 32.1.1.1/24
user@host# set book1 attach zone trust
user@host# set book2 address Chicago-lan 33.1.1.1/24
user@host# set book2 attach zone untrust
```

10. Create security policies.

```
[edit security security-policies from-zone trust to-zone untrust]
user@host# set policy to-chicago match source-address Sunnyvale-lan
user@host# set policy to-chicago match destination-address Chicago-lan
user@host# set policy to-chicago match application any
user@host# set policy to-chicago then permit

[edit security security-policies from-zone untrust to-zone trust]
user@host# set policy from-chicago match source-address Chicago-lan
user@host# set policy from-chicago match destination-address Sunnyvale-lan
user@host# set policy from-chicago match application any
user@host# set policy from-chicago then permit
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, **show security address-book**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
  unit 0 {
    family inet {
      address 71.1.1.1/24;
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 32.1.1.1/24;
    }
  }
}
st0 {
  unit 1 {
    family inet {
      address 31.1.1.1/24
    }
  }
}
```

```
[edit]
user@host# show routing-options
static {
  route 0.0.0.0/0 next-hop 71.1.1.2;
  route 33.1.1.0/24 next-hop st0.1;
}
```

```
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      ike;
    }
  }
  interfaces {
    ge-0/0/2.0;
    st0.1;
  }
}
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
    all;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  interfaces {
    ge-0/0/3.0;
  }
}
[edit]
user@host# show security address-book
  book1 {
    address Sunnyvale-lan 32.1.1.1/24;
    attach {
      zone trust;
    }
  }
  book2 {
    address Chicago-lan 33.1.1.1/24;
    attach {
      zone untrust;
    }
  }
[edit]
user@host# show security policies
  from-zone trust to-zone untrust {
    policy to-chicago {
      match {
        source-address Sunnyvale-lan;
        destination-address Chicago-lan;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone untrust to-zone trust {
    policy from-chicago {
      match {
        source-address Chicago-lan;
        destination-address Sunnyvale-lan;
        application any;
      }
      then {
        permit;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set security ike proposal ike_prop authentication-method pre-shared-keys
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_pol mode main
set security ike policy ike_pol proposals ike_prop
set security ike policy ike_pol pre-shared-key ascii-text "$ABC123"
set security ike gateway gw1 ike-policy ike_pol
set security ike gateway gw1 address 1.0.0.1
set security ike gateway gw1 local-identity user-at-hostname
  responder_natt1@example.net
set security ike gateway gw1 remote-identity user-at-hostname branch_natt1@example.net
set security ike gateway gw1 external-interface ge-0/0/2.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike_prop
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike_prop]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-algorithm sha1
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike]
   user@host# set policy ike_pol
   ```

7. Set the IKE Phase 1 policy mode.

   ```
   [edit security ike policy ike_pol]
   user@host# set mode main
   ```

8. Specify a reference to the IKE proposal.

   ```
   [edit security ike policy ike_pol]
   user@host# set proposals ike_prop
   ```

9. Define the IKE Phase 1 policy authentication method.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike policy ike_pol]
user@host# set pre-shared-key ascii-text “$ABC123”
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike gateway gw1]
user@host# set external-interface ge-0/0/2.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway gw1]
user@host# set ike-policy ike_pol
```

12. Define the IKE Phase 1 gateway address.

```
[edit security ike gateway gw1]
user@host# set address 1.0.0.1
```

13. Set **local-identity** of the responder.

```
[edit security ike gateway gw1]
user@host# set local-identity user-at-hostname responder_natt1@example.net
```

14. Set **remote-identity** of the responder. This is the IKE identifier.

```
[edit security ike gateway gw1]
user@host# set remote-identity user-at-hostname branch_natt1@example.net
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ike
proposal ike_prop {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm sha1;
   encryption-algorithm 3des-cbc;
}
policy ike_pol {
   mode main;
   proposals ike_prop;
   pre-shared-key ascii-text “$ABC123”;
}
gateway gw1 {
   ike-policy ike_pol;
   address 1.0.0.1;
   local-identity user-at-hostname "responder_natt1@example.net";
   remote-identity user-at-hostname "branch_natt1@example.net";
   external-interface ge-0/0/2.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring IPsec for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec policy ipsec_pol perfect-forward-secrecy keys group2
set security ipsec policy ipsec_pol proposals ipsec_prop
set security ipsec vpn vpn1 bind-interface st0.1
set security ipsec vpn vpn1 ike gateway gw1
set security ipsec vpn vpn1 ike ipsec-policy ipsec_pol
set security ipsec vpn vpn1 establish-tunnels immediately
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec_prop
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set authentication-algorithm hmac-sha1-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec_prop ]
   user@host# set encryption-algorithm 3des-cbc
   ```

5. Create the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@host# set policy ipsec_pol
   ```

6. Specify IPsec Phase 2 to use perfect forward secrecy (PFS).

   ```
   [edit security ipsec policy ipsec_pol]
   user@host# set perfect-forward-secrecy keys group2
   ```

7. Specify the IPsec Phase 2 proposal reference.

   ```
   [edit security ipsec policy ipsec_pol]
   user@host# set proposals ipsec_prop
   ```

8. Specify the IKE gateway.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ipsec]
user@host# set security ipsec vpn vpn1 ike gateway gw1
```

9. Specify the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set vpn vpn1 ike ipsec-policy ipsec_pol
```

10. Specify the interface to bind.

```
[edit security ipsec]
user@host# set vpn vpn1 bind-interface st0.1
```

11. Specify that the tunnel be brought up immediately without waiting for a verification packet to be sent.

```
[edit security ipsec]
user@host# set vpn vpn1 establish-tunnels immediately
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ipsec
proposal ipsec_prop {
   protocol esp;
   authentication-algorithm hmac-sha1-96;
   encryption-algorithm 3des-cbc;
}
policy ipsec_pol {
   perfect-forward-secrecy {
      keys group2;
   }
   proposals ipsec_prop;
}
vpn vpn1 {
   bind-interface st0.1;
   ike {
      gateway gw1;
      ipsec-policy ipsec_pol;
      }
   establish-tunnels immediately;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the IKE Phase 1 Status for the Initiator on page 6266
- Verifying IPsec Security Associations for the Initiator on page 6267
- Verifying the IKE Phase 1 Status for the Responder on page 6268
- Verifying IPsec Security Associations for the Responder on page 6270

Copyright © 2016, Juniper Networks, Inc.

Verifying the IKE Phase 1 Status for the Initiator

Purpose Verify the IKE Phase 1 status.

Action

NOTE: Before starting the verification process, you must send traffic from a host in the 33.1.1.0 network to a host in the 32.1.1.0 network. For route-based VPNs, traffic can be initiated by the SRX Series device through the tunnel. We recommend that when testing IPsec tunnels, test traffic be sent from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate a ping operation from 33.1.1.2 to 32.1.1.2.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

106321  UP     d31d6833108fd69f  9ddfe2ce133086aa  Main           1.1.1.1

user@host> show security ike security-associations index 1 detail
IKE peer 1.1.1.1, Index
  Initiator cookie: d31d6833108fd69f, Responder cookie: 9ddfe2ce133086aa
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 1.0.0.1:500, Remote: 1.1.1.1:500
  Lifetime: Expires in 28785 seconds
  Peer ike-id: responder_natt1@example.net
  Xauth assigned IP: responder_natt1@example.net
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Flags: IKE SA is created
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 1.0.0.1:500, Remote: 1.1.1.1:500
    Local identity: branch_natt1@example.net
    Remote identity: responder_natt1@example.net
    Flags: IKE SA is created
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

Copyright © 2016, Juniper Networks, Inc.

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote address—Verify that the remote IP address is correct and that port 500 is being used for peer-to-peer communication.
- Role initiator state
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.
  - Both peers in the IPsec SA pair are using port 500.
  - Peer IKE ID—Verify the remote address is correct.
  - Local identity and remote identity—Verify these are correct.
- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations** command lists additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying IPsec Security Associations for the Initiator

Purpose Verify the IPsec status.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <131073 ESP:3des/sha1 ac23df79 2532/ unlim   -   root 500  1.1.1.1
  >131073 ESP:3des/sha1 cbc9281a 2532/ unlim   -   root 500  1.1.1.1

user@host> show security ipsec security-associations detail
Virtual-system: root
  Local Gateway: 1.0.0.1, Remote Gateway: 1.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
    DF-bit: clear
    Direction: inbound, SPI: ac23df79, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3186 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2578 seconds
    Mode: Tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: cbc9281a, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3186 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2578 seconds
    Mode: Tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has a NAT address of 1.1.1.1.
- Both peers in the IPsec SA pair are using port 500.
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 2532/ unlim value indicates that the Phase 2 lifetime expires in 2532 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

Verifying the IKE Phase 1 Status for the Responder

Purpose Verify the IKE Phase 1 status.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5802591 UP     d31d6833108fd69f  9ddfe2ce133086aa  Main           1.0.0.1

user@host> show security ike security-associations index 1 detail
IKE peer 1.0.0.1, Index 5802591,
  Role: Responder, State: UP
  Initiator cookie: d31d6833108fd69f, Responder cookie: 9ddfe2ce133086aa
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 71.1.1.1:500, Remote: 1.0.0.1:500
  Lifetime: Expires in 25704 seconds
  Peer ike-id: branch_natt1@example.net
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Flags: IKE SA is created
  IPSec security associations: 8 created, 2 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 71.1.1.1:500, Remote: 1.0.0.1:500
    Local identity: responder_natt1@example.net
    Remote identity: branch_natt1@example.net
    Flags: IKE SA is created
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote address—Verify that the remote IP address is correct and that port 500 is being used for peer-to-peer communication.
- Role responder state
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.
  - Peer IKE ID—Verify the address is correct.
  - Local identity and remote identity—Verify these addresses are correct.
- Mode—Verify that the correct mode is being used.

Copyright © 2016, Juniper Networks, Inc.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations** command lists additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying IPsec Security Associations for the Responder

Purpose Verify the IPsec status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <131073 ESP:3des/sha1 a5224cd9 3571/ unlim   -   root 500  1.0.0.1
  >131073 ESP:3des/sha1 82a86a07 3571/ unlim   -   root 500  1.0.0.1

user@host> show security ipsec security-associations detail
Virtual-system: root
  Local Gateway: 71.1.1.1, Remote Gateway: 1.0.0.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
    DF-bit: clear
    Direction: inbound, SPI: a5224cd9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3523 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2923 seconds
```


Copyright © 2016, Juniper Networks, Inc.

```
Mode: Tunnel, Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: counter-based enabled, Replay window size: 64

Direction: outbound, SPI: 82a86a07, AUX-SPI: 0
                          , VPN Monitoring: -
Hard lifetime: Expires in 3523 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 2923 seconds
Mode: Tunnel, Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has an ip address of 1.0.0.1.
- Both peers in the IPsec SA pair are using port 500.
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3571/ unlim value indicates that the Phase 2 lifetime expires in 3571 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

The output from the **show security ipsec security-associations index *index_id*detail** command lists the following information:

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, are correct for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be specified. For some third-party vendors, the proxy ID must be manually entered to match.
- Another common reason for Phase 2 failure is not specifying the ST interface binding. If IPsec cannot complete, check the kmd log or set traceoptions.

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding NAT-T on page 6243
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device

This example shows how to configure a policy-based VPN with both an initiator and a responder behind a NAT device to allow data to be securely transferred between a branch office and the corporate office.

- Requirements on page 6272
- Overview on page 6272
- Configuration on page 6277
- Verification on page 6292

Requirements

Before you begin, read "IPsec VPN Overview" on page 6041.

Overview

In this example, you configure a policy-based VPN for a branch office in Chicago, Illinois, because you want to conserve tunnel resources but still get granular restrictions on VPN traffic. Users in the branch office will use the VPN to connect to their corporate headquarters in Sunnyvale, California.

In this example, you configure interfaces, routing options, security zones, security policies for both an initiator and a responder.

Figure 264 shows an example of a topology for a VPN with both an initiator and a responder behind a NAT device.

Copyright © 2016, Juniper Networks, Inc.

Figure 264: Policy-Based VPN Topology with Both an Initiator and a Responder Behind a NAT Device

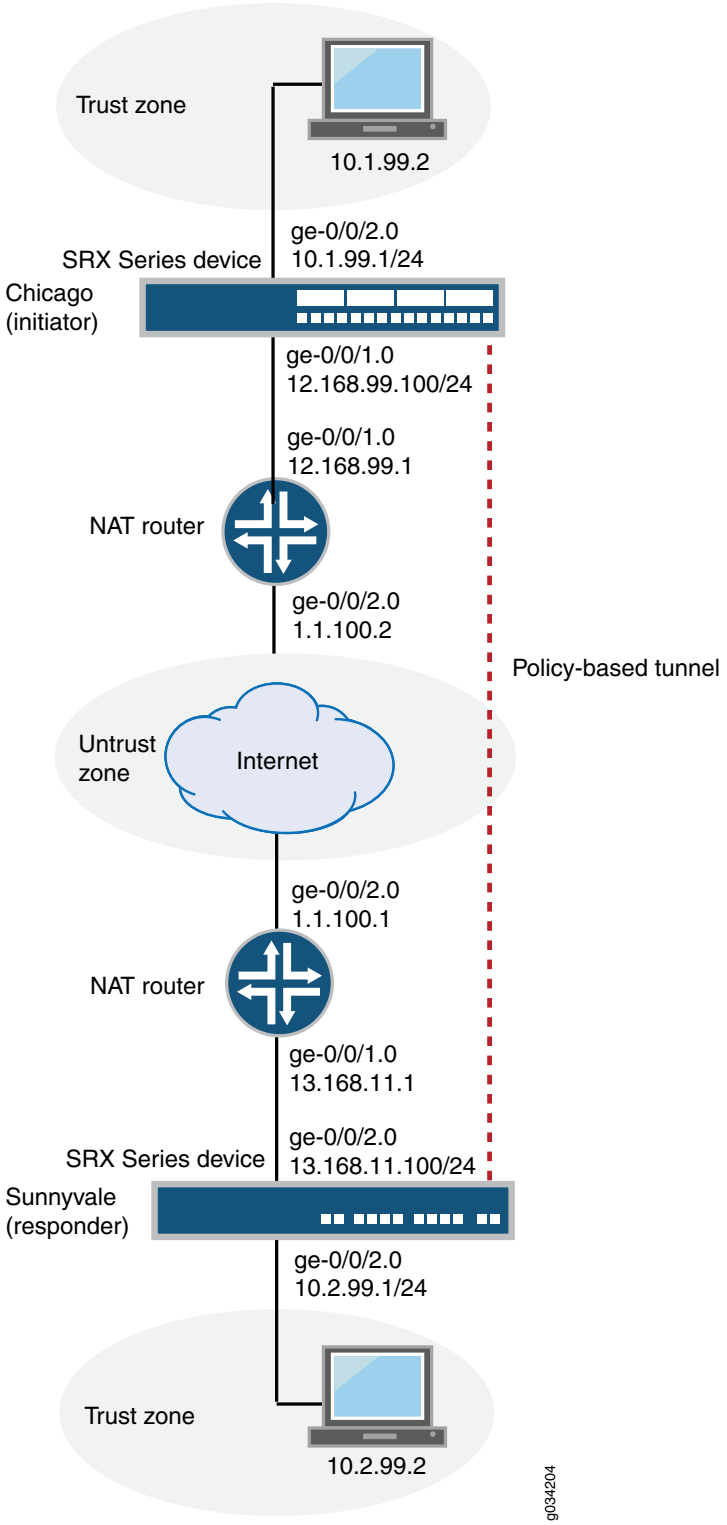

In this example, you configure interfaces, an IPv4 default route, and security zones. Then you configure IKE Phase 1, including local and remote peers, IPsec Phase 2, and the security policy. Note in the example above, the responder's private IP address 13.168.11.1 is hidden by the NAT device and mapped to public IP address 1.1.100.1.

See Table 537 through Table 540 for specific configuration parameters used for the initiator in the examples.

Table 537: Interface, Routing Options, and Security Zones for the Initiator

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/1	12.168.99.100/24
	ge-0/0/2	10.1.99.1/24
Static routes	10.2.99.0/24 (default route)	The next hop is 12.168.99.1.
	13.168.11.0/24	The next hop is 12.168.99.1.
	1.1.100.0/24	12.168.99.1
Security zones	trust	• All system services are allowed. • All protocols are allowed. • The ge-0/0/2.0 interface is bound to this zone.
	untrust	• All system services are allowed. • All protocols are allowed. • The ge-0/0/1.0 interface is bound to this zone.

Table 538: IKE Phase 1 Configuration Parameters for the Initiator

Feature	Name	Configuration Parameters
Proposal	ike_prop	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: md5 • Encryption algorithm: 3des-cbc
Policy	ike_pol	• Mode: main • Proposal reference: ike_prop • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gate	• IKE policy reference: ike_pol • External interface: ge-0/0/1.0 • Gateway address: 1.1.100.23 • Local peer is inet 11.11.11.11 • Remote peer is inet 44.44.44.44

Copyright © 2016, Juniper Networks, Inc.

Table 539: IPsec Phase 2 Configuration Parameters for the Initiator

Feature	Name	Configuration Parameters
Proposal	ipsec_prop	• Protocol: esp • Authentication algorithm: hmac-md5-96 • Encryption algorithm: 3des-cbc
Policy	ipsec_pol	• Proposal reference: ipsec_prop • Perfect forward secrecy (PFS): group1
VPN	first_vpn	• IKE gateway reference: gate • IPsec policy reference: ipsec_pol

Table 540: Security Policy Configuration Parameters for the Initiator

Purpose	Name	Configuration Parameters

The security policy permits tunnel traffic from the trust zone to the untrust zone.	pol1	• Match criteria:
		  ◦ source-address any
		  ◦ destination-address any
		  ◦ application any
		• Action: permit tunnel ipsec-vpn first_vpn

The security policy permits tunnel traffic from the untrust zone to the trust zone.	pol1	• Match criteria:
		  ◦ source-address any
		  ◦ destination-address any
		  ◦ application any
		• Action: permit tunnel ipsec-vpn first_vpn

See Table 541 through Table 544 for specific configuration parameters used for the responder in the examples.

Table 541: Interface, Routing Options, and Security Zones for the Responder

Feature	Name	Configuration Parameters
Interfaces	ge-0/0/2	13.168.11.100/24
	ge-0/0/3	10.2.99.1/24
Static routes	10.1.99.0/24 (default route)	The next hop is 13.168.11.1.
	12.168.99.0/24	The next hop is 13.168.11.1.
	1.1.100.0/24	13.168.11.1

Copyright © 2016, Juniper Networks, Inc.

Table 541: Interface, Routing Options, and Security Zones for the Responder *(continued)*

Feature	Name	Configuration Parameters
Security zones	trust	• All system services are allowed. • All protocols are allowed. • The ge-0/0/3.0 interface is bound to this zone.
	untrust	• All system services are allowed. • All protocols are allowed. • The ge-0/0/2.0 interface is bound to this zone.

Table 542: IKE Phase 1 Configuration Parameters for the Responder

Feature	Name	Configuration Parameters
Proposal	ike_prop	• Authentication method: pre-shared-keys • Diffie-Hellman group: group2 • Authentication algorithm: md5 • Encryption algorithm: 3des-cbc
Policy	ike_pol	• Mode: main • Proposal reference: ike_prop • IKE Phase 1 policy authentication method: pre-shared-key ascii-text
Gateway	gate	• IKE policy reference: ike_pol • External interface: ge-0/0/2.0 • Gateway address: 1.1.100.22 • Always send dead-peer detection • Local peer is inet 44.44.44.44 • Remote peer is inet 11.11.11.11

Table 543: IPsec Phase 2 Configuration Parameters for the Responder

Feature	Name	Configuration Parameters
Proposal	ipsec_prop	• Protocol: esp • Authentication algorithm: hmac-md5-96 • Encryption algorithm: 3des-cbc
Policy	ipsec_pol	• Proposal reference: ipsec_prop • Perfect forward secrecy (PFS): group1
VPN	first_vpn	• IKE gateway reference: gate • IPsec policy reference: ipsec_pol • Establish tunnels immediately

Copyright © 2016, Juniper Networks, Inc.

Table 544: Security Policy Configuration Parameters for the Responder

Purpose	Name	Configuration Parameters

The security policy permits tunnel traffic from the trust zone to the untrust zone.	pol1	• Match criteria:
		  • source-address any
		  • destination-address any
		  • application any
		• Action: permit tunnel ipsec-vpn first_vpn

The security policy permits tunnel traffic from the untrust zone to the trust zone.	pol1	• Match criteria:
		  • source-address any
		  • destination-address any
		  • application any
		• Action: permit tunnel ipsec-vpn first_vpn

Configuration

- Configuring Interface, Routing Options, and Security Zones for the Initiator on page 6277
- Configuring IKE for the Initiator on page 6279
- Configuring IPsec for the Initiator on page 6281
- Configuring Security Policies for the Initiator on page 6283
- Configuring Interface, Routing Options, and Security Zones for the Responder on page 6284
- Configuring IKE for the Responder on page 6287
- Configuring IPsec for the Responder on page 6289
- Configuring Security Policies for the Responder on page 6291

Configuring Interface, Routing Options, and Security Zones for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces ge-0/0/1 unit 0 family inet address 12.168.99.100/24
set interfaces ge-0/0/2 unit 0 family inet address 10.1.99.1/24
set routing-options static route 10.2.99.0/24 next-hop 12.168.99.1
set routing-options static route 13.168.11.0/24 next-hop 12.168.99.1
set routing-options static route 1.1.100.0/24 next-hop 12.168.99.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/2.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interfaces, static routes, and security zones:

1. Configure Ethernet interface information.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family inet address 12.168.99.100/24
user@host# set interfaces ge-0/0/2 unit 0 family inet address 10.1.99.1/24
```

2. Configure static route information.

```
[edit]
user@host# set routing-options static route 10.2.99.0/24 next-hop 12.168.99.1
user@host# set routing-options static route 13.168.11.0/24 next-hop 12.168.99.1
```

3. Configure the trust security zone.

```
[edit ]
user@host# set security zones security-zone trust host-inbound-traffic protocols
  all
```

4. Assign an interface to the trust security zone.

```
[edit security zones security-zone trust]
user@host# set interfaces ge-0/0/2.0
```

5. Specify system services for the trust security zone.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
```

6. Configure the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic protocols all
```

7. Assign an interface to the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set interfaces ge-0/0/1.0
```

8. Specify system services for the untrust security zone.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security zones** commands If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 12.168.99.100/24;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.99.1/24;
    }
  }
}
```

```
[edit]
user@host# show routing-options
static {
  route 10.2.99.0/24 next-hop 12.168.99.1;
  route 13.168.11.0/24 next-hop 12.168.99.1;
  route 1.1.100.0/24 next-hop 12.168.99.1;
}
```

```
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols{
    all;
    }
  }
  interfaces {
    ge-0/0/1.0.;
  }
}
security-zone trust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
    all;
    }
  }
  interfaces {
    ge-0/0/2.0;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike_prop authentication-method pre-shared-keys
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm md5
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_pol mode main
set security ike policy ike_pol proposals ike_prop
set security ike policy ike_pol pre-shared-key ascii-text "$ABC123"
set security ike gateway gate ike-policy ike_pol
set security ike gateway gate address 1.1.100.23
set security ike gateway gate external-interface ge-0/0/1.0
set security ike gateway gate local-identity inet 11.11.11.11
set security ike gateway gate remote-identity inet 44.44.44.44
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike_prop
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-method pre-shared-keys
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike_prop]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-algorithm md5
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

6. Create an IKE Phase 1 policy.

   ```
   [edit security ike policy ]
   user@host# set policy ike_pol
   ```

7. Set the IKE Phase 1 policy mode.

   ```
   [edit security ike policy ike_pol]
   user@host# set mode main
   ```

8. Specify a reference to the IKE proposal.

   ```
   [edit security ike policy ike_pol]
   user@host# set proposals ike_prop
   ```

9. Define the IKE Phase 1 policy authentication method.

   ```
   [edit security ike policy ike_pol pre-shared-key]
   ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ascii-text "$ABC123”
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike ]
user@host# set gateway gate external-interface ge-0/0/1.0
```

11. Create an IKE Phase 1 gateway address.

```
[edit security ike gateway]
set gate address 1.1.100.23
```

12. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway]
set gate ike-policy ike_pol
```

13. Set **local-identity** for the local peer.

```
[edit security ike gateway gate]
user@host# set local-identity inet 11.11.11.11
```

14. Set **remote-identity** for the responder. This is the responder's local identity.

```
[edit security ike gateway gate ]
user@host# set remote-identity inet 44.44.44.44
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ike
proposal ike_prop {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm md5;
   encryption-algorithm 3des-cbc;
}
policy ike_pol {
   mode main;
   proposals ike_prop;
   pre-shared-key ascii-text “$ABC123”;
}
gateway gate {
   ike-policy ike_pol;
   address 1.1.100.23;
   local-identity 11.11.11.11;
   remote-identity 44.44.44.44;
   external-interface ge-0/0/1.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec for the Initiator

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

Copyright © 2016, Juniper Networks, Inc.

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec policy ipsec_pol perfect-forward-secrecy keys group1
set security ipsec policy ipsec_pol proposals ipsec_prop
set security ipsec vpn first_vpn ike gateway gate
set security ipsec vpn first_vpn ike ipsec-policy ipsec_pol
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec_prop
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set authentication-algorithm hmac-md5-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

5. Specify the IPsec Phase 2 proposal reference.

   ```
   [edit security ipsec policy ipsec_pol]
   user@host# set proposals ipsec_prop
   ```

6. Specify IPsec Phase 2 to use perfect forward secrecy (PFS) group1.

   ```
   [edit security ipsec policy ipsec_pol ]
   user@host# set perfect-forward-secrecy keys group1
   ```

7. Specify the IKE gateway.

   ```
   [edit security ipsec]
   user@host# set vpn first_vpn ike gateway gate
   ```

8. Specify the IPsec Phase 2 policy.

   ```
   [edit security ipsec]
   user@host# set vpn first_vpn ike ipsec-policy ipsec_pol
   ```

Results

From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security ipsec
proposal ipsec_prop {
   protocol esp;
   authentication-algorithm hmac-md5-96;
   encryption-algorithm 3des-cbc;
}
policy ipsec_pol {
perfect-forward-secrecy {
keys group1;
   proposals ipsec_prop;
}
   vpn first_vpn {
      ike {
         gateway gate;
         ipsec-policy ipsec_pol;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Security Policies for the Initiator

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone untrust policy pol1 match source-address
   any
set security policies from-zone trust to-zone untrust policy pol1 match destination-address
   any
set security policies from-zone trust to-zone untrust policy pol1 match application any
set security policies from-zone trust to-zone untrust policy pol1 then permit tunnel
   ipsec-vpn first_vpn
set security policies from-zone untrust to-zone trust policy pol1 match source-address
   any
set security policies from-zone untrust to-zone trust policy pol1 match destination-address
   any
set security policies from-zone untrust to-zone trust policy pol1 match application any
set security policies from-zone untrust to-zone trust policy pol1 then permit tunnel
   ipsec-vpn first_vpn
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies:

1. Create the security policy to permit traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy pol1 match source-address any
user@host# set policy pol1 match destination-address any
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set policy pol1 match application any
user@host# set policy pol1 then permit tunnel ipsec-vpn first_vpn
```

2. Create the security policy to permit traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy pol1 match source-address any
user@host# set policy pol1 match destination-address any
user@host# set policy pol1 match application any
user@host# set policy pol1 then permit tunnel ipsec-vpn first_vpn
```

Results From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
    policy pol1 {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
                tunnel {
            ipsec-vpn first_vpn;
        }
    }
}
from-zone untrust to-zone trust {
    policy pol1 {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
                tunnel {
            ipsec-vpn first_vpn;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Interface, Routing Options, and Security Zones for the Responder

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/2 unit 0 family inet address 13.168.11.100/24
set interfaces ge-0/0/3 unit 0 family inet address 10.2.99.1/24
set routing-options static route 10.1.99.0/24 next-hop 13.168.11.1
set routing-options static route 12.168.99.0/24 next-hop 13.168.11.1
set routing-options static route 1.1.100.0/24 next-hop 13.168.11.1
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/2.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interfaces, static routes, security zones, and security policies:

1. Configure Ethernet interface information.

   ```
   [edit]
   user@host# set interfaces ge-0/0/2 unit 0 family inet address 13.168.11.100/24
   user@host# set interfaces ge-0/0/3 unit 0 family inet address 10.2.99.1/24
   ```

2. Configure static route information.

   ```
   [edit]
   user@host# set routing-options static route 10.1.99.0/24 next-hop 13.168.11.1
   user@host# set routing-options static route 12.168.99.0/24 next-hop 13.168.11.1
   user@host# set routing-options static route 1.1.100.0/24 next-hop 13.168.11.1
   ```

3. Configure the untrust security zone.

   ```
   [edit ]
   user@host# set security zones security-zone untrust host-inbound-traffic protocols
     all
   ```

4. Assign an interface to the untrust security zone.

   ```
   [edit security zones security-zone untrust]
   user@host# set interfaces ge-0/0/2.0
   ```

5. Specify allowed system services for the untrust security zone.

   ```
   [edit security zones security-zone untrust]
   user@host# set host-inbound-traffic system-services all
   ```

6. Configure the trust security zone.

   ```
   [edit]
   user@host# set security zones security-zone trust host-inbound-traffic protocols
     all
   ```

7. Assign an interface to the trust security zone.

   ```
   [edit security zones security-zone trust]
   user@host# set interfaces ge-0/0/3.0
   ```

8. Specify allowed system services for the trust security zone.

   ```
   [edit security zones security-zone trust]
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set host-inbound-traffic system-services all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security zones** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
   unit 0 {
      family inet {
         address 13.168.11.100/24;
      }
   }
}
ge-0/0/3 {
   unit 0 {
      family inet {
         address 10.2.99.1/244;
      }
   }
}
```

```
[edit]
user@host# show routing-options
static {
   route 10.1.99.0/24 next-hop 13.168.11.1;
   route 12.168.99.0/24 next-hop 13.168.11.1;
   route 1.1.100.0/24 next-hop 13.168.11.1;
}
```

```
[edit]
user@host# show security zones
security-zone untrust {
   host-inbound-traffic {
      system-services {
         all;
         protocols {
         all;
      }
   }
   interfaces {
      ge-0/0/2.0;
   }
}
security-zone trust {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
      all;
      }
   }
   interfaces {
```

 Copyright © 2016, Juniper Networks, Inc.

```
        ge-0/0/3.0;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IKE for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike_prop authentication-method pre-shared-keys
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm md5
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_pol mode main
set security ike policy ike_pol proposals ike_prop
set security ike policy ike_pol pre-shared-key ascii-text "$ABC123"
set security ike gateway gate ike-policy ike_pol
set security ike gateway gate address 1.1.100.22
set security ike gateway gate dead-peer-detection probe-idle-tunnel
set security ike gateway gate external-interface ge-0/0/2.0
set security ike gateway gate local-identity inet 44.44.44.44
set security ike gateway gate remote-identity inet 11.11.11.11
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IKE:

1. Create the IKE Phase 1 proposal.

   ```
   [edit security ike]
   user@host# set proposal ike-phase1-proposal
   ```

2. Define the IKE proposal authentication method.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-method pre-shared-key
   ```

3. Define the IKE proposal Diffie-Hellman group.

   ```
   [edit security ike proposal ike_prop]
   user@host# set dh-group group2
   ```

4. Define the IKE proposal authentication algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set authentication-algorithm md5
   ```

5. Define the IKE proposal encryption algorithm.

   ```
   [edit security ike proposal ike_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

6. Create an IKE Phase 1 policy.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike]
user@host# set policy ike_pol
```

7. Set the IKE Phase 1 policy mode.

```
[edit security ike policy ike_pol]
user@host# set mode main
```

8. Specify a reference to the IKE proposal.

```
[edit security ike policy ike_pol]
user@host# set proposals ike_prop
```

9. Define the IKE Phase 1 policy authentication method.

```
[edit security ike policy ike_pol proposals ike_prop set security ike policy ike_pol
  pre-shared-key]
user@host# set ascii-text "$ABC123"
```

10. Create an IKE Phase 1 gateway and define its external interface.

```
[edit security ike]
user@host# set security ike gateway gate external-interface ge-0/0/2.0
```

11. Define the IKE Phase 1 policy reference.

```
[edit security ike gateway]
user@host# set gate ike-policy ike_pol
```

12. Create an IKE Phase 1 gateway address.

```
[edit security ike gateway]
user@host# set gate address 1.1.100.22
```

13. Set **local-identity** for the local peer (initiator).

```
[edit security ike gateway gate]
user@host# set local-identity inet 44.44.44.44
```

14. Set **remote-identity** for the responder. This is the responder's local identity.

```
[edit security ike gateway gate]
user@host# set remote-identity inet 11.11.11.11
```

15. Set dead peer detection to detect whether the peer is up or down.

```
[edit security ike gateway gate]
user@host# set dead-peer-detection probe-idle-tunnel
```

Results From configuration mode, confirm your configuration by entering the **show security ike** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ike
proposal ike_prop {
   authentication-method pre-shared-keys;
   dh-group group2;
   authentication-algorithm md5;
   encryption-algorithm 3des-cbc;
}
policy ike_pol {
```


Copyright © 2016, Juniper Networks, Inc.

```
    mode main;
    proposals ike_prop;
    pre-shared-key ascii-text "$ABC123";
  }
  gateway gate {
    ike-policy ike_pol;
    address 1.1.100.22;
    dead-peer-detection probe-idle-tunnel;
    external-interface ge-0/0/2.0;
    local-identity inet 44.44.44.44;
    remote-identity inet 11.11.11.11;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring IPsec for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec policy ipsec_pol perfect-forward-secrecy keys group1
set security ipsec policy ipsec_pol proposals ipsec_prop
set security ipsec vpn first_vpn ike gateway gate
set security ipsec vpn first_vpn ike ipsec-policy ipsec_pol
set security ipsec vpn first_vpn establish-tunnels immediately
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec:

1. Create an IPsec Phase 2 proposal.

   ```
   [edit]
   user@host# set security ipsec proposal ipsec_prop
   ```

2. Specify the IPsec Phase 2 proposal protocol.

   ```
   [edit security security ipsec proposal ipsec_prop]
   user@host# set protocol esp
   ```

3. Specify the IPsec Phase 2 proposal authentication algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set authentication-algorithm hmac-md5-96
   ```

4. Specify the IPsec Phase 2 proposal encryption algorithm.

   ```
   [edit security ipsec proposal ipsec_prop]
   user@host# set encryption-algorithm 3des-cbc
   ```

5. Set IPsec Phase 2 to use perfect forward secrecy (PFS) group1.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ipsec policy ipsec_pol]
user@host# set perfect-forward-secrecy keys group1
```

6. Create the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set policy ipsec_pol
```

7. Specify the IPsec Phase 2 proposal reference.

```
[edit security ipsec policy ipsec_pol]
user@host# set proposals ipsec_prop
```

8. Specify the IKE gateway.

```
[edit security ipsec]
user@host# set vpn first_vpn ike gateway gate
```

9. Specify the IPsec Phase 2 policy.

```
[edit security ipsec]
user@host# set vpn first_vpn ike ipsec-policy ipsec_pol
```

10. Specify that the tunnel be brought up immediately without a verification packet.

```
[edit security ipsec]
user@host# set security ipsec vpn first_vpn establish-tunnels immediately
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security ipsec
proposal ipsec_prop {
    protocol esp;
    authentication-algorithm hmac-md5-96;
    encryption-algorithm 3des-cbc;
}
policy ipsec_pol {
    perfect-forward-secrecy {
        keys group1;
    }
    proposals ipsec_prop;
}
vpn first_vpn {
    ike {
        gateway gate;
        ipsec-policy ipsec_pol;
        establish-tunnels immediately;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring Security Policies for the Responder

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security policies from-zone trust to-zone untrust policy pol1 match source-address
  any
set security policies from-zone trust to-zone untrust policy pol1 match destination-address
  any
set security policies from-zone trust to-zone untrust policy pol1 match application any
set security policies from-zone trust to-zone untrust policy pol1 then permit tunnel
  ipsec-vpn first_vpn
set security policies from-zone untrust to-zone trust policy pol1 match source-address
  any
set security policies from-zone untrust to-zone trust policy pol1 match destination-address
  any
set security policies from-zone untrust to-zone trust policy pol1 match application any
set security policies from-zone untrust to-zone trust policy pol1 then permit tunnel
  ipsec-vpn first_vpn
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security policies:

1. Create the security policy to permit traffic from the trust zone to the untrust zone.

   ```
   [edit security policies from-zone trust to-zone untrust]
   user@host# set policy pol1 match source-address any
   user@host# set policy pol1 match destination-address any
   user@host# set policy pol1 match application any
   user@host# set policy pol1 then permit tunnel ipsec-vpn first_vpn
   ```

2. Create the security policy to permit traffic from the untrust zone to the trust zone.

   ```
   [edit security policies from-zone untrust to-zone trust]
   user@host# set policy pol1 match source-address any
   user@host# set policy pol1 match destination-address any
   user@host# set policy pol1 match application any
   user@host# set policy pol1 then permit tunnel ipsec-vpn first_vpn
   ```

Results

From configuration mode, confirm your configuration by entering the **show security policies** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy pol1 {
    match {
      source-address any;
      destination-address any;
```

Copyright © 2016, Juniper Networks, Inc.

```
        application any;
      }
      then {
        permit;
          tunnel {
        ipsec-vpn first_vpn;
      }
    }
  }
  from-zone untrust to-zone trust {
    policy pol1 {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
          tunnel {
        ipsec-vpn first_vpn;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the IKE Phase 1 Status for the Initiator on page 6292
- Verifying IPsec Security Associations for the Initiator on page 6294
- Verifying the IKE Phase 1 Status for the Responder on page 6295
- Verifying IPsec Security Associations for the Responder on page 6297

Verifying the IKE Phase 1 Status for the Initiator

Purpose Verify the IKE Phase 1 status.

Action

NOTE: Before starting the verification process, you must send traffic from a host in the 10.1.99.0 network to a host in the 10.2.99.0 network. For route-based VPNs, traffic can be initiated by the SRX Series device through the tunnel. We recommend that when testing IPsec tunnels, test traffic be sent from a separate device on one side of the VPN to a second device on the other side of the VPN. For example, initiate a ping operation from 10.1.99.2 to 10.2.99.2.

Copyright © 2016, Juniper Networks, Inc.

From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5137403 UP     b3a24bc00e963c51  7bf96bcc6230e484  Main           1.1.100.23

user@host> show security ike security-associations index 1 detail
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

1400579286 UP  487cfb570908425c  7710c8487f9ff20c  Main           1.1.100.22


{primary:node0}[edit]
root@poway# run show security ike security-associations detail
node0:

IKE peer 1.1.100.22, Index 1400579286,
  Location: FPC 5, PIC 0, KMD-Instance 4
  Role: Initiator, State: UP
  Initiator cookie: 487cfb570908425c, Responder cookie: 7710c8487f9ff20c
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 13.168.11.100:4500, Remote: 1.1.100.22:4500
  Lifetime: Expires in 28622 seconds
  Peer ike-id: 44.44.44.44
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-md5-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-md5
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 0
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote address—Verify that the remote IP address is correct and that port 4500 is being used for peer-to-peer communication.
- Role initiator state
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.

Copyright © 2016, Juniper Networks, Inc.

- Both peers in the IPsec SA pair are using port 4500, which indicates that NAT-T is implemented. (NAT-T uses port 4500 or another random high-numbered port.)
- Peer IKE ID—Verify the remote (responder) address is correct. In this example, the address is 44.44.44.44.
- Local identity and remote identity—Verify these are correct.

- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations** command lists additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying IPsec Security Associations for the Initiator

Purpose Verify the IPsec status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <2    ESP:3des/md5    2bf24122 3390/ unlim   -   root 4500  1.1.100.23
  >2    ESP:3des/md5    2baef146 3390/ unlim   -   root 4500  1.1.100.23

user@host> show security ipsec security-associations detail
```


Copyright © 2016, Juniper Networks, Inc.

```
Local Gateway: 12.168.99.100, Remote Gateway: 1.1.100.23
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
    DF-bit: clear
    Policy-name: pol1

    Location: FPC 5, PIC 0, KMD-Instance 4
    Direction: inbound, SPI: 2bf24122, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3388 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2801 seconds
    Mode: Tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Location: FPC 5, PIC 0, KMD-Instance 4
    Direction: outbound, SPI: 2baef146, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3388 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2801 seconds
    Mode: Tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-md5-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has a NAT address of 1.1.100.23.
- Both peers in the IPsec SA pair are using port 4500, which indicates that NAT-T is implemented. (NAT-T uses port 4500 or another random high-numbered port.).
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3390/ unlimited value indicates that the Phase 2 lifetime expires in 3390 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

Verifying the IKE Phase 1 Status for the Responder

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index *index_number* detail** command.

```
user@host> show security ike security-associations
```

```
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5802591 UP     d31d6833108fd69f  9ddfe2ce133086aa  Main           1.0.0.1

user@host> show security ike security-associations index 1 detail
IKE peer 1.1.100.23, Index 1400579287,
  Location: FPC 5, PIC 0, KMD-Instance 4
  Role: Responder, State: UP
  Initiator cookie: 487cfb570908425c, Responder cookie: 7710c8487f9ff20c
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 12.168.99.100:4500, Remote: 1.1.100.23:4500
  Lifetime: Expires in 28587 seconds
  Peer ike-id: 11.11.11.11
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-md5-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-md5
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 71.1.1.1:4500, Remote: 1.0.0.1:4500
    Local identity: branch_natt1@example.net
    Remote identity: limits_natt1@example.net
    Flags: IKE SA is created
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index detail** command to get more information about the SA.
- Remote address—Verify that the remote IP address is correct and that port 4500 is being used for peer-to-peer communication.
- Role responder state
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.
  - Peer IKE ID—Verify the local (initiator) address for the peer is correct. In this example, the address is 11.11.11.11.
  - Local identity and remote identity—Verify these are correct.
- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

 Copyright © 2016, Juniper Networks, Inc.

- External interfaces (the interface must be the one that receives IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations** command lists additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: **Troubleshooting is best performed on the peer using the responder role.**

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying IPsec Security Associations for the Responder

Purpose Verify the IPsec status.

Action From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index *index_number* detail** command.

```
user@host> show security ipsec security-associations
Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <131073 ESP:3des/sha1 a5224cd9 3571/ unlim   -   root 4500  1.0.0.1
  >131073 ESP:3des/sha1 82a86a07 3571/ unlim   -   root 4500  1.0.0.1

user@host> show security ipsec security-associations detail
Virtual-system: root
  Local Gateway: 71.1.1.1, Remote Gateway: 1.0.0.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
    DF-bit: clear
    Direction: inbound, SPI: a5224cd9, AUX-SPI: 0
                               , VPN Monitoring: -
    Hard lifetime: Expires in 3523 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2923 seconds
    Mode: Tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
```

Copyright © 2016, Juniper Networks, Inc.

```
Anti-replay service: counter-based enabled, Replay window size: 64

Direction: outbound, SPI: 82a86a07, AUX-SPI: 0
                          , VPN Monitoring: -
Hard lifetime: Expires in 3523 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 2923 seconds
Mode: Tunnel, Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has a NAT address of 1.0.0.1.
- Both peers in the IPsec SA pair are using port 4500, which indicates that NAT-T is implemented. (NAT-T uses port 4500 or another random high-numbered port.)
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The 3571/ unlim value indicates that the Phase 2 lifetime expires in 3571 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- VPN monitoring is not enabled for this SA, as indicated by a hyphen in the Mon column. If VPN monitoring is enabled, U indicates that monitoring is up, and D indicates that monitoring is down.
- The virtual system (vsys) is the root system, and it always lists 0.

Related Documentation

- IPsec VPN Overview on page 6041
- Understanding NAT-T on page 6243
- Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device on page 6244

Example: Configuring NAT-T with Dynamic Endpoint VPN

This example shows how to configure a route-based VPN where the IKEv2 initiator is a dynamic endpoint behind a NAT device.

- Requirements on page 6299
- Overview on page 6299
- Configuration on page 6300
- Verification on page 6312

Copyright © 2016, Juniper Networks, Inc.

Requirements

This example uses the following hardware and software components:

- Two SRX Series devices configured in a chassis cluster
- One SRX Series device providing NAT
- One SRX Series device providing branch office network access
- Junos OS Release 12.1X46-D10 or later for IKEv2 NAT-T support

Overview

In this example, an IPsec VPN is configured between the branch office (IKEv2 initiator) and headquarters (IKEv2 responder) to secure network traffic between the two locations. The branch office is located behind the NAT device. The branch office address is assigned dynamically and is unknown to the responder. The initiator is configured with the remote identity of the responder for tunnel negotiation. This configuration establishes a dynamic endpoint VPN between the peers across the NAT device.

Figure 265 shows an example of a topology with NAT-Traversal (NAT-T) and dynamic endpoint VPN.

Figure 265: NAT-T with Dynamic Endpoint VPN

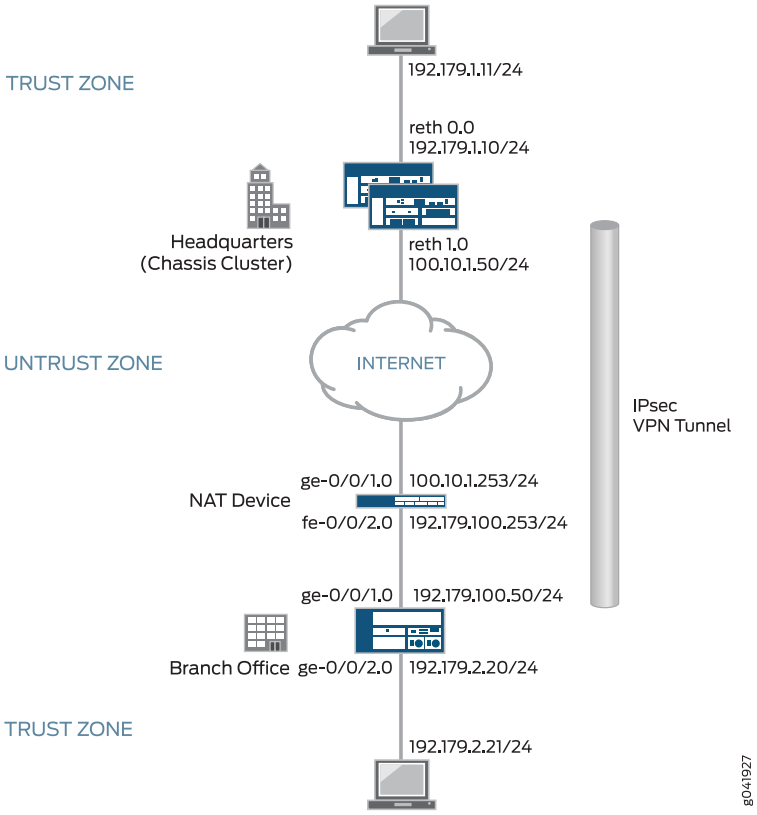

In this example, the initiator's IP address, 192.179.100.50, which has been dynamically assigned to the device, is hidden by the NAT device and translated to 100.10.1.253.

The following configuration options apply in this example:

- The local identity configured on the initiator must match the remote gateway identity configured on the responder.
- Phase 1 and Phase 2 options must match between the initiator and responder.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

NOTE: Starting with Junos OS 12.1X46-D10, the default value for the **nat-keepalive** option configured at the **[edit security ike gateway *gateway-name*]** hierarchy level has been changed from 5 seconds to 20 seconds.

NOTE: In SRX5600 and SRX5800 devices, IKE negotiations involving NAT traversal do not work if the IKE peer is behind a NAT device that will change the source IP address of the IKE packets during the negotiation. For example, if the NAT device is configured with DIP, it changes the source IP because the IKE protocol switches the UDP port from 500 to 4500.

Configuration

- Configuring the Branch Office Device (IKEv2 Initiator) on page 6300
- Configuring the NAT Device on page 6304
- Configuring the Headquarters Device (IKEv2 Responder) on page 6307

Configuring the Branch Office Device (IKEv2 Initiator)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 192.179.100.50/24
set interfaces ge-0/0/2 unit 0 family inet address 192.179.2.20/24
set interfaces st0 unit 0 family inet address 172.168.100.1/16
set routing-options static route 192.179.1.0/24 next-hop st0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/2.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security ike proposal IKE_PROP authentication-method pre-shared-keys
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL pre-shared-key ascii-text "$ABC123"
set security ike gateway HQ_GW ike-policy IKE_POL
set security ike gateway HQ_GW address 100.10.1.50
set security ike gateway HQ_GW local-identity hostname branch.example.net
set security ike gateway HQ_GW external-interface ge-0/0/1.0
set security ike gateway HQ_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn HQ_VPN bind-interface st0.0
set security ipsec vpn HQ_VPN ike gateway HQ_GW
set security ipsec vpn HQ_VPN ike ipsec-policy IPSEC_POL
set security ipsec vpn HQ_VPN establish-tunnels immediately
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the branch office device:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 192.179.100.50/24
user@host# set ge-0/0/2 unit 0 family inet address 192.179.2.20/24
user@host# set st0 unit 0 family inet address 172.168.100.1/16
```

2. Configure routing options.

```
[edit routing-options]
user@host# set static route 192.179.1.0/24 next-hop st0.0
```

3. Configure zones.

```
[edit security zones security-zones trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/2.0

[edit security zones security-zones untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host#set interfaces st0.0
```

4. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
user@host# set proposals IKE_PROP
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security ike gateway HQ_GW]
user@host# set ike-policy IKE_POL
user@host# set address 100.10.1.50
user@host# set local-identity hostname branch.example.net
user@host# set external-interface ge-0/0/1.0
user@host# set version v2-only
```

5. Configure Phase 2 options.

```
[edit security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy IPSEC_POL]
user@host# set proposals IPSEC_PROP
user@host# set perfect-forward-secrecy keys group5

[edit security ipsec vpn HQ_VPN]
user@host# set bind-interface st0.0
user@host# set ike gateway HQ_GW
user@host# set ike ipsec-policy IPSEC_POL
user@host# set establish-tunnels immediately
```

6. Configure the security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security zones**, **show security ike**, **show security ipsec**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            address 192.179.100.50/24;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
```

 Copyright © 2016, Juniper Networks, Inc.

```
        address 192.179.2.20/24;
      }
    }
  }
  st0 {
    unit 0 {
      family inet {
        address 172.168.100.1/16;
      }
    }
  }
  [edit]
  user@host# show routing-options
  static {
    route 192.179.1.0/24 next-hop st0.0;
  }
  [edit]
  user@host# show security zones
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/2.0;
    }
  }
  security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/1.0;
      st0.0;
    }
  }
  [edit]
  user@host# show security ike
  proposal IKE_PROP {
    authentication-method pre-shared-keys;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
  }
  policy IKE_POL {
    proposals IKE_PROP;
    pre-shared-key ascii-text "$ABC123”
```

Copyright © 2016, Juniper Networks, Inc.

```
}
gateway HQ_GW{
   ike-policy IKE_POL;
   address 100.10.1.50;
   local-identity hostname branch.example.net;
   external-interface ge-0/0/1.0;
   version v2-only;
}
[edit]
user@host# show security ipsec
proposal IPSEC_PROP {
   protocol esp;
   authentication-algorithm hmac-sha1-96;
   encryption-algorithm aes-256-cbc;
}
policy IPSEC_POL {
   perfect-forward-secrecy {
      keys group5;
   }
   proposals IPSEC_PROP;
}
vpn HQ_VPN {
   bind-interface st0.0;
   ike {
      gateway HQ_GW;
      ipsec-policy IPSEC_POL;
   }
   establish-tunnels immediately;
}
[edit]
user@host# show security policies
default-policy {
   permit-all;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the NAT Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 100.10.1.253/24
set interfaces fe-0/0/2 unit 0 family inet address 192.179.100.253/24
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/1.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/2.0
set security nat source rule-set DYNAMIC from zone untrust
set security nat source rule-set DYNAMIC to zone trust
set security nat source rule-set DYNAMIC rule R2R3 match source-address 0.0.0.0/0
```


Copyright © 2016, Juniper Networks, Inc.

```
set security nat source rule-set DYNAMIC rule R2R3 then source-nat interface
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the intermediate router providing NAT:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 100.10.1.253/24
user@host# set fe-0/0/2 unit 0 family inet address 192.179.100.253/24
```

2. Configure zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/2.0
```

3. Configure NAT.

```
[edit security nat source rule-set DYNAMIC]
user@host# set from zone untrust
user@host# set to zone trust
user@host# set rule R2R3 match source-address 0.0.0.0/0
user@host# set rule R2R3 then source-nat interface
```

4. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show security zones**, **show security nat source**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            address 100.10.1.253/24;
        }
    }
}
fe-0/0/2 {
    unit 0 {
        family inet {
```

Copyright © 2016, Juniper Networks, Inc.

```
        address 192.179.100.253/24;
      }
    }
  }
  [edit]
  user@host# show security zones
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/1.0;
    }
  }
  security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      fe-0/0/2.0;
    }
  }
  [edit]
  user@host# show security nat source
  rule-set DYNAMIC {
    from zone untrust;
    to zone trust;
    rule R2R3 {
      match {
        source-address 0.0.0.0/0;
      }
      then {
        source-nat {
          interface;
        }
      }
    }
  }
  [edit]
  user@host# show security policies
  default-policy {
    permit-all;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring the Headquarters Device (IKEv2 Responder)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis cluster reth-count 5
set chassis cluster redundancy-group 1 node 0 priority 220
set chassis cluster redundancy-group 1 node 1 priority 149
set chassis cluster redundancy-group 1 interface-monitor ge-0/0/1 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/0/1 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-0/0/2 weight 255
set chassis cluster redundancy-group 1 interface-monitor ge-8/0/2 weight 255
set interfaces ge-0/0/1 gigether-options redundant-parent reth0
set interfaces ge-0/0/2 gigether-options redundant-parent reth1
set interfaces ge-8/0/1 gigether-options redundant-parent reth0
set interfaces ge-8/0/2 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 192.179.1.10/24
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 100.10.1.50/24
set interfaces st0 unit 0 family inet address 172.168.100.2/16
set routing-options static route 192.179.2.0/24 next-hop st0.0
set routing-options static route 192.179.100.0/24 next-hop 100.10.1.253
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces reth1.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces reth0.0
set security ike proposal IKE_PROP authentication-method pre-shared-keys
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL pre-shared-key ascii-text "$ABC123"
set security ike gateway Branch_GW ike-policy IKE_POL
set security ike gateway Branch_GW dynamic hostname branch.example.net
set security ike gateway Branch_GW dead-peer-detection optimized
set security ike gateway Branch_GW external-interface reth1.0
set security ike gateway Branch_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn Branch_VPN bind-interface st0.0
set security ipsec vpn Branch_VPN ike gateway Branch_GW
set security ipsec vpn Branch_VPN ike ipsec-policy IPSEC_POL
set security policies default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure two nodes as the chassis cluster.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 1 node 0 priority 220
user@host# set redundancy-group 1 node 1 priority 149
user@host# set redundancy-group 1 interface-monitor ge-0/0/1 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/0/1 weight 255
user@host# set redundancy-group 1 interface-monitor ge-0/0/2 weight 255
user@host# set redundancy-group 1 interface-monitor ge-8/0/2 weight 255
```

2. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 gigether-options redundant-parent reth0
user@host# set ge-0/0/2 gigether-options redundant-parent reth1
user@host# set ge-8/0/1 gigether-options redundant-parent reth0
user@host# set ge-8/0/2 gigether-options redundant-parent reth1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 192.179.1.10/24
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 100.10.1.50/24
user@host# set st0 unit 0 family inet address 172.168.100.2/16
```

3. Configure routing options.

```
[edit routing-options]
user@host# set static route 192.179.2.0/24 next-hop st0.0
user@host# set static route 192.179.100.0/24 next-hop 100.10.1.253
```

4. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic protocols all
user@host# set host-inbound-traffic system-services all
user@host# set interfaces st0.0
user@host# set interfaces reth1.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces reth0.0
```

5. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
user@host# set proposals IKE_PROP
user@host# set pre-shared-key ascii-text "$ABC123"
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike gateway Branch_GW]
user@host# set ike-policy IKE_POL
user@host# set dynamic hostname branch.example.net
user@host# set dead-peer-detection optimized
user@host# set external-interface reth1.0
user@host# set version v2-only
```

6. Configure Phase 2 options.

```
[edit security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy IPSEC_POL]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals IPSEC_PROP

[edit security ipsec vpn Branch_VPN]
user@host# set bind-interface st0.0
user@host# set ike gateway Branch_GW
user@host# set ike ipsec-policy IPSEC_POL
```

7. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show chassis cluster**, **show interfaces**, **show routing-options**, **show security zones**, **show security ike**, **show security ipsec**, and **show security policies** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis cluster
reth-count 5;
redundancy-group 1 {
  node 0 priority 220;
  node 1 priority 149;
  interface-monitor {
    ge-0/0/1 weight 255;
    ge-8/0/1 weight 255;
    ge-0/0/2 weight 255;
    ge-8/0/2 weight 255;
  }
}
[edit]
user@host# show interfaces
ge-0/0/1 {
  gigether-options {
    redundant-parent reth0;
  }
}
ge-0/0/2 {
  gigether-options {
    redundant-parent reth1;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
  ge-8/0/1 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-8/0/2 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  reth0 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 192.179.1.10/24;
      }
    }
  }
  reth1 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 100.10.1.50/24;
      }
    }
  }
  st0 {
    unit 0{
      family inet {
        address 172.168.100.2/16;
      }
    }
  }
  [edit]
  user@host# show routing-options
  static {
    route 192.179.2.0/24 next-hop st0.0;
    route 192.179.100.0/24 next-hop 100.10.1.253;
  }
  [edit]
  user@host# show security zones
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    interfaces {
      reth0.0;
    }
  }
  security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      st0.0;
      reth1.0;
    }
  }
  [edit]
  user@host# show security ike
  proposal IKE_PROP {
    authentication-method pre-shared-keys;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
  }
  policy IKE_POL {
    proposals IKE_PROP;
    pre-shared-key ascii-text "$ABC123"
  }
  gateway Branch_GW {
    ike-policy IKE_POL;
    dynamic hostname branch.example.net;
    dead-peer-detection optimized;
    external-interface reth1.0;
    version v2-only;
  }
  [edit]
  user@host# show security ipsec
  proposal IPSEC_PROP {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
  }
  policy IPSEC_POL {
    perfect-forward-secrecy {
      keys group5;
    }
    proposals IPSEC_PROP;
  }
  vpn Branch_VPN {
    bind-interface st0.0;
    ike {
      gateway Branch_GW;
      ipsec-policy IPSEC_POL;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
```

Verification

Confirm that the configuration is working properly.

- Verifying the IKE Phase 1 Status for the Responder on page 6312
- Verifying IPsec Security Associations for the Responder on page 6313

Verifying the IKE Phase 1 Status for the Responder

Purpose Verify the IKE Phase 1 status.

Action From operational mode on node 0, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations detail** command.

```
user@host# show security ike security-associations
node0:
Index      State  Initiator cookie   Responder cookie  Mode    Remote Address
1367024684 UP     f82c54347e2f3fb1   020e28e1e4cae003  IKEv2    100.10.1.253

user@host# show security ike security-associations detail
node0:
IKE peer 100.10.1.253, Index 1367024684, Gateway Name: Branch_GW
  Location: FPC 5, PIC 0, KMD-Instance 2
  Role: Responder, State: UP
  Initiator cookie: f82c54347e2f3fb1, Responder cookie: 020e28e1e4cae003
  Exchange type: IKEv2, Authentication method: Pre-shared-keys
  Local: 100.10.1.50:4500, Remote: 100.10.1.253:2541
  Lifetime: Expires in 3593 seconds
  Peer ike-id: branch.example.net
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                  683
   Output bytes  :                  400
   Input  packets:                    2
   Output packets:                    1
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration.

If SAs are listed, review the following information:

Copyright © 2016, Juniper Networks, Inc.

- Index—This value is unique for each IKE SA, which you can use in the **show security ike security-associations index *index_id* detail** command to get more information about the SA.
- Remote address—Verify that the local IP address is correct and that port 4500 is being used for peer-to-peer communication.
- Role responder state
  - Up—The Phase 1 SA has been established.
  - Down—There was a problem establishing the Phase 1 SA.
  - Peer IKE ID—Verify the address is correct.
  - Local identity and remote identity—Verify these addresses are correct.
- Mode—Verify that the correct mode is being used.

Verify that the following are correct in your configuration:

- External interfaces (the interface must be the one that sends IKE packets)
- IKE policy parameters
- Preshared key information
- Phase 1 proposal parameters (must match on both peers)

The **show security ike security-associations** command lists additional information about security associations:

- Authentication and encryption algorithms used
- Phase 1 lifetime
- Traffic statistics (can be used to verify that traffic is flowing properly in both directions)
- Role information

NOTE: Troubleshooting is best performed on the peer using the responder role.

- Initiator and responder information
- Number of IPsec SAs created
- Number of Phase 2 negotiations in progress

Verifying IPsec Security Associations for the Responder

Purpose Verify the IPsec status.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode on node 0, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations detail** command.

```
user@host# show security ipsec security-associations
node0
  Total active tunnels: 1
  ID        Algorithm            SPI      Life:sec/kb  Mon lsys Port Gateway
  <77856771 ESP:aes-cbc-256/sha1 4ad5af40 7186/unlim   - root   2541 100.10.1.253

  >77856771 ESP:aes-cbc-256/sha1 5bb0a5ee 7186/unlim   - root   2541 100.10.1.253

user@host# show security ipsec security-associations detail
node0
  ID: 77856771 Virtual-system: root, VPN Name: Branch_VPN
  Local Gateway: 100.10.1.50, Remote Gateway: 100.10.1.253
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
    DF-bit: clear
    Bind-interface: st0.0

  Port: 2541, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 608a29
  Tunnel Down Reason: SA not initiated
    Location: FPC 5, PIC 0, KMD-Instance 2
    Direction: inbound, SPI: 4ad5af40, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 7182 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 6587 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output from the **show security ipsec security-associations** command lists the following information:

- The remote gateway has an IP address of 100.10.1.253.
- The SPIs, lifetime (in seconds), and usage limits (or lifesize in KB) are shown for both directions. The lifetime value indicates that the Phase 2 lifetime expires in 7186 seconds, and that no lifesize has been specified, which indicates that it is unlimited. Phase 2 lifetime can differ from Phase 1 lifetime, as Phase 2 is not dependent on Phase 1 after the VPN is up.
- The virtual system (vsys) is the root system, and it always lists 0.

The output from the **show security ipsec security-associations index *index_id* detail** command lists the following information:

- The local identity and remote identity make up the proxy ID for the SA.

  A proxy ID mismatch is one of the most common causes for a Phase 2 failure. If no IPsec SA is listed, confirm that Phase 2 proposals, including the proxy ID settings, match for both peers. For route-based VPNs, the default proxy ID is local=0.0.0.0/0, remote=0.0.0.0/0, and service=any. Issues can occur with multiple route-based VPNs from the same peer IP. In this case, a unique proxy ID for each IPsec SA must be

 Copyright © 2016, Juniper Networks, Inc.

specified. For some third-party vendors, the proxy ID must be manually entered to match.

- Another common reason for Phase 2 failure is not specifying the ST interface binding. If IPsec cannot complete, check the kmd log or set traceoptions.

Related Documentation

- IPsec VPN Overview on page 6041
- Security Policies Overview on page 1085
- Understanding NAT-T on page 6243
- Example: Configuring a Route-Based VPN with Only the Responder Behind a NAT Device on page 6244
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 83

Configuring IPsec VPN Tunnels with Chassis Clusters

- Configuring IPsec VPN Tunnels with Chassis Clusters on page 6319

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 262

Configuring IPsec VPN Tunnels with Chassis Clusters

- Understanding Dual Active-Backup IPsec VPN Chassis Clusters on page 6319
- Understanding Loopback Interface for a High Availability VPN on page 6320
- Example: Configuring Redundancy Groups for Loopback Interfaces on page 6321

Understanding Dual Active-Backup IPsec VPN Chassis Clusters

In an active/passive chassis cluster, all VPN tunnels terminate on the same node, as shown in Figure 266.

Figure 266: Active/Passive Chassis Cluster with IPsec VPN Tunnels

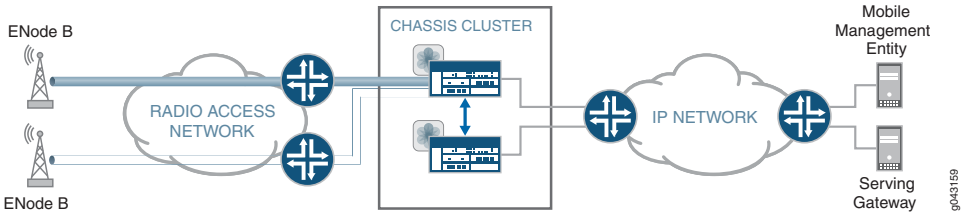

In an active/active chassis cluster, VPN tunnels can terminate on either node. Both nodes in the chassis cluster can actively pass traffic through VPN tunnels on both nodes at the same time, as shown in Figure 267. This deployment is known as *dual active-backup IPsec VPN chassis clusters*.

Figure 267: Dual Active-Backup IPsec VPN Chassis Clusters

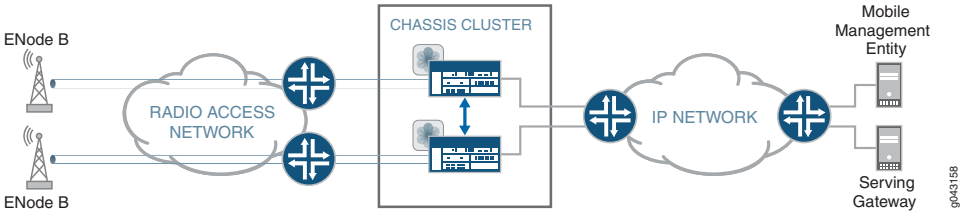

Copyright © 2016, Juniper Networks, Inc.

The following features are supported with dual active-backup IPsec VPN chassis clusters:

- Route-based VPNs only. Policy-based VPNs are not supported.
- IKEv2 only. IKEv1 is not supported.
- IKE and secure tunnel interfaces (st0) in virtual routers.
- Network Address Translation-Traversal (NAT-T).
- VPN monitoring.
- Dead peer detection.
- In-service software upgrade (ISSU).
- Insertion of Services Processing Cards (SPCs) on a chassis cluster device without disrupting the traffic on the existing VPN tunnels. See "Understanding VPN Support for Inserting Services Processing Cards" on page 6068.
- Dynamic routing protocols.
- Secure tunnel interfaces (st0) configured in point-to-multipoint mode.
- AutoVPN with st0 interfaces in point-to-point mode with traffic selectors.
- IPv4-in-IPv4, IPv6-in-IPv4, IPv6-in-IPv6 and IPv4-in-IPv6 tunnel modes.
- Fragmented traffic.
- The loopback interface can be configured as the external interface for the VPN.

Dual active-backup IPsec VPN chassis clusters are only supported on high-end SRX Series chassis clusters. Dual active-backup IPsec VPN chassis clusters cannot be configured with the following features:

- VPNs with manual or preshared keys.
- Z-mode flows. Z-mode flows occur when traffic enters an interface on a chassis cluster node, passes through the fabric link, and exits through an interface on the other cluster node.

Related Documentation

- *Chassis Cluster Overview*
- *Preparing Your Equipment for Chassis Cluster Formation*

Understanding Loopback Interface for a High Availability VPN

An Internet Key Exchange (IKE) gateway needs an external interface to communicate with a peer device. In a high availability chassis cluster setup, the node on which the external interface is active selects a Services Processing Unit (SPU) to support the VPN tunnel. IKE and IPsec packets are processed on that SPU. Therefore, the active external interface decides the anchor SPU.

In a chassis cluster setup, the external interface is a redundant Ethernet interface. A redundant Ethernet interface can go down when its physical (child) interfaces are down.

Copyright © 2016, Juniper Networks, Inc.

You can configure a loopback interface as an alternate physical interface to reach the peer gateway.

This feature allows the loopback interface to be configured for any redundancy group. This redundancy group configuration is only checked for VPN packets, because only VPN packets must find the anchor SPU through the active interface.

On branch SRX Series devices, the lo0 pseudointerface can be configured in any redundancy group; for example, RG0, RG1, RG2, and so on. However, on high-end SRX Series devices, the lo0 pseudointerface cannot be configured in RG0 when it is used as an IKE gateway external interface. Because a VPN is only supported in an active-passive chassis cluster environment on high-end SRX Series devices, the lo0 pseudointerface can be configured in such a setup for RG1. In a chassis cluster setup, the node on which the external interface is active selects an SPU to anchor the VPN tunnel. IKE and IPsec packets are processed on that SPU. Thus an active external interface decides the anchor SPU.

Related Documentation

- IPsec VPN Overview on page 6041

Example: Configuring Redundancy Groups for Loopback Interfaces

This example shows how to configure a redundancy group (RG) for a loopback interface in order to prevent VPN failure. Redundancy groups are used to bundle interfaces into a group for failover purpose in a chassis cluster setup.

- Requirements on page 6321
- Overview on page 6321
- Configuration on page 6323
- Verification on page 6326

Requirements

This example uses the following hardware and software:

- A pair of supported chassis cluster SRX Series devices
- An SSG140 device or equivalent
- Two switches
- Junos OS Release 12.1x44-D10 or later for SRX Series Services Gateways

Before you begin:

Understand chassis cluster redundant Ethernet interfaces. See *Understanding Chassis Cluster Redundant Ethernet Interfaces*.

Overview

An Internet Key Exchange (IKE) gateway needs an external interface to communicate with a peer device. In a chassis cluster setup, the node on which the external interface is active selects a Services Processing Unit (SPU) to support the VPN tunnel. IKE and IPsec

Copyright © 2016, Juniper Networks, Inc.

packets are processed on that SPU. Therefore, the active external interface decides the anchor SPU.

In a chassis cluster setup, the external interface is a redundant Ethernet interface. A redundant Ethernet interface can go down when its physical (child) interfaces are down. You can configure a loopback interface as an alternative physical interface to reach the peer gateway. Loopback interfaces can be configured on any redundancy group. This redundancy group configuration is only checked for VPN packets, because only VPN packets must find the anchor SPU through the active interface.

NOTE: You must configure lo0.x in a custom virtual router, since lo0.0 is in the default virtual router and only one loopback interface is allowed in a virtual router.

Figure 268 shows an example of a loopback chassis cluster VPN topology. In this topology, the SRX Series chassis cluster device is located in Sunnyvale, California. The SRX Series chassis cluster device works as a single gateway in this setup. The SSG Series device (or a third-party device) is located in Chicago, Illinois. This device acts as a peer device to the SRX chassis cluster and it helps to build a VPN tunnel.

Copyright © 2016, Juniper Networks, Inc.

Figure 268: Loopback Interface for Chassis Cluster VPN

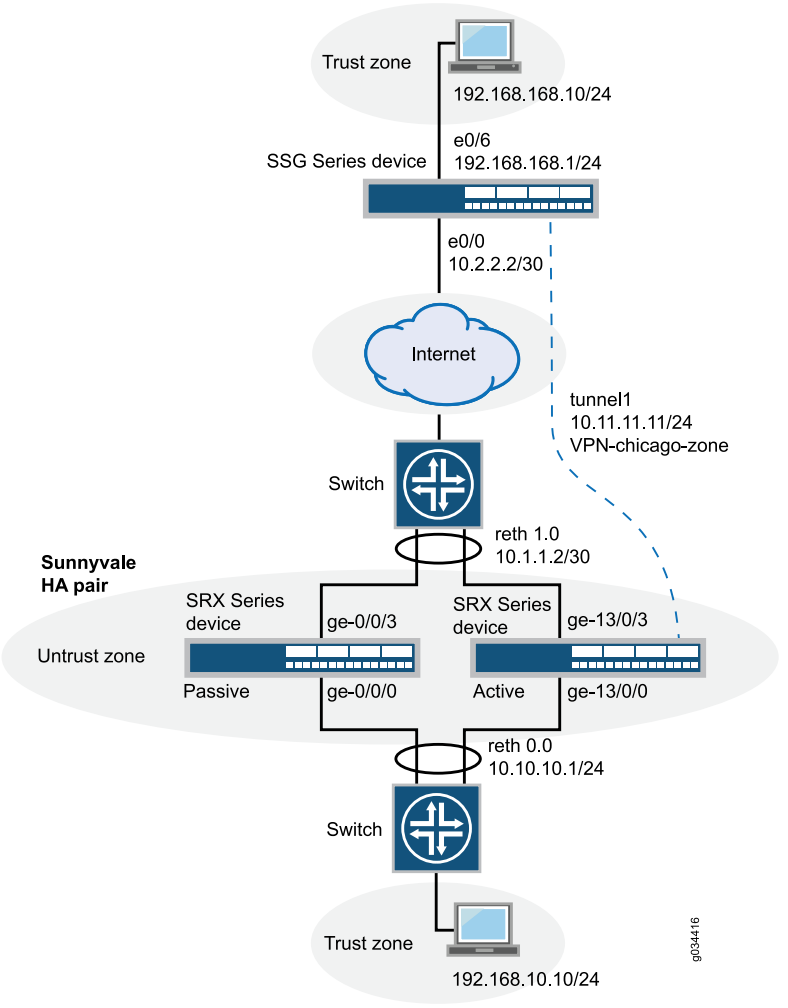

lo0.1 10.3.3.3/30 Untrust zone	A logical interface on loopback that may be active on either node in the cluster depending on the activeness of its RG.
st0.0 10.11.11.10/24 vpn-chicago-zone	A logical interface on the secure tunnel interface for IPSec VPN tunnel.

Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces lo0 redundant-pseudo-interface-options redundancy-group 1
set interfaces lo0 unit 1 family inet address 10.3.3.3/30
```

Copyright © 2016, Juniper Networks, Inc.

```
set routing-instances vr1 instance-type virtual-router
set routing-instances vr1 interface lo0.1
set routing-instances vr1 interface reth0.0
set routing-instances vr1 interface reth1.0
set routing-instances vr1 interface st0.0
set routing-instances vr1 routing-options static route 192.168.168.1/24 next-hop st0.0
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposal-set standard
set security ike policy ike-policy1 pre-shared-key ascii-text "$ABC123t"
set security ike gateway t-ike-gate ike-policy ike-policy1
set security ike gateway t-ike-gate address 10.2.2.2
set security ike gateway t-ike-gate external-interface lo0.1
set security ipsec proposal p2-std-p1 authentication-algorithm hmac-sha1-96
set security ipsec proposal p2-std-p1 encryption-algorithm 3des-cbc
set security ipsec proposal p2-std-p1 lifetime-seconds 180
set security ipsec proposal p2-std-p2 authentication-algorithm hmac-sha1-96
set security ipsec proposal p2-std-p2 encryption-algorithm aes-128-cbc
set security ipsec proposal p2-std-p2 lifetime-seconds 180
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group2
set security ipsec policy vpn-policy1 proposals p2-std-p1
set security ipsec policy vpn-policy1 proposals p2-std-p2
set security ipsec vpn t-ike-vpn bind-interface st0.0
set security ipsec vpn t-ike-vpn ike gateway t-ike-gate
set security ipsec vpn t-ike-vpn ike proxy-identity local 10.10.10.1/24
set security ipsec vpn t-ike-vpn ike proxy-identity remote 192.168.168.1/24
set security ipsec vpn t-ike-vpn ike ipsec-policy vpn-policy1
```

Step-by-Step Procedure

To configure a redundancy group for a loopback interface:

1. Configure the loopback interface in one redundancy group.

   ```
   [edit interfaces]
   user@host# set lo0 redundant-pseudo-interface-options redundancy-group 1
   ```

2. Configure the IP address for the loopback interface.

   ```
   [edit interfaces]
   user@host# set lo0 unit 1 family inet address 10.3.3.3/30
   ```

3. Configure routing options.

   ```
   [edit routing-instances]
   user@host# set vr1 instance-type virtual-router
   user@host# set vr1 interface lo0.1
   user@host# set vr1 interface reth0.0
   user@host# set vr1 interface reth1.0
   user@host# set vr1 interface st0.0
   user@host# set vr1 routing-options static route 192.168.168.1/24 next-hop st0.0
   ```

4. Configure the loopback interface as an external interface for the IKE gateway.

   ```
   [edit security ike]
   user@host# set policy ike-policy1 mode main
   user@host# set policy ike-policy1 proposal-set standard
   user@host# set policy ike-policy1 pre-shared-key ascii-text "$ABC123"
   user@host# set gateway t-ike-gate ike-policy ike-policy1
   user@host# set gateway t-ike-gate address 10.2.2.2
   user@host# set gateway t-ike-gate external-interface lo0.1
   ```


Copyright © 2016, Juniper Networks, Inc.

5. Configure an IPsec proposal.

```
[edit security ipsec]
user@host# set proposal p2-std-p1 authentication-algorithm hmac-sha1-96
user@host# set proposal p2-std-p1 encryption-algorithm 3des-cbc
user@host# set proposal p2-std-p1 lifetime-seconds 180
user@host# set proposal p2-std-p2 authentication-algorithm hmac-sha1-96
user@host# set proposal p2-std-p2 encryption-algorithm aes-128-cbc
user@host# set proposal p2-std-p2 lifetime-seconds 180
user@host# set policy vpn-policy1 perfect-forward-secrecy keys group2
user@host# set policy vpn-policy1 proposals p2-std-p1
user@host# set policy vpn-policy1 proposals p2-std-p2
user@host# set vpn t-ike-vpn bind-interface st0.0
user@host# set vpn t-ike-vpn ike gateway t-ike-gate
user@host# set vpn t-ike-vpn ike proxy-identity local 10.10.10.1/24
user@host# set vpn t-ike-vpn ike proxy-identity remote 192.168.168.1/24
user@host# set vpn t-ike-vpn ike ipsec-policy vpn-policy1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lo0**, **show routing-instances**, **show security ike**, and **show security ipsec** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces lo0
    unit 1 {
      family inet {
        address 10.3.3.3/30;
      }
    }
    redundant-pseudo-interface-options {
      redundancy-group 1;
    }
```

```
[edit]
user@host# show routing-instances
  vr1 {
    instance-type virtual-router;
    interface lo0.1;
    interface reth0.0;
    interface reth1.0;
    interface st0.0;
    routing-options {
      static {
        route 192.168.168.1/24 next-hop st0.0;
      }
    }
  }
```

```
[edit]
user@host# show security ike
  policy ike-policy1 {
    mode main;
    proposal-set standard;
    pre-shared-key ascii-text "$ABC123";
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    gateway t-ike-gate {
      ike-policy ike-policy1;
      address 10.2.2.2;
      external-interface lo0.1;
    }
```

```
[edit]
user@host# show security ipsec
  proposal p2-std-p1 {
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 180;
  }
    proposal p2-std-p2 {
      authentication-algorithm hmac-sha1-96;
      encryption-algorithm aes-128-cbc;
      lifetime-seconds 180;
    }
    policy vpn-policy1 {
      perfect-forward-secrecy {
        keys group2;
      }
      proposals [ p2-std-p1 p2-std-p2 ];
    }
  policy vpn-policy2 {
    perfect-forward-secrecy {
      keys group2;
    }
    proposals [ p2-std-p1 p2-std-p2 ];
  }
    vpn t-ike-vpn {
      bind-interface st0.0;
      ike {
        gateway t-ike-gate;
        proxy-identity {
          local 10.10.10.1/24;
          remote 192.168.168.1/24;
        }
        ipsec-policy vpn-policy1;
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configuration

Purpose Verify that the configuration for redundancy groups for loopback interfaces is correct.

Action From operational mode, enter the **show chassis cluster interfaces** command.

```
user@host> show chassis cluster interfaces
Control link status: Up
  Control interfaces:
```


Copyright © 2016, Juniper Networks, Inc.

```
Index Interface Status
0 em0 Up
1 em1 Down
Fabric link status: Up
Fabric interfaces:
Name Child-interface Status
fab0 ge-0/0/7 Up / Up
fab0
fab1 ge-13/0/7 Up / Up
fab1
Redundant-ethernet Information:
Name Status Redundancy-group
reth0 Up 1
reth1 Up 1
reth2 Up 1
reth3 Down Not configured
reth4 Down Not configured
Redundant-pseudo-interface Information:
Name Status Redundancy-group
lo0 Up 1
```

Meaning The **show chassis cluster interfaces** command displays the chassis cluster interfaces information. If the status of the Redundant-pseudo-interface Information field shows the lo0 interface as Up and the status of the Redundant-ethernet Information field shows reth0, reth1, and reth2 fields as Up then your configuration is correct.

Related Documentation

- Understanding Loopback Interface for a High Availability VPN on page 6320

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 84

Configuring IPv6 IPsec VPNs

- Configuring IPv6 IPsec VPNs on page 6331

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 263

Configuring IPv6 IPsec VPNs

- VPN Feature Support for IPv6 Addresses on page 6331
- Understanding IPv6 IKE and IPsec Packet Processing on page 6335
- IPv6 IPsec Configuration Overview on page 6341
- Example: Configuring an IPv6 IPsec Manual VPN on page 6341

VPN Feature Support for IPv6 Addresses

A route-based site-to-site VPN tunnel with a point-to-point secure tunnel interface can operate in IPv4-in-IPv4, IPv6-in-IPv6, IPv6-in-IPv4, or IPv4-in-IPv6 tunnel modes. IPv6 addresses can be in the outer IP header, which represents the tunnel endpoint, or in the inner IP header, which represents the final source and destination addresses for a packet.

Table 545 defines the support for IPv6 addresses in VPN features.

Table 545: IPv6 Address Support in VPN Features

Feature	Supported	Not Supported	Exceptions
IKE and IPsec Support:			
IKEv1 and IKEv2	X		Unless specified, all supported features are applicable for IKEv1 and IKEv2.
Route-based VPN	X		–
Policy-based VPN	X	X	IPv6 policy-based VPNs are not supported on high-end SRX Series devices or on branch SRX Series devices in chassis cluster configurations. IPv6 policy-based VPNs are only supported with IPv6-in-IPv6 tunnels on standalone branch SRX Series devices.
Site-to-site VPN	X		Only one-to-one, site-to-site VPN is supported. Many-to-one, site-to-site VPN (NHTB) is not supported. NHTB configuration cannot be committed for tunnel modes other than IPv4-in-IPv4 tunnels.

Copyright © 2016, Juniper Networks, Inc.

Table 545: IPv6 Address Support in VPN Features *(continued)*

Feature	Supported	Not Supported	Exceptions
Dynamic endpoint VPN		X	IPv6 dynamic endpoint VPNs are blocked during negotiation.
Dialup VPN		X	IPv6 dialup VPNs are blocked during negotiation.
AutoVPN		X	–
Group VPN		X	–
Point-to-point tunnel interfaces	X		–
Point-to-multipoint tunnel interfaces		X	–
Hub-and-spoke scenario for site-to-site VPNs	X		–
Numbered and unnumbered tunnel interfaces	X		–
Unicast static and dynamic (RIP, OSPF, BGP) routing	X		–
Multicast dynamic routing (PIM)		X	–
Virtual router	X		–
Logical system		X	–
Automatic and manual SA and key management	X		–
Multiple SPUs	X		–
Chassis cluster	X		IPsec VPN with active-active mode is supported only on branch SRX Series devices for route-based IPv6 tunnels. IPsec VPN with active-active mode is not supported on high-end SRX Series devices.
Statistics, logs, per-tunnel debugging	X		–
SNMP MIB	X		–
Local address selection	X		When multiple addresses in the same address family are configured on a physical external interface to a VPN peer, we recommend that you also configure **local-address** at the [**edit security ike gateway *gateway-name***] hierarchy level.
Loopback address termination	X		–

Copyright © 2016, Juniper Networks, Inc.

Table 545: IPv6 Address Support in VPN Features *(continued)*

Feature	Supported	Not Supported	Exceptions
Xauth or modecfg over IPv6		X	–
SPC insert	X		–
ISSU	X		–
DNS name as IKE gateway address	X		As with IPv4 tunnels, peer gateway address changes in the DNS name are not supported with IPv6 tunnels.
Preshared key or certificate authentication	X		–
NAT-Traversal (NAT-T) for IPv4 IKE peers	X		NAT-T is supported only for IPv6-in-IPv4 and IPv4-in-IPv4 tunnel modes with IKEv1. IPv6-in-IPv6 and IPv4-in-IPv6 tunnel modes are not supported. IKEv2 is not supported for NAT-T. NAT-T from IPv6 to IPv4 or from IPv4 to IPv6 is not supported.
Dead peer detection (DPD) and DPD gateway failover	X		DPD gateway failover is only supported for different gateway addresses within the same family. Failover from an IPv6 gateway address to an IPv4 gateway address, or vice versa, is not supported.
Encryption sets, authentication algorithms, and DH groups supported in Junos OS Release 12.1X45-D10 release for SRX Series devices.	X		–
Generic proposals and policies for IPv6 and IPv4	X		–
General IKE ID	X		–
ESP and AH transport modes		X	These modes are not supported for IPv4.
ESP and AH tunnel modes	X		AH tunnel mode with mutable extension headers and options is not supported.
Extended sequence number		X	–
Single proxy ID pairs	X		–
Multiple traffic selector pairs	X		Supported with IKEv1 only.
Lifetime of IKE or IPsec SA, in seconds	X		–
Lifetime of IKE SA, in kilobytes	X		–

Copyright © 2016, Juniper Networks, Inc.

Table 545: IPv6 Address Support in VPN Features *(continued)*

Feature	Supported	Not Supported	Exceptions
VPN monitoring		X	Configuration with IPv6 tunnels cannot be committed.
DF bit	X		For IPv6-in-IPv6 tunnels, the DF bit is set only if configured at the [**edit security ipsec vpn** ***vpn-name***] hierarchy level. **df-bit clear** is the default.
Dual-stack over tunnel external interface	X		–
IPv6 extension headers	X		IPv6 extension headers and IPv4 options for IKE and IPsec packets are accepted but are not processed. AH with mutable EHs and options is not supported.
Fragmentation and reassembly	X		–
VPN session affinity	X		–
Multicast traffic		X	–
Tunnel IP services (Screen, NAT, ALG, IPS, AppSecure)	X		–
Packet reordering for IPv6 fragments over tunnel		X	–
PKI Support:			
PKI in virtual router	X		–
RSA signature authentication (512-, 1024-, 2048-, or 4096-bit key size)	X		–
DSA signature authentication (512-, 1024-, 2048-, or 4096-bit key size)	X		–
ECDSA signatures	X		–
Certificate chain authentication		X	–
Automatic or manual enrollment over IPv4	X		–
Automatic or manual revocation over IPv4	X		–
Automatic or manual enrollment over IPv6		X	–
Automatic or manual revocation over IPv6		X	–

Copyright © 2016, Juniper Networks, Inc.

Table 545: IPv6 Address Support in VPN Features *(continued)*

Feature	Supported	Not Supported	Exceptions
IPv6 addresses within PKI certificate fields		X	–

Related Documentation

- Understanding VPN Tunnel Modes on page 6181
- IPsec VPN Overview on page 6041

Understanding IPv6 IKE and IPsec Packet Processing

This topic includes the following sections:

- IPv6 IKE Packet Processing on page 6335
- IPv6 IPsec Packet Processing on page 6336

IPv6 IKE Packet Processing

Internet Key Exchange (IKE) is part of the IPsec suite of protocols. It automatically enables two tunnel endpoints to set up security associations (SAs) and negotiate secret keys with each other. There is no need to manually configure the security parameters. IKE also provides authentication for communicating peers.

IKE packet processing in IPv6 networks involves the following elements:

- Internet Security Association and Key Management Protocol (ISAKMP) Identification Payload

  ISAKMP identification payload is used to identify and authenticate the communicating IPv6 peers. Two ID types (ID_IPV6_ADDR and ID_IPV6_ADDR_SUBNET) are enabled for IPv6. The ID type indicates the type of identification to be used. The ID_IPV6_ADDR type specifies a single 16-octet IPv6 address. This ID type represents an IPv6 address. The ID_IPV6_ADDR_SUBNET type specifies a range of IPv6 addresses represented by two 16-octet values. This ID type represents an IPv6 network mask. Table 546 lists the ID types and their assigned values in the identification payload.

Table 546: ISAKMP ID Types and Their Values

ID Type	Value
RESERVED	0
ID_IPV4_ADDR	1
ID_FQDN	2
ID_USER_FQDN	3
ID_IPV4_ADDR_SUBNET	4

Copyright © 2016, Juniper Networks, Inc.

Table 546: ISAKMP ID Types and Their Values *(continued)*

ID Type	Value
ID_IPV6_ADDR	5
ID_IPV6_ADDR_SUBNET	6
ID_IPV4_ADDR_RANGE	7
ID_IPV6_ADDR_RANGE	8
ID_DER_ASN1_DN	9
ID_DER_ASN1_GN	10
ID_KEY_ID	11
ID_LIST	12

The ID_IPV6_ADDR_RANGE type specifies a range of IPv6 addresses represented by two 16-octet values. The first octet value represents the starting IPv6 address and the second octet value represents the ending IPv6 address in the range. All IPv6 addresses falling between the first and last IPv6 addresses are considered to be part of the list.

NOTE: **Two ID types in ISAKMP identification payload (ID_IPV6_ADDR_RANGE and ID_IPV4_ADDR_RANGE) are not supported in this release.**

- Proxy ID

  A proxy ID is used during Phase 2 of IKE negotiation. It is generated before an IPsec tunnel is established. A proxy ID identifies the SA to be used for the VPN. Two proxy IDs are generated—local and remote. The local proxy ID refers to the local IPv4 or IPv6 address/network and subnet mask. The remote proxy ID refers to the remote IPv4 or IPv6 address/network and subnet mask.

- Security Association

  An SA is an agreement between VPN participants to support secure communication. SAs are differentiated based on three parameters—security parameter index (SPI), destination IPv6 address, and security protocol (either AH or ESP). The SPI is a unique value assigned to an SA to help identify an SA among multiple SAs. In an IPv6 packet, the SA is identified from the destination address in the outer IPv6 header and the security protocol is identified from either the AH or the ESP header.

IPv6 IPsec Packet Processing

After IKE negotiations are completed and the two IKE gateways have established Phase 1 and Phase 2 SAs, IPv6 IPsec employs authentication and encryption technologies to

Copyright © 2016, Juniper Networks, Inc.

secure the IPv6 packets. Because IPv6 addresses are 128 bits long compared to IPv4 addresses, which are 32-bits long, IPv6 IPsec packet processing requires more resources.

NOTE: Packet reordering for IPv6 fragments over a tunnel is not supported.

Devices with IPv6 addressing do not perform fragmentation. IPv6 hosts should either perform path MTU discovery or send packets smaller than the IPv6 minimum MTU size of 1280 bytes.

This topic includes the following sections:

- AH Protocol in IPv6 on page 6337
- ESP Protocol in IPv6 on page 6337
- IPv4 Options and IPv6 Extension Headers with AH and ESP on page 6338
- Integrity Check Value Calculation in IPv6 on page 6338
- Header Construction in Tunnel Modes on page 6339

AH Protocol in IPv6

The AH protocol provides data integrity and data authentication for IPv6 packets. IPv6 IPsec uses extension headers (for example, hop-by-hop and routing options) that must be arranged in a particular way in the IPv6 datagram. In AH tunnel mode, the AH header immediately follows the new outer IPv6 header similar to that in IPv4 AH tunnel mode. The extension headers are placed after the original inner header. Therefore, in AH tunnel mode, the entire packet is encapsulated by adding a new outer IPv6 header, followed by an authentication header, an inner header, extension headers, and the rest of the original datagram as shown in Figure 269.

Figure 269: IPv6 AH Tunnel Mode

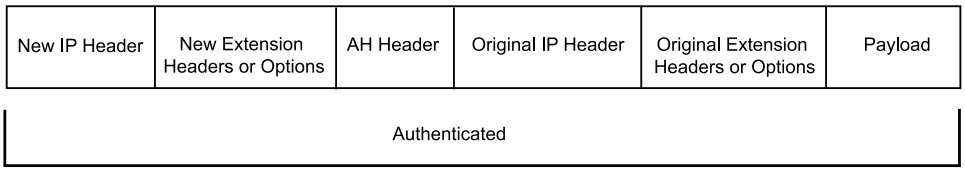

Unlike ESP, the AH authentication algorithm covers the outer header as well as any new extension headers and options.

NOTE: AH tunnel mode on SRX Series devices does not support IPv4 mutable options or IPv6 mutable extension headers. See Table 547.

ESP Protocol in IPv6

ESP protocol provides both encryption and authentication for IPv6 packets. Because IPv6 IPsec uses extension headers (for example, hop-by-hop and routing options) in the IPv6 datagram, the most important difference between IPv6 ESP tunnel mode and IPv4 ESP tunnel mode is the placement of extension headers in the packet layout. In ESP

Copyright © 2016, Juniper Networks, Inc.

tunnel mode, the ESP header immediately follows the new outer IPv6 header similar to that in IPv4 ESP tunnel mode. Therefore, in ESP tunnel mode, the entire packet is encapsulated by adding a new outer IPv6 header, followed by an ESP header, an inner header, extension headers, and the rest of the original datagram as shown in Figure 270.

Figure 270: IPv6 ESP Tunnel Mode

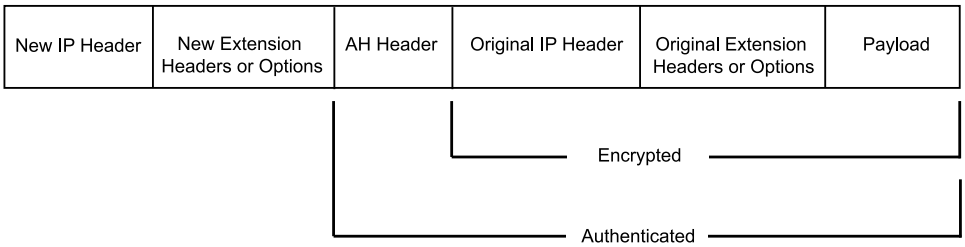

IPv4 Options and IPv6 Extension Headers with AH and ESP

IPsec packets with IPv4 options or IPv6 extension headers can be received for decapsulation on SRX Series devices. Table 547 shows the IPv4 options or IPv6 extension headers that are supported with the ESP or AH protocol on SRX Series devices. If an unsupported IPsec packet is received, ICV calculation fails and the packet is dropped.

Table 547: Support for IPv4 Options or IPv6 Extension Headers

Options or Extension Headers	Branch SRX Series Devices	High-End SRX Series Devices
ESP with IPv4 options	Supported	Supported
ESP with IPv6 extension headers	Supported	Supported
AH with IPv4 immutable options	Supported	Supported
AH with IPv6 immutable extension headers	Supported	Supported
AH with IPv4 mutable options	Not supported	Not supported
AH with IPv6 mutable extension headers	Not supported	Not supported

Integrity Check Value Calculation in IPv6

The AH protocol verifies the integrity of the IPv6 packet by computing an Integrity Check Value (ICV) on the packet contents. ICV is usually built over an authentication algorithm such as MD5 or SHA-1. The IPv6 ICV calculations differ from that in IPv4 in terms of two header fields—mutable header and optional extension header.

You can calculate the AH ICV over the IPv6 header fields that are either immutable in transit or predictable in value upon arrival at the tunnel endpoints. You can also calculate the AH ICV over the AH header and the upper level protocol data (considered to be immutable in transit). You can calculate the ESP ICV over the entire IPv6 packet, excluding the new outer IPv6 header and the optional extension headers.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Unlike IPv4, IPv6 has a method for tagging options as mutable in transit. IPv6 optional extension headers contain a flag that indicates mutability. This flag determines the appropriate processing.

IPv4 mutable options and IPv6 extension headers are not supported with the AH protocol.

Header Construction in Tunnel Modes

In tunnel mode, the source and destination addresses of the outer IPv4 or IPv6 header represent the tunnel endpoints, while the source and destination addresses of the inner IPv4 or IPv6 header represent the final source and destination addresses. Table 548 summarizes how the outer IPv6 header relates to the inner IPv6 or IPv4 header for IPv6-in-IPv6 or IPv4-in-IPv6 tunnel modes. In outer header fields, "Constructed" means that the value of the outer header field is constructed independently of the value in the inner header field.

Table 548: IPv6 Header Construction for IPv6-in-IPv6 and IPv4-in-IPv6 Tunnel Modes

Header Fields	Outer Header at Encapsulator	Inner Header at Decapsulator
version	6.	No change.
DS field	Copied from the inner header.	No change.
ECN field	Copied from the inner header.	Constructed.
flow label	0.	No change.
payload length	Constructed.	No change.
next header	AH, ESP, and routing header.	No change.
hop limit	64.	Decrement.
src address	Constructed.	No change.
dest address	Constructed.	No change.
Extension headers	Never copied.	No change.

Table 549 summarizes how the outer IPv4 header relates to the inner IPv6 or IPv4 header for IPv6-in-IPv4 or IPv4-in-IPv4 tunnel modes. In outer header fields, "Constructed" means that the value of the outer header field is constructed independently of the value in the inner header field.

Copyright © 2016, Juniper Networks, Inc.

Table 549: IPv4 Header Construction for IPv6-in-IPv4 and IPv4-in-IPv4 Tunnel Modes

Header Fields	Outer Header	Inner Header
version	4.	No change.
header length	Constructed.	No change.
DS field	Copied from the inner header.	No change.
ECN field	Copied from the inner header.	Constructed.
total length	Constructed.	No change.
ID	Constructed.	No change.
flags (DF, MF)	Constructed.	No change.
fragment offset	Constructed.	No change.
TTL	64.	Decrement.
protocol	AH, ESP	No change.
checksum	Constructed.	Constructed.
src address	Constructed.	No change.
dest address	Constructed.	No change.
options	Never copied.	No change.

For IPv6-in-IPv4 tunnel mode, the Don't Fragment (DF) bit is cleared by default. If the **df-bit set** or **df-bit copy** options are configured at the [**edit security ipsec vpn** *vpn-name*] hierarchy level for the corresponding IPv4 VPN, the DF bit is set in the outer IPv4 header.

For IPv4-in-IPv4 tunnel mode, the DF bit in the outer IPv4 header is based on the **df-bit** option configured for the inner IPv4 header. If **df-bit** is not configured for the inner IPv4 header, the DF bit is cleared in the outer IPv4 header.

Related Documentation

- IPsec VPN Overview on page 6041
- IPv6 IPsec Configuration Overview on page 6341
- Example: Configuring an IPv6 IPsec Manual VPN on page 6341

Copyright © 2016, Juniper Networks, Inc.

IPv6 IPsec Configuration Overview

Juniper Networks supports manual and autokey IKE with preshared keys configurations for IPv6 IPsec VPN.

- Manual VPN—In a manual VPN configuration, the secret keys and security associations (SAs) are manually configured on the tunnel endpoints using the manual key mechanism. To create an IPv6 IPsec manual VPN, see "Example: Configuring an IPv6 IPsec Manual VPN" on page 6341.
- AutoKey IKE VPN—In an autoKey IKE VPN configuration, the secret keys and SAs are automatically created using the autoKey IKE mechanism. To set up an IPv6 autoKey IKE VPN, two phases of negotiations are required—Phase 1 and Phase 2.
  - Phase 1—In this phase, the participants establish a secure channel for negotiating the IPsec SAs. For more information on Phase 1 negotiations, see "Understanding Phase 1 of IKE Tunnel Negotiation" on page 6055.
  - Phase 2—In this phase, the participants negotiate the IPsec SAs for authenticating and encrypting the IPv6 data packets. For more information on Phase 2 negotiations, see "Understanding Phase 2 of IKE Tunnel Negotiation" on page 6057.

Related Documentation

- Understanding IPv6 IKE and IPsec Packet Processing on page 6335
- Example: Configuring an IPv6 IPsec Manual VPN on page 6341
- *Example: Configuring an IPv6 AutoKey IKE Policy-Based VPN*

Example: Configuring an IPv6 IPsec Manual VPN

This example shows how to configure an IPv6 IPsec manual VPN.

- Requirements on page 6341
- Overview on page 6341
- Configuration on page 6342
- Verification on page 6343

Requirements

Before you begin:

- Understand how VPNs work. See "IPsec VPN Overview" on page 6041.
- Understand IPv6 IPsec packet processing. See "Understanding IPv6 IKE and IPsec Packet Processing" on page 6335.

Overview

In a Manual VPN configuration, the secret keys are manually configured on the two IPsec endpoints.

Copyright © 2016, Juniper Networks, Inc.

In this example, you:

- Configure the authentication parameters for a VPN named vpn-sunnyvale.
- Configure the encryption parameters for vpn-sunnyvale.
- Specify the outgoing interface for the SA.
- Specify the IPv6 address of the peer.
- Define the IPsec protocol. Select the ESP protocol because the configuration includes both authentication and encryption.
- Configure a security parameter index (SPI).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ipsec vpn vpn-sunnyvale manual authentication algorithm hmac-md5–96
  key ascii-text “$ABC123”
set security ipsec vpn vpn-sunnyvale manual encryption algorithm 3des-cbc key ascii-text
  “$ABC123”
set security ipsec vpn vpn-sunnyvale manual external-interface ge-0/0/14.0
set security ipsec vpn vpn-sunnyvale manual gateway 2001:db8:1212::1112
set security ipsec vpn vpn-sunnyvale manual protocol esp
set security ipsec vpn vpn-sunnyvale manual spi 12435
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure security algorithms:

1. Configure the authentication parameters.

   ```
   [edit security ipsec vpn vpn-sunnyvale manual]
   user@host# set authentication algorithm hmac-md5–96 key ascii-text “$ABC123”
   ```

2. Configure the encryption parameters.

   ```
   [edit security ipsec vpn vpn-sunnyvale manual]
   user@host# set encryption algorithm 3des-cbc key ascii-text “$ABC123”
   ```

3. Specify the outgoing interface for the SA.

   ```
   [edit security ipsec vpn vpn-sunnyvale manual]
   user@host# set external-interface ge-0/0/14.0
   ```

4. Specify the IPv6 address of the peer.

   ```
   [edit security ipsec vpn vpn-sunnyvale manual]
   user@host# set gateway 2001:db8:1212::1112
   ```

5. Define the IPsec protocol.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ipsec vpn vpn-sunnyvale manual]
user@host# set protocol esp
```

6. Configure an SPI.

```
[edit security ipsec vpn vpn-sunnyvale manual]
user@host# set spi 12435
```

Results From configuration mode, confirm your configuration by entering the **show security ipsec vpn vpn-sunnyvale** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
[user@host]show security ipsec vpn vpn-sunnyvale
manual {
gateway 2001:db8:1212::1112 ;
external-interface ge-0/0/14.0 ;
protocol esp ;
spi 12435 ;
authentication {
  algorithm hmac-md5-96 ;
  key ascii-text $ABC123” ;## SECRET DATA
}
  encryption {
    algorithm 3des-cbc ;
    key ascii-text $ABC123”; ## SECRET DATA
    }
  }
```

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying Security Algorithms on page 6343

Verifying Security Algorithms

Purpose Determine if security algorithms are applied or not.

Action From operational mode, enter the **show security ipsec security-associations** command.

Related Documentation

- Understanding IPv6 IKE and IPsec Packet Processing on page 6335
- IPv6 IPsec Configuration Overview on page 6341

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 85

Configuring Public Key Infrastructure

- Managing Digital Certificates with PKI on page 6347
- Configuring Digital Certificate Validation on page 6357
- Generating a Public-Private Key Pair on page 6367
- Configuring Certificate Authority Profiles on page 6369
- Configuring CA and Local Certificates on page 6371
- Managing Certificate Revocation on page 6409
- Generating Self-Signed Certificates on page 6433
- Configuring a Device for Certificate Chains on page 6437

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 264

Managing Digital Certificates with PKI

- Understanding Certificates and PKI on page 6347
- Cryptographic Key Handling Overview on page 6351
- Understanding CMPv2 and SCEP Certificate Enrollment on page 6352
- Understanding Certificate Enrollment with CMPv2 on page 6353
- Digital Certificates Configuration Overview on page 6355

Understanding Certificates and PKI

A digital certificate is an electronic means for verifying your identity through a trusted third party, known as a certificate authority (CA). Alternatively, you can use a self-signed certificate to attest to your identity.

The CA server you use can be owned and operated by an independent CA or by your own organization, in which case you become your own CA. If you use an independent CA, you must contact them for the addresses of their CA and certificate revocation list (CRL) servers (for obtaining certificates and CRLs) and for the information they require when submitting personal certificate requests. When you are your own CA, you determine this information yourself.

The Public Key Infrastructure (PKI) provides an infrastructure for digital certificate management.

This topic includes the following sections:

- Certificate Signatures and Verification on page 6347
- Public Key Infrastructure on page 6348
- PKI Management and Implementation on page 6350
- Internet Key Exchange on page 6351

Certificate Signatures and Verification

The CA that issues a certificate uses a hash algorithm to generate a digest, and then “signs” the certificate by encrypting the digest with its private key. The result is a digital signature. The CA then makes the digitally signed certificate available for download to the person who requested it. Figure 271 illustrates this process.

Copyright © 2016, Juniper Networks, Inc.

The recipient of the certificate generates another digest by applying the same hash algorithm to the certificate file, then uses the CA's public key to decrypt the digital signature. By comparing the decrypted digest with the digest just generated, the recipient can confirm the integrity of the CA's signature and, by extension, the integrity of the accompanying certificate. Figure 271 illustrates this process.

NOTE: A certificate is considered valid if the digital signature can be verified and the serial number of the certificate is not listed in a certificate revocation list.

Figure 271: Digital Signature Verification

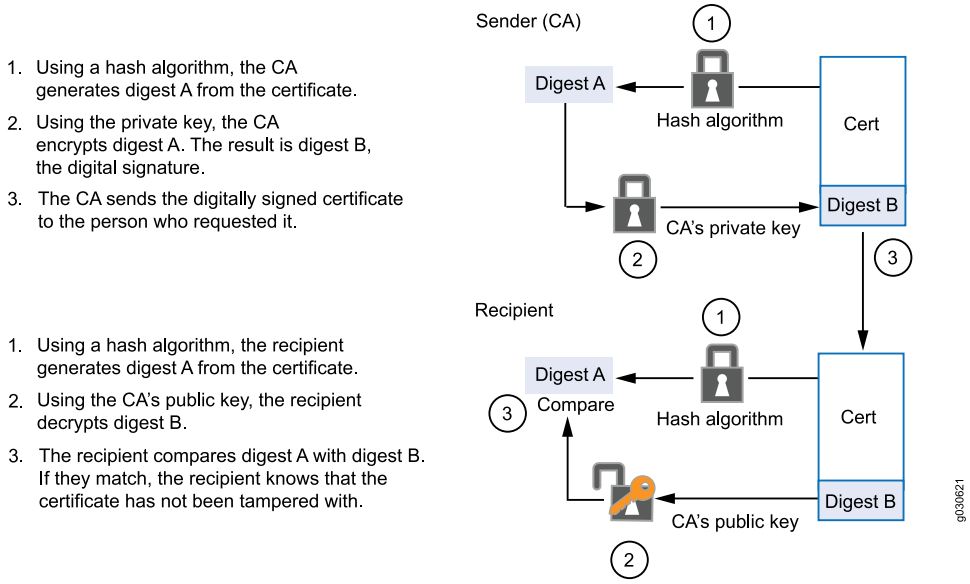

When Digital Signature Algorithm (DSA) signatures are used, the SHA-1 hash algorithm is used to generate the digest. When Rivest-Shamir-Adleman (RSA) signatures are used, SHA-1 is the default hash algorithm used to generate the digest; you can specify the SHA-256 hash algorithm with the **digest** option of the **request security pki generate-certificate-request** or **request security pki local-certificate generate-self-signed** commands. When Elliptic Curve Digital Signature Algorithm (ECDSA) signatures are used, the SHA-256 hash algorithm is used for ECDSA-256 signatures and the SHA-384 hash algorithm is used for ECDSA-384 signatures.

Public Key Infrastructure

To verify the trustworthiness of a certificate, you must be able to track a path of certified certificate authorities (CAs) from the one issuing your local certificate to the root authority of a CA domain. Public key infrastructure (PKI) refers to the hierarchical structure of trust required for the successful implementation of public key cryptography.

Figure 272 shows the structure of a single-domain certificate authority with multiple hierarchy levels.

Copyright © 2016, Juniper Networks, Inc.

Figure 272: PKI Hierarchy of Trust—CA Domain

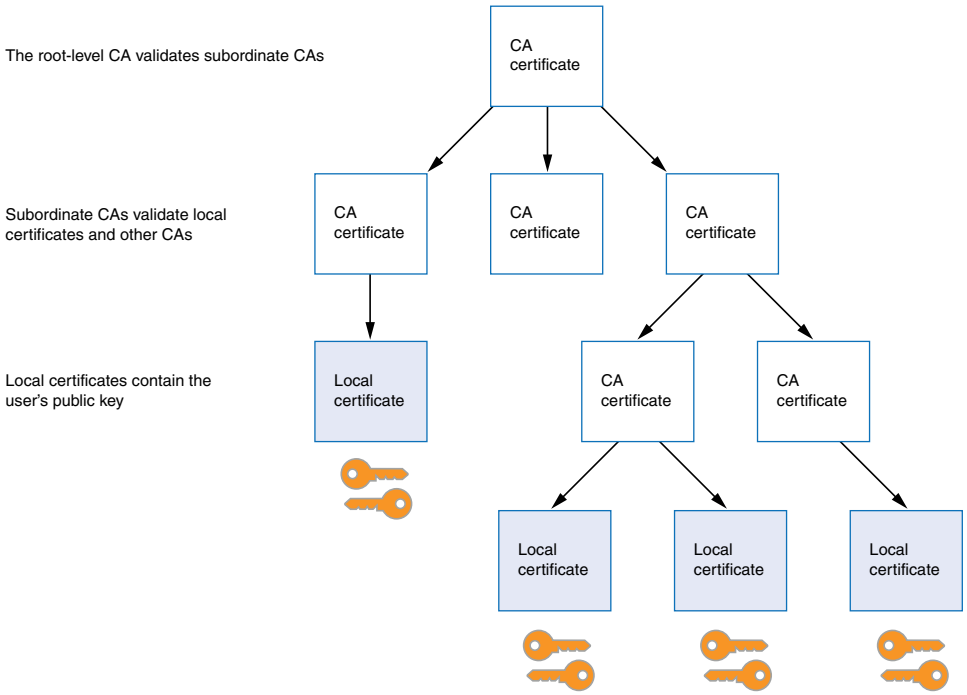

If certificates are used solely within an organization, that organization can have its own CA domain within which a company CA issues and validates certificates for its employees. If that organization later wants its employees to exchange their certificates with certificates from another CA domain (for example, with employees at another organization that has its own CA domain), the two CAs can develop cross-certification by agreeing to trust the authority of each other. In this case, the PKI structure does not extend vertically but does extend horizontally. See Figure 273.

Copyright © 2016, Juniper Networks, Inc.

Figure 273: Cross-Certification

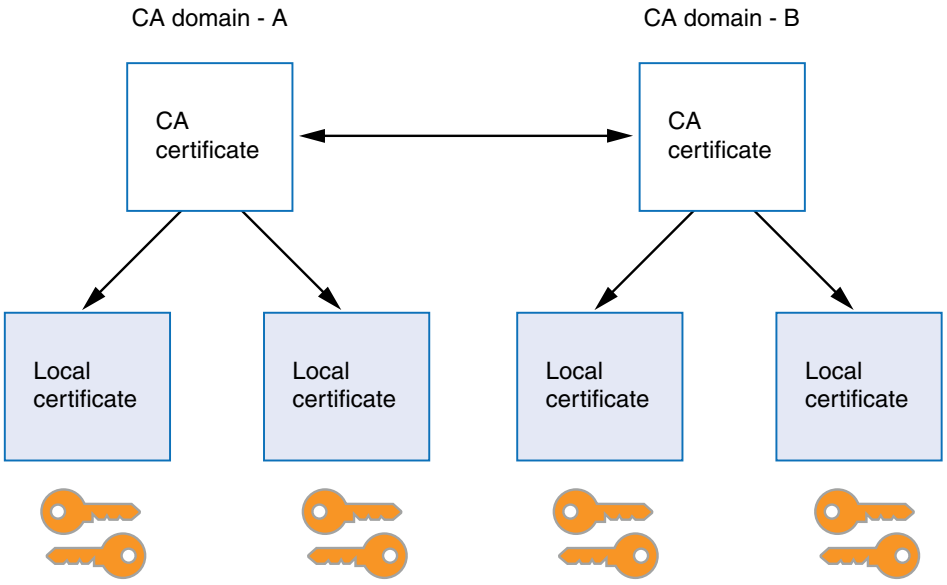

Users in the CA domain A can use their certificates and key pairs with users in CA domain B because the CA's have cross-certified each other.

PKI Management and Implementation

The minimum PKI elements required for certificate-based authentication in Junos OS are:

- CA certificates and authority configuration.
- Local certificates including the device's identity (example: IKE ID type and value) and private and public keys
- Certificate validation through a CRL.

Junos OS supports three different types of PKI objects:

- Private/public key pair
- Certificates
  - Local certificate—The local certificate contains the public key and identity information for the Juniper Networks device. The Juniper Networks device owns the associated private key. This certificate is generated based on a certificate request from the Juniper Networks device.
  - Pending certificate — A pending certificate contains a key pair and identity information that is generated into a PKCS10 certificate request and manually sent to a certificate authority (CA). While the Juniper Networks device waits for the certificate from the CA, the existing object (key pair and the certificate request) is tagged as a certificate request or pending certificate.

Copyright © 2016, Juniper Networks, Inc.

  - CA certificate — When the certificate is issued by the CA and loaded into the Junos device, the pending certificate is replaced by the newly generated local certificate. All other certificates loaded into the device are considered CA certificates.

- Certificate revocation lists (CRLs)

For convenience and practicality, PKI must be transparently managed and implemented. Toward this goal, Junos OS supports the following features:

- Generates a public-private key pair.
- Loads multiple local certificates from different CAs.
- Uses either Simple Certificate Enrollment Protocol (SCEP) or CMPv2 for certificate enrollment.
- Delivers a certificate when establishing an IPsec tunnel.
- Validates a certificate path upward through a single level of CA authorities.
- Supports the Public-Key Cryptography Standards #7 (PKCS #7) cryptographic . As a result, the device can accept X.509 certificates and certificate revocation lists (CRLs) packaged within a PKCS #7 envelope.

  NOTE: **Junos OS supports a PKCS #7 file size of up to 7 KB.**

- Retrieves CRLs online retrieval through Lightweight Directory Access Protocol (LDAP) or Hypertext Transfer Protocol (HTTP).

Internet Key Exchange

The procedure for digitally signing messages sent between two participants in an Internet Key Exchange (IKE) session is similar to digital certificate verification, with the following differences:

- Instead of making a digest from the CA certificate, the sender makes it from the data in the IP packet payload.
- Instead of using the CA's public-private key pair, the participants use the sender's public-private key pair.

Related Documentation

- Digital Certificates Configuration Overview on page 6355
- Understanding Certificate Chains on page 6437
- IPsec VPN Overview on page 6041

Cryptographic Key Handling Overview

With cryptographic key handling, persistent keys are stored in the memory of the device without any attempt to alter them. While the internal memory device is not directly accessible to a potential adversary, those who require a second layer of defence, may

Copyright © 2016, Juniper Networks, Inc.

enable special handling for cryptographic keys. When enabled, the cryptographic key handling encrypts keys when not immediately in use, performs error detection when copying a key from one memory location to another, and overwrites the memory location of a key with a random bit pattern when the key is no longer in use. Keys are also protected when they are stored in the flash memory of the device. Enabling cryptographic key handling feature does not cause any externally observable change in the behavior of the device, and the device continues to interoperate with the other devices.

NOTE: A cryptographic administrator can enable and disable the cryptographic self-test functions, however the security administrator can modify the behavior of the cryptographic self test functions like configuring periodic self-test or selecting a subset of cryptographic self-tests.

The following persistent keys are currently under the management of IKE and PKI:

- IKE preshared keys (IKE PSKs)
- PKI private keys
- Manual VPN keys

Related Documentation

- Understanding Certificates and PKI on page 6347

Understanding CMPv2 and SCEP Certificate Enrollment

Based on your deployment environment, you can use either Certificate Management Protocol version 2 (CMPv2) or Simple Certificate Enrollment Protocol (SCEP) for online certificate enrollment. This topic describes some of the basic differences between the two protocols.

Table 550 describes the differences between the CMPv2 and SCEP certificate enrollment protocols.

Table 550: Comparison of CMPv2 and SCEP Certificate Enrollment

Attribute	CMPv2	SCEP
Supported certificate types:	DSA, ECDSA, and RSA	RSA only
Supported standards	RFCs 4210 and 4211	Internet Engineering Task Force draft

Certificate enrollment and reenrollment requests and responses differ between CMPv2 and SCEP. With CMPv2, there is no separate command to enroll CA certificates. With SCEP, you enroll CA certificates with the **request security pki ca-certificate enroll** command and specify the CA profile. A CA profile must be configured with either CMPv2 or SCEP.

Related Documentation

- Understanding Certificate Enrollment with CMPv2 on page 6353
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

- Enrolling Digital Certificates Online: Configuration Overview on page 6355

Understanding Certificate Enrollment with CMPv2

The **request security pki local-certificate enroll cmpv2** command uses CMPv2 to enroll a local digital certificate online. This command loads both end-entity and CA certificates based on the CA server configuration. The CA profile must be created prior to CA certificate enrollment because the enrollment URL is extracted from the CA profile.

This topic describes certificate enrollment with the CMPv2 protocol.

- Certificate Enrollment and Reenrollment Messages on page 6353
- End-Entity Certificate with Issuer CA Certificate on page 6353
- End-Entity Certificate with CA Certificate Chain on page 6354

Certificate Enrollment and Reenrollment Messages

The CMPv2 protocol mainly involves certificate enrollment and reenrollment operations. The certificate enrollment process includes Initialization Request and Initialization Response messages, while certificate reenrollment includes Key Update Request and Key Update Response messages.

NOTE: If the Initialization Response message needs to be authenticated by a CA certificate, the CA certificate must be enrolled prior to any end-entity certificate enrollment.

The Initialization Response or Key Update Response message can contain an issuer CA certificate or a chain of CA certificates. The CA certificates received in the responses are treated as trusted CA certificates and stored in the receiving device if they are not already present in the trusted CA store. These CA certificates are later used for end-entity certificate validation.

NOTE: CA certificate reenrollment is not supported.

A CA might issue a new CA certificate prior to the expiration of the current CA certificate. If a new CA certificate arrives during certificate reenrollment with a new public key, the new CA certificate is not saved in the device.

End-Entity Certificate with Issuer CA Certificate

In a simple scenario, the Initialization Response message might contain only an end-entity certificate, in which case the CA information is provided separately. The certificate is stored in the end-entity certificate store.

The Initialization Response message can contain an end-entity certificate as well as a self-signed issuer CA certificate. The end-entity certificate is first stored in the certificate store, and then the CA certificate is checked. If the CA certificate is found and the subject

Copyright © 2016, Juniper Networks, Inc.

distinguished name (DN) of the CA certificate in the Initialization Response message matches the issuer DN of the end-entity certificate, the CA certificate is stored in the CA certificate store for the CA profile name specified in the CMPv2 certificate enrollment command. If the CA certificate already exists in the CA certificate store, no action is taken.

End-Entity Certificate with CA Certificate Chain

In many deployments, the end-entity certificate is issued by an intermediate CA in a certificate chain. In this case, the Initialization Response message can contain the end-entity certificate along with a list of CA certificates in the chain. The intermediate CA certificates and the self-signed root CA certificates are all required to validate the end-entity certificate. The CA chain might also be needed to validate certificates received from peer devices with similar hierarchies. The following section describes how certificates in the CA chain are stored.

In Figure 274, the Initialization Response message includes the end-entity certificate and three CA certificates in a certificate chain.

Figure 274: End-Entity Certificate with CA Certificate Chain

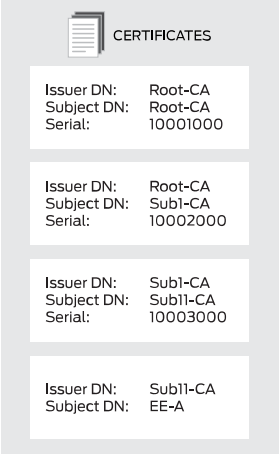

The end-entity certificate is stored in the end-entity certificate store. Each CA certificate needs a CA profile. The CA certificate with the subject DN Sub11-CA is the first CA in the chain and is the issuer of the end-entity certificate. It is stored in the CA profile that is specified with the CMPv2 certificate enrollment command.

Each of the remaining CA certificates in the chain is checked for its presence in the CA store. If a CA certificate is not present in the CA store, it is saved and a CA profile is created for it. The new CA profile name is created using the least significant 16 digits of the CA certificate serial number. If the serial number is longer than 16 digits, the most significant digits beyond 16 digits are truncated. If the serial number is shorter than 16 digits, the remaining most significant digits are filled with **0**s. For example, if the serial number is 11111000100010001000, then the CA profile name is **1000100010001000**. If the serial number is 10001000, then the CA profile name is **0000000010001000**.

It is possible that multiple certificate serial numbers can have the same least significant 16 digits. In that case, **-00** is appended to the profile name to create a unique CA profile

Copyright © 2016, Juniper Networks, Inc.

name; additional CA profile names are created by incrementing the appended number, from **-01** up to **-99**. For example, CA profile names can be **1000100010001000**, **1000100010001000-00**, and **1000100010001000-01**.

Related Documentation

- Understanding Certificate Authority Profiles on page 6369
- Understanding Certificate Chains on page 6437
- Understanding CMPv2 and SCEP Certificate Enrollment on page 6352

Digital Certificates Configuration Overview

You can obtain CA and local certificates manually, or online using Simple Certificate Enrollment Protocol (SCEP) or CMPv2. Certificates are verifiable and renewable, and you can delete them when they are no longer needed.

Manual certificate processing includes generation of a PKCS10 request, submission to the CA, retrieval of the signed certificate, and manually loading of the certificate into the Juniper Networks device. Based on your deployment environment, you can use either SCEP or CMPv2 for online certificate enrollment.

To use a digital certificate to authenticate your identity when establishing a secure VPN connection, you must first do the following:

- Obtain a CA certificate from which you intend to obtain a local certificate, and then load the CA certificate onto the device. The CA certificate can contain a CRL to identify invalid certificates.
- Obtain a local certificate from the CA whose CA certificate you have previously loaded, and then load the local certificate in the device. The local certificate establishes the identity of the Juniper Networks device with each tunnel connection.

This topic includes the following sections:

- Enrolling Digital Certificates Online: Configuration Overview on page 6355
- Manually Generating Digital Certificates: Configuration Overview on page 6356

Enrolling Digital Certificates Online: Configuration Overview

You can use either Certificate Management Protocol version 2 (CMPv2) or Simple Certificate Enrollment Protocol (SCEP) to enroll digital certificates. To enroll a certificate online:

1. Generate a key pair on the device. See "Example: Generating a Public-Private Key Pair" on page 6368.
2. Create a CA profile or profiles containing information specific to a CA. See "Example: Configuring a CA Profile" on page 6369.
3. For SCEP only, enroll the CA certificate. See "Enrolling a CA Certificate Online Using SCEP" on page 6372.

Copyright © 2016, Juniper Networks, Inc.

4. Enroll the local certificate from the CA whose CA certificate you have previously loaded. See "Example: Enrolling a Local Certificate Online Using SCEP" on page 6373.
5. Configure automatic reenrollment. See "Example: Using SCEP to Automatically Renew a Local Certificate" on page 6375.

Manually Generating Digital Certificates: Configuration Overview

To obtain digital certificates manually:

1. Generate a key pair on the device. See "Example: Generating a Public-Private Key Pair" on page 6368.
2. Create a CA profile or profiles containing information specific to a CA. See "Example: Configuring a CA Profile" on page 6369.
3. Generate the CSR for the local certificate and send it to the CA server. See "Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server" on page 6376.
4. Load the certificate onto the device. See "Example: Loading CA and Local Certificates Manually" on page 6378.
5. Configure automatic reenrollment. See "Example: Using SCEP to Automatically Renew a Local Certificate" on page 6375.
6. If necessary, load the certificate's CRL on the device. See "Example: Manually Loading a CRL onto the Device" on page 6427.
7. If necessary, configure the CA profile with CRL locations. See "Example: Configuring a Certificate Authority Profile with CRL Locations" on page 6428

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Verifying Certificate Validity on page 6429
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428
- Deleting Certificates (CLI Procedure) on page 6379

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 265

Configuring Digital Certificate Validation

- Understanding Digital Certificate Validation on page 6357
- Example: Improving Digital Certificate Validation by Configuring Policy OIDs on an SRX Series Device on page 6362

Understanding Digital Certificate Validation

During IKE negotiation, the PKI daemon on an SRX Series device validates X509 certificates received from VPN peers. The certificate validation performed is specified in RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. Basic certificate and certificate chain validations include signature and date validation as well as revocation checks. This topic describes additional digital certificate validations performed by the PKI daemon.

- Policy Validation on page 6357
- Path Length Validation on page 6359
- Key Usage on page 6360
- Issuer and Subject Distinguished Name Validation on page 6361

Policy Validation

X509 certificates can include optional policy validation fields. If a policy validation field is present, policy validation is performed for the entire certificate chain including the end entity (EE) certificate and intermediate certificate authority (CA) certificates. Policy validation is not applicable to the root certificate. Policy validation ensures that the EE and intermediate CA certificates have a common policy. If no common policy exists for the certificate chain being validated, certificate validation fails.

Prior to policy validation, a certificate chain containing the self-signed root certificate, intermediate CA certificates, and EE certificate must be built. The policy validation starts with the intermediate CA certificate issued by the self-signed root certificate and continues through the EE certificate.

The following optional certificate fields are used for policy validation:

- **policy-oids**
- **requireExplicitPolicy**

Copyright © 2016, Juniper Networks, Inc.

- **skipCerts**

These fields are described in the following sections.

Policy OIDs Configured on SRX Series Devices

In some situations, it might be desirable to only accept certificates with known policy object identifiers (OIDs) from peers. This optional configuration allows certificate validation to succeed only if the certificate chain received from the peer contains at least one policy OID that is configured on the SRX Series device.

On the SRX Series device, policy OIDs are configured in an IKE policy with the **policy-oids** configuration statement at the [**edit security ike policy *policy-name* certificate**] hierarchy level. You can configure up to five policy OIDs. For a peer's certificate to be validated successfully, the peer's certificate chain must contain at least one of the policy OIDs configured on the SRX Series device. Note that the **policy-oids** field in a certificate is optional. If you configure policy OIDs on the SRX Series device but the peer's certificate chain does not contain any policy OIDs, certificate validation fails.

No Policy OIDs Configured on SRX Series Devices

If no policy OID is configured on the SRX Series device, policy validation starts whenever the **requireExplicitPolicy** field is encountered in the certificate chain. A certificate may contain one or more certificate policy OIDs. For policy validation to succeed, there must be a common policy OID in the certificate chain.

Figure 275 shows a certificate chain that consists of certificates for a root CA, three intermediate CAs, and an EE. The CA certificate for Int-CA-2 contains the **requireExplicitPolicy** field; therefore, policy validation starts with Int-CA-2 and continues through EE-1. The certificate for Int-CA-2 contains policy OIDs P1, P2, and P3. The certificate for Int-CA-3 contains policy OIDs P2, P3, and P4. The certificate for EE-1 contains policy OIDs P2 and P5. Because the policy OID P2 is common to the certificates being validated, policy validation succeeds.

Figure 275: Policy Validation with requireExplicitPolicy Field

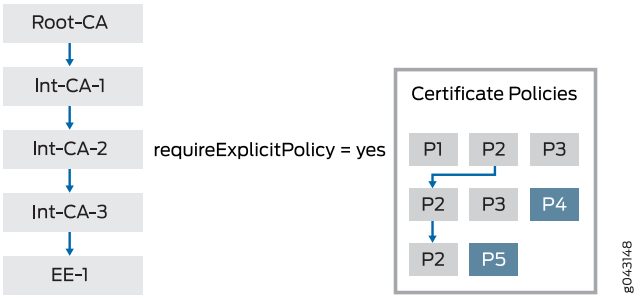

The optional **skipCerts** field in an intermediate CA certificate indicates the number of certificates, including the current CA certificate, that are to be excluded from policy validation. If **skipCerts** is 0, policy validation starts from the current certificate. If **skipCerts** is 1, the current certificate is excluded from policy validation. The value of the **skipCerts**

Copyright © 2016, Juniper Networks, Inc.

field is checked in every intermediate CA certificate. If a **skipCerts** value is encountered that is lower than the current number of certificates being excluded, the lower **skipCerts** value is used.

Figure 276 shows a certificate chain consisting of a root CA, four intermediate CAs, and an EE. The **skipCerts** value in Int-CA-1 is 12, which skips 12 certificates including the certificate for Int-CA-1. However, the **skipCerts** value is checked in every intermediate CA certificate in the chain. The **skipCerts** value in Int-CA-2 is 2, which is lower than 12, so now 2 certificates are skipped. The **skipCerts** value in Int-CA-4 is 5, which is greater than 2, so the Int-CA-4 **skipCerts** value is ignored.

Figure 276: Policy Validation with skipCerts Field

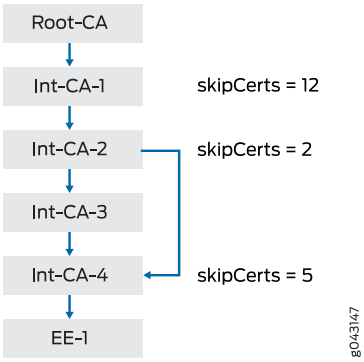

When policy OIDs are configured on the SRX Series device, the certificate fields **requireExplicitPolicy** and **skipCerts** are ignored.

Path Length Validation

Certificate validation can involve a certificate chain that includes a root CA, one or more optional intermediate CAs, and an EE certificate. The number of intermediate CAs can grow depending upon the deployment scenario. Path length validation provides a mechanism to limit the number of intermediate certificates involved in certificate validation. **path-length** is an optional field in an X509 certificate. The value of **path-length** indicates the number of non-self-signed intermediate CA certificates allowed for certificate validation. The last certificate, which is generally the EE certificate, is not included in the path limit. If the root certificate contains a **path-length** value of 0, no intermediate CA certificates are allowed. If the **path-length** value is 1, there can be 0 or 1 intermediate CA certificates.

path-length can be present in multiple CA certificates in the certificate chain. The path length validation always begins with the self-signed root certificate. The path limit is decremented by 1 at each intermediate certificate in the chain. If an intermediate certificate contains a **path-length** value less than the current path limit, the new limit is enforced. On the other hand, if the **path-length** value is larger than the current path limit, it is ignored.

Figure 277 shows a certificate chain that consists of a root CA, four intermediate CAs, and an EE. The **path-length** value in Root-CA is 10, therefore the initial path limit of non-self-signed intermediate CA certificates allowed for certificate validation is 10. At

Copyright © 2016, Juniper Networks, Inc.

Int-CA-1, the path limit is 10-1 or 9. The **path-length** value in Int-CA-1 is 4, which is less than the path limit of 9, so the new path limit becomes 4. At Int-CA-2, the path limit is 4-1 or 3. The **path-length** value in Int-CA-2 is 5, which is larger than the path limit of 3, so it is ignored. At Int-CA-3, the path limit is 3-1 or 2. The **path-length** value in Int-CA-3 is 20, which is larger than the path limit of 2, so it is also ignored.

Figure 277: Path Length Validation

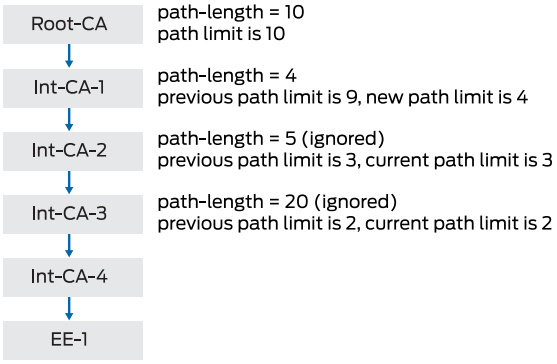

Key Usage

The key usage field in an EE or CA certificate defines the purpose of the key contained in the certificate.

EE Certificates

For EE certificates, if the key usage field is present but the certificate does not contain **digitalSignature** or **nonrepudiation** flags, the certificate is rejected. If the key usage field is not present, then key usage is not checked.

CA Certificates

The key can be used for certificate or CRL signature validation. Because the PKI daemon is responsible for both X509 certificate validation and CRL downloads, key usage must be checked before validating the certificate or CRL.

Certificate Signature Validation

The **keyCertSign** flag indicates that a CA certificate can be used for certificate signature validation. If this flag is not set, certificate validation is aborted.

CRL Signature Validation

In Phase 1 negotiations, participants check the certificate revocation list (CRL) to see if certificates received during an IKE exchange are still valid. The CRL is periodically downloaded for CA profiles configured with CRL as the certificate revocation check. Downloaded CRL files must be verified before they are downloaded into the device. One of the verification steps is to validate the CRL signature using a CA certificate. The downloaded CRL is signed with the CA certificate's private key and it must be verified with the CA certificate's public key stored in the device. The key usage field in the CA

Copyright © 2016, Juniper Networks, Inc.

certificate must contain the **CRLSign** flag to verify the downloaded CRL. If this flag is not present, the CRL is discarded.

Issuer and Subject Distinguished Name Validation

Signature validation is performed for certificates received from a peer as well as for the CRL file downloaded from a CA server. Signature validation involves looking up the CA certificate in a CA database based on the issuer's distinguished name (DN) in the certificate or the CRL being verified.

Figure 278 shows the lookup for CA certificates based on the issuer DN. In the EE certificate, the issuer DN is CA-1, which is the subject DN of the intermediate CA certificate in the chain. In the intermediate CA certificate, the issuer DN is CA-Root, which is the subject DN of the self-signed Root-CA certificate in the chain. In the CRL, the issuer DN is CA-Root, which is the subject DN of the self-signed Root-CA certificate.

Figure 278: Issuer and Subject DN Validation

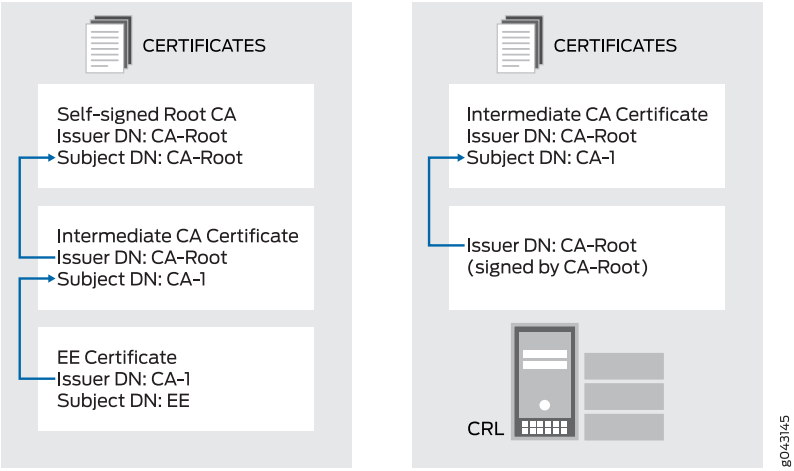

The lookup for the issuer or subject DN must follow these rules for attribute values:

- Attribute values encoded in different ASN.1 types (for example, PrintableString and BMPString) are assumed to represent different strings.
- Attribute values encoded in PrintableString types are not case-sensitive. These attribute values are compared after removing leading and trailing white spaces and converting internal substrings of one or more consecutive white spaces to a single space.
- Attribute values encoded in types other than PrintableString are case-sensitive.

Related Documentation

- Example: Improving Digital Certificate Validation by Configuring Policy OIDs on an SRX Series Device on page 6362
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

Example: Improving Digital Certificate Validation by Configuring Policy OIDs on an SRX Series Device

In some situations, it might be desirable to only accept certificates with known policy object identifiers (OIDs) from peers. This optional configuration allows certificate validation to succeed only if the certificate chain received from the peer contains at least one policy OID that is configured on the SRX Series device. This example shows how to configure policy OIDs in the IKE policy on an SRX Series device.

NOTE: You must ensure that at least one of the policy OIDs configured on the SRX Series device is included in a peer's certificate or certificate chain. Note that the **policy-oids** field in a peer's certificate is optional. If you configure policy OIDs in an IKE policy and the peer's certificate chain does not contain any policy OIDs, certificate validation for the peer fails.

- Requirements on page 6362
- Overview on page 6362
- Configuration on page 6362
- Verification on page 6363

Requirements

Before you begin:

- Ensure that you are using Junos OS Release 12.3X48-D10 or later for SRX Series devices.
- Configure an IPsec VPN tunnel. See "IPsec VPN with Autokey IKE Configuration Overview" on page 6059. The complete IKE phase 1 and phase 2 VPN tunnel configuration is not shown in this example.

Overview

This example shows an IKE policy configuration where policy OIDs 2.16.840.1.101.3.1.48.2 and 5.16.40.1.101.3.1.55.2 are specified. The IKE policy ike_cert_pol references the IKE proposal ike_cert_prop, which is not shown. The local certificate on the SRX Series device is lc-igloo-root.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike policy ike_cert_pol mode main
set security ike policy ike_cert_pol proposals ike_cert_prop
set security ike policy ike_cert_pol certificate local-certificate lc-igloo-root
set security ike policy ike_cert_pol certificate policy-oids 2.16.840.1.101.3.1.48.2
set security ike policy ike_cert_pol certificate policy-oids 5.16.40.1.101.3.1.55.2
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure policy OIDs for certificate validation:

- Configure the IKE policy:

```
[edit security ike policy ike_cert_pol]
user@host# set mode main
user@host# set proposals ike_cert_prop
user@host# set certificate local-certificate lc-igloo-root
user@host# set certificate policy-oids 2.16.840.1.101.3.1.48.2
user@host# set certificate policy-oids 5.16.40.1.101.3.1.55.2
```

Results

From configuration mode, confirm your configuration by entering the **show security ike policy ike_cert_pol** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show security ike policy ike_cert_pol
mode main;
proposals ike_cert_prop;
certificate {
  local-certificate lc-igloo-root;
  policy-oids [ 2.16.840.1.101.3.1.48.2 5.16.40.1.101.3.1.55.2 ];
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the CA Certificate

Purpose

Display the CA certificate configured on the device.

Action

From operational mode, enter the **show security pki ca-certificate ca-profile ca-tmp** command.

```
user@host> show security pki ca-certificate ca-profile ca-tmp detail
                  Certificate identifier: ca-tmp
                         Certificate version: 3
                         Serial number: 00000047
                         Issuer:
                           Organization: U.S. Government,
                        Organizational unit: DoD, Organizational unit: Testing,
 Country: US,
                           Common name: Trust Anchor
                         Subject:
                           Organization: U.S. Government,
                        Organizational unit: Dod, Organizational unit: Testing,
 Country: US,
                           Common name: CA1-PP.01.03
                         Subject string:
                           C=US, O=U.S. Government, OU=Dod, OU=Testing,
```

Copyright © 2016, Juniper Networks, Inc.

```
CN=CA1-PP.01.03
                          Validity:
                            Not before: 01- 1-1998 12:01 UTC
                            Not after: 01- 1-2048 12:01 UTC

?Public key algorithm: rsaEncryption(1024 bits)
  30:81:89:02:81:81:00:cb:fd:78:0c:be:87:ac:cd:c0:33:66:a3:18
  9e:fd:40:b7:9b:bc:dc:66:ff:08:45:f7:7e:fe:8e:d6:32:f8:5b:75
  db:76:f0:4d:21:9a:6e:4f:04:21:4c:7e:08:a1:f9:3d:ac:8b:90:76
  44:7b:c4:e9:9b:93:80:2a:64:83:6e:6a:cd:d8:d4:23:dd:ce:cb:3b
  b5:ea:da:2b:40:8d:ad:a9:4d:97:58:cf:60:af:82:94:30:47:b7:7d
  88:c3:76:c0:97:b4:6a:59:7e:f7:86:5d:d8:1f:af:fb:72:f1:b8:5c
  2a:35:1e:a7:9e:14:51:d4:19:ae:c7:5c:65:ea:f5:02:03:01:00:01
Signature algorithm: sha1WithRSAEncryption
Certificate Policy:
  Policy Identifier = 2.16.840.1.101.3.1.48.2
Use for key: CRL signing, Certificate signing
Fingerprint:
  e0:b3:2f:2e:a1:c5:ee:ad:af:dd:96:85:f6:78:24:c5:89:ed:39:40 (sha1)
  f3:47:6e:55:bc:9d:80:39:5a:40:70:8b:10:0e:93:c5 (md5)
```

Verifying Policy OID Validation

Purpose If the peer's certificate is successfully validated, IKE and IPsec security associations are established. If the validation of the peer's certificate fails, no IKE security association is established.

Action From operational mode, enter the **show security ike security-associations** and **show security ipsec security-associations** commands.

```
user@host> show security ike security-associations
node0:
--------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

821765168 UP   88875c981252c1d8  b744ac9c21bde57e  IKEv2          31.1.1.2

1106977837 UP  1a09e32d1e6f20f1  e008278091060acb  IKEv2          12.1.1.202

user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <213909506 ESP:aes-cbc-192/sha256 8cb9e40a 1295/ unlim - root 500 31.1.1.2

  >213909506 ESP:aes-cbc-192/sha256 8271d2b2 1295/ unlim - root 500 31.1.1.2

  <218365954 ESP:aes-cbc-192/sha256 d0153bc0 1726/ unlim - root 1495 12.1.1.202

  >218365954 ESP:aes-cbc-192/sha256 97611813 1726/ unlim - root 1495 12.1.1.202
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. In this case, check for the PKID_CERT_POLICY_CHECK_FAIL message in the system logs. This message indicates that the peer's certificate chain does not contain a policy OID that is configured

Copyright © 2016, Juniper Networks, Inc.

on the SRX Series device. Check the **policy-oids** values in the peer's certificate chain with the values configured on the SRX Series device.

It might also be that the peer's certificate chain does not contain any **policy-oids** fields, which are optional fields. If this is the case, certificate validation fails if there are any policy OIDs configured on the SRX Series device.

Related Documentation

- Understanding Digital Certificate Validation on page 6357

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 266

Generating a Public-Private Key Pair

- Understanding Public Key Cryptography on page 6367
- Example: Generating a Public-Private Key Pair on page 6368

Understanding Public Key Cryptography

The public-private key pairs used in public key cryptography play an important role in the use of digital certificates. A public-private key pair encrypts and decrypts data. Data encrypted with a public key, which the owner makes available to the public, can be decrypted with the corresponding private key only, which the owner keeps secret and protected. For example, if Alice wants to send Bob an encrypted message, Alice can encrypt it with Bob's public key and send it to him. Bob then decrypts the message with his private key.

The reverse process is also useful: encrypting data with a private key and decrypting it with the corresponding public key. This process is known as creating a digital signature. For example, if Alice wants to present her identity as the sender of a message, she can encrypt the message with her private key and send the message to Bob. Bob then decrypts the message with Alice's public key, thus verifying that Alice is indeed the sender.

When you generate a public-private key pair, the device automatically saves the key pair in a file in the certificate store, where it is subsequently used in certificate request commands. The generated key pair is saved as *certificate-id*.

NOTE: The default RSA and DSA key size is 1024 bits. Simple Certificate Enrollment Protocol (SCEP) supports RSA certificates only. CMPv2 supports RSA, DSA, and ECDSA certificate types.

NOTE: If the device renews a great number of certificates at once, thus using up keys rapidly, it might run out of pregenerated keys and have to generate them promptly for each new request. In this case, the generation of keys might affect the performance of the device, especially in a high-availability environment where the performance of the device might slow down for a number of minutes.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Generating a Public-Private Key Pair on page 6368
- Digital Certificates Configuration Overview on page 6355

Example: Generating a Public-Private Key Pair

This example shows how to generate a public-private key pair.

- Requirements on page 6368
- Overview on page 6368
- Configuration on page 6368
- Verification on page 6368

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you generate a public-private key pair named ca-ipsec.

Configuration

Step-by-Step Procedure

To generate a public-private key pair:

1. Create a certificate key pair.

```
[edit]
user@host> request security pki generate-key-pair certificate-id ca-ipsec
```

Verification

After the public-private key pair is generated, the Juniper Networks device displays the following:

```
generated key pair ca-ipsec, key size 1024 bits
```

Related Documentation

- Understanding Public Key Cryptography on page 6367
- Example: Verifying Certificate Validity on page 6429
- Digital Certificates Configuration Overview on page 6355

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 267

Configuring Certificate Authority Profiles

- Understanding Certificate Authority Profiles on page 6369
- Example: Configuring a CA Profile on page 6369

Understanding Certificate Authority Profiles

A certificate authority (CA) profile configuration contains information specific to a CA. You can have multiple CA profiles on the device. For example, you might have one profile for Microsoft and one for Entrust. Each profile is associated with a CA certificate. If you want to load a new CA certificate without removing the older one, you must create a new CA profile (for example, Microsoft-2008).

NOTE: **The following CAs are supported: Entrust, Microsoft, and Verisign. SCEP only supports the Microsoft CA.**

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Configuring a CA Profile on page 6369

Example: Configuring a CA Profile

This example shows how to configure a CA profile.

- Requirements on page 6369
- Overview on page 6370
- Configuration on page 6370
- Verification on page 6370

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you create a CA profile called ca-profile-ipsec with CA identity microsoft-2008. The configuration specifies that the CRL be refreshed every 48 hours, and the location to retrieve the CRL is http://www.my-ca.com. Within the example, you set the enrollment retry value to 20. (The default retry value is 10.)

Automatic certificate polling is set to every 30 minutes. If you configure retry only without configuring a retry interval, then the default retry interval is 900 seconds (or 15 minutes). If you do not configure retry or a retry interval, then there is no polling.

Configuration

Step-by-Step Procedure

To configure a CA profile:

1. Create a CA profile.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile-ipsec ca-identity microsoft-2008
     revocation-check crl refresh-interval 48 url http://www.my-ca.com/my-crl.crl
   ```

2. Specify the enrollment retry value.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile-ipsec enrollment retry 20
   ```

3. Specify the time interval in seconds between attempts to automatically enroll the CA certificate online.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile-ipsec enrollment retry-interval
     1800
   ```

4. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security pki** command.

Related Documentation

- Understanding Certificate Authority Profiles on page 6369
- Digital Certificates Configuration Overview on page 6355

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 268

Configuring CA and Local Certificates

- Understanding Online CA Certificate Enrollment on page 6371
- Understanding Local Certificate Requests on page 6371
- Enrolling a CA Certificate Online Using SCEP on page 6372
- Example: Enrolling a Local Certificate Online Using SCEP on page 6373
- Example: Using SCEP to Automatically Renew a Local Certificate on page 6375
- Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server on page 6376
- Understanding Certificate Loading on page 6377
- Example: Loading CA and Local Certificates Manually on page 6378
- Deleting Certificates (CLI Procedure) on page 6379
- Example: Configuring PKI on page 6380

Understanding Online CA Certificate Enrollment

With Simple Certificate Enrollment Protocol (SCEP), you can configure your Juniper Networks device to obtain a certificate authority (CA) certificate online and start the online enrollment for the specified certificate ID. The CA public key verifies certificates from remote peers.

Related Documentation

- Understanding Public Key Cryptography on page 6367
- Understanding Certificates and PKI on page 6347
- Enrolling a CA Certificate Online Using SCEP on page 6372
- Example: Enrolling a Local Certificate Online Using SCEP on page 6373

Understanding Local Certificate Requests

When you create a local certificate request, the device generates a CA certificate in PKCS #10 format from a key pair you previously generated using the same certificate ID.

Copyright © 2016, Juniper Networks, Inc.

A subject name is associated with the local certificate request in the form of a common name (CN), organizational unit (OU), organization (O), locality (L), state (ST), country (C), and domain component (DC). Additionally, a subject alternative name is associated in the following form:

- IP address
- E-mail address
- Fully qualified domain name (FQDN)

NOTE: Some CAs do not support an e-mail address as the domain name in a certificate. If you do not include an e-mail address in the local certificate request, you cannot use an e-mail address as the local IKE ID when configuring the device as a dynamic peer. Instead, you can use a fully qualified domain name (if it is in the local certificate), or you can leave the local ID field empty. If you do not specify a local ID for a dynamic peer, enter the *hostname.domain-name* of that peer on the device at the other end of the IPsec tunnel in the peer ID field.

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server on page 6376

Enrolling a CA Certificate Online Using SCEP

Before you begin:

1. Generate a public and private key pair. See "Example: Generating a Public-Private Key Pair" on page 6368.
2. Create a CA profile. See "Example: Configuring a CA Profile" on page 6369.

To enroll a CA certificate online:

1. Retrieve the CA certificate online using SCEP. (The attributes required to reach the CA server are obtained from the defined CA profile.)

   ```
   user@host> request security pki ca-certificate enroll ca-profile ca-profile-ipsec
   ```

   The command is processed synchronously to provide the fingerprint of the received CA certificate.

   ```
   Fingerprint:
   e6:fa:d6:da:e8:8d:d3:00:e8:59:12:e1:2c:b9:3c:c0:9d:6c:8f:8d (sha1)
   82:e2:dc:ea:48:4c:08:9a:fd:b5:24:b0:db:c3:ba:59 (md5)
   Do you want to load the above CA certificate ? [yes,no]
   ```

2. Confirm that the correct certificate is loaded. The CA certificate is loaded only when you type **yes** at the CLI prompt.

Copyright © 2016, Juniper Networks, Inc.

For more information on the certificate, such as the bit length of the key pair, use the command **show security pki ca-certificate**.

Related Documentation

- Understanding Online CA Certificate Enrollment on page 6371
- Digital Certificates Configuration Overview on page 6355
- Example: Enrolling a Local Certificate Online Using SCEP on page 6373
- Example: Using SCEP to Automatically Renew a Local Certificate on page 6375

Example: Enrolling a Local Certificate Online Using SCEP

This example shows how to enroll a local certificate online using Simple Certificate Enrollment Protocol (SCEP).

- Requirements on page 6373
- Overview on page 6373
- Configuration on page 6374
- Verification on page 6374

Requirements

Before you begin:

- Generate a public and private key pair. See "Example: Generating a Public-Private Key Pair" on page 6368.
- Configure a certificate authority profile. See "Example: Configuring a CA Profile" on page 6369.
- For SCEP, enroll the CA certificate. See "Enrolling a CA Certificate Online Using SCEP" on page 6372.

Overview

In this example, you configure your Juniper Networks device to obtain a local certificate online and start the online enrollment for the specified certificate ID with SCEP. You specify the URL path to the CA server in the CA profile name **ca-profile-ipsec**.

You use the **request security pki local-certificate enroll scep** command to start the online enrollment for the specified certificate ID. You must specify the CA profile name (for example, **ca-profile-ipsec**), the certificate ID corresponding to a previously generated key-pair (for example, **qqq**), and the following information:

- The challenge password provided by the CA administrator for certificate enrollment and reenrollment.
- At least one of the following values:
  - The domain name to identify the certificate owner in IKE negotiations—for example, **qqq.example.net**.

Copyright © 2016, Juniper Networks, Inc.

  - The identity of the certificate owner for IKE negotiation with the e-mail statement—for example, **qqq@example.net**.
  - The IP address if the device is configured for a static IP address—for example, **10.10.10.10**.
- Specify the subject name in the distinguished name format in quotation marks, including the domain component (DC), common name (CN), serial number (SN), organizational unit name (OU), organization name (O), locality (L), state (ST), and country (C).

Once the device certificate is obtained and the online enrollment begins for the certificate ID. The command is processed asynchronously.

Configuration

Step-by-Step Procedure

To enroll a local certificate online:

1. Specify the CA profile.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile-ipsec enrollment url
     path-to-ca-server
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

3. Initiate the enrollment process by running the operational mode command.

   ```
   user@host> request security pki local-certificate enroll scep ca-profile
     ca-profile-ipsec certificate-id qqq challenge-password ca-provided-password
     domain-name qqq.example.net email qqq@example.net ip-address 10.10.10.10
     subject DC=example, CN=router3, SN, OU=marketing, O=example, L=sunnyvale,
     ST=california, C=us
   ```

NOTE: If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR).

Verification

To verify the configuration is working properly, enter the **show security pki** command.

Related Documentation

- Digital Certificates Configuration Overview on page 6355
- Enrolling Digital Certificates Online: Configuration Overview on page 6355

Copyright © 2016, Juniper Networks, Inc.

Example: Using SCEP to Automatically Renew a Local Certificate

You can use either Certificate Management Protocol version 2 (CMPv2) or Simple Certificate Enrollment Protocol (SCEP) to enroll digital certificates. This example shows how to renew the local certificates automatically using SCEP.

- Requirements on page 6375
- Overview on page 6375
- Configuration on page 6376
- Verification on page 6376

Requirements

Before you begin:

- Obtain a certificate either on line or manually. See "Enrolling Digital Certificates Online: Configuration Overview" on page 6355.
- Obtain a local certificate. See "Example: Enrolling a Local Certificate Online Using SCEP" on page 6373.

Overview

You can enable the device to automatically renew certificates that were acquired by online enrollment or loaded manually. Automatic certificate renewal saves you from having to remember to renew certificates on the device before they expire, and helps to maintain valid certificates at all times.

Automatic certificate renewal is disabled by default. You can enable automatic certificate renewal and configure the device to automatically send out a request to reenroll a certificate before it expires. You can specify when the certificate reenrollment request is to be sent; the trigger for reenrollment is the percentage of the certificate's lifetime that remains before expiration. For example, if the renewal request is to be sent when the certificate's remaining lifetime is 10 percent, then configure 10 for the reenrollment trigger.

For this feature to work, the device must be able to reach the CA server, and the certificate must be present on the device during the renewal process. Furthermore, you must also ensure that the CA issuing the certificate can return the same DN. The CA must not modify the subject name or alternate subject name extension in the new certificate.

You can enable and disable automatic SCEP certificate renewal either for all SCEP certificates or on a per-certificate basis. You use the **set security pki auto-re-enrollment scep** command to enable and configure certificate reenrollment. In this example, you specify the certificate ID of the CA certificate as **ca-ipsec** and set the CA profile name associated with the certificate to **ca-profile-ipsec**. You set the challenge password for the CA certificate to the challenge password provided by the CA administrator; this password must be the same one configured previously for the CA. You also set the percentage for the reenrollment trigger to **10**. During automatic reenrollment, the Juniper Networks device by default uses the existing key pair. A good security practice is to

Copyright © 2016, Juniper Networks, Inc.

regenerate a new key pair for reenrollment. To generate a new key pair, use the **re-generate-keypair** command.

Configuration

Step-by-Step Procedure

To enable and configure local certificate reenrollment:

1. To enable and configure certificate reenrollment.

```
[edit]
user@host# set security pki auto-re-enrollment scep certificate-id ca-ipsec
  ca-profile-name ca-profile-ipsec challenge-password ca-provided-password
  re-enroll-trigger-time-percentage 10 re-generate-keypair
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show security pki local-certificate detail** operational mode command.

Related Documentation

- Enrolling Digital Certificates Online: Configuration Overview on page 6355

Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server

This example shows how to generate a certificate signing request manually.

- Requirements on page 6376
- Overview on page 6376
- Configuration on page 6377
- Verification on page 6377

Requirements

Generate a public and private key. See "Example: Generating a Public-Private Key Pair" on page 6368.

Overview

In this example, you generate a certificate request using the certificate ID of a public-private key pair you previously generated (ca-ipsec). Then you specify the domain name (example.net) and the associated common name (abc). The certificate request is displayed in PEM format.

You copy the generated certificate request and paste it into the appropriate field at the CA website to obtain a local certificate. (Refer to the CA server documentation to determine where to paste the certificate request.) When the PKCS #10 content is displayed, the MD5 hash and SHA-1 hash of the PKCS #10 file is also displayed.

Copyright © 2016, Juniper Networks, Inc.

Configuration

Step-by-Step Procedure

To generate a local certificate manually:

1. Specify certificate ID, domain name, and common name.

```
user@host> request security pki generate-certificate-request certificate-id ca-ipsec
  domain-name example.net subject CN=abc
```

Verification

To view the certificate signing request, enter the **show security pki certificate-request detail** command.

```
Certificate identifier: ca-ipsec
Certificate version: 1
Issued to: CN = abc
Public key algorithm: rsaEncryption(1024 bits)
30:81:89:02:81:81:00:da:ea:cd:3a:49:1f:b7:33:3c:c5:50:fb:57
de:17:34:1c:51:9b:7b:1c:e9:1c:74:86:69:a4:36:77:13:a7:10:0e
52:f4:2b:52:39:07:15:3f:39:f5:49:d6:86:70:4b:a6:2d:73:b6:68
39:d3:6b:f3:11:67:ee:b4:40:5b:f4:de:a9:a4:0e:11:14:3f:96:84
03:3c:73:c7:75:f5:c4:c2:3f:5b:94:e6:24:aa:e8:2c:54:e6:b5:42
c7:72:1b:25:ca:f3:b9:fa:7f:41:82:6e:76:8b:e6:d7:d2:93:9b:38
fe:fd:71:01:2c:9b:5e:98:3f:0c:ed:a9:2b:a7:fb:02:03:01:00:01
Fingerprint:
0f:e6:2e:fc:6d:52:5d:47:6e:10:1c:ad:a0:8a:4c:b7:cc:97:c6:01 (sha1)
f8:e6:88:53:52:c2:09:43:b7:43:9c:7a:a2:70:98:56 (md5)
```

Related Documentation

- Understanding Local Certificate Requests on page 6371
- Digital Certificates Configuration Overview on page 6355

Understanding Certificate Loading

After you download certificates from a CA, you transfer them to the device (for example, using FTP), and then load them.

You can load the following certificate files onto a device running Junos OS:

- A local or end-entity (EE) certificate that identifies your local device. This certificate is your public key.
- A CA certificate that contains the CA's public key.
- A CRL that lists any certificates revoked by the CA.

NOTE: **You can load multiple EE certificates onto the device.**

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Loading CA and Local Certificates Manually on page 6378

Copyright © 2016, Juniper Networks, Inc.

Example: Loading CA and Local Certificates Manually

This example shows how to load CA and local certificates manually.

- Requirements on page 6378
- Overview on page 6378
- Configuration on page 6378
- Verification on page 6379

Requirements

Before you begin:

- Generate a public-private key pair. See "Example: Generating a Public-Private Key Pair" on page 6368.
- Create a CA profile. See "Understanding Certificate Authority Profiles" on page 6369.

NOTE: CA Profile is only required for the CA certificate and not for the local certificate

- Generate a certificate request. See "Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server" on page 6376.

Overview

In this example, you download the local.cert and ca.cert certificates and save them to the /var/tmp/ directory on the device.

Configuration

Step-by-Step Procedure

To load the certificate files onto a device:

1. Load the local certificate.

```
[edit]
user@host> request security pki local-certificate load certificate-id local.cert
  filename /var/tmp/local.cert
```

2. Load the CA certificate.

```
[edit]
user@host> request security pki ca-certificate load ca-profile ca-profile-ipsec
  filename /var/tmp/ca.cert
```

3. Examine the fingerprint of the CA certificate, if it is correct for this CA certificate say yes to accept.

Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the certificates loaded properly, enter the **show security pki local-certificate** and **show security pki ca-certificate** commands in operational mode.

```
Fingerprint:
e8:bf:81:6a:cd:26:ad:41:b3:84:55:d9:10:c4:a3:cc:c5:70:f0:7f (sha1)
19:b0:f8:36:e1:80:2c:30:a7:31:79:69:99:b7:56:9c (md5)
Do you want to load this CA certificate ? [yes,no] (no) yes
```

Related Documentation

- Understanding Certificate Loading on page 6377
- Digital Certificates Configuration Overview on page 6355
- Example: Using SCEP to Automatically Renew a Local Certificate on page 6375
- Example: Verifying Certificate Validity on page 6429
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428

Deleting Certificates (CLI Procedure)

You can delete a local or trusted CA certificate that is automatically or manually generated.

Use the following command to delete a local certificate:

user@host> clear security pki local certificate certificate-id (*certificate-id*| all | system-generated)

Specify a certificate ID to delete a local certificate with a specific ID, use **all** to delete all local certificates, or specify **system-generated** to delete the automatically generated self-signed certificate.

When you delete an automatically generated self-signed certificate, the device generates a new one.

To delete a CA certificate:

user@host> clear security pki ca-certificate ca-profile (*ca-profile-name* | all)

Specify a CA profile to delete a specific CA certificate, or use **all** to delete all CA certificates present in the persistent store.

NOTE: You are asked for confirmation before a CA certificate can be deleted.

Related Documentation

- Digital Certificates Configuration Overview on page 6355

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring PKI

This example shows how to configure, verify, and troubleshoot the PKI. This topic includes the following sections:

- Requirements on page 6380
- Overview on page 6380
- Configuration on page 6384
- Verification on page 6392
- Troubleshooting IKE, PKI, and IPsec Issues on page 6398

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later
- Juniper Networks security devices

Before you begin:

- Ensure that the internal LAN interface of the SRX Series device is ge-0/0/0 in zone trust and has a private IP subnet.
- Ensure that the Internet interface of the device is ge-0/0/3 in zone untrust and has a public IP.
- Ensure that all traffic between the local and remote LANs is permitted, and traffic can be initiated from either side.
- Ensure that the SSG5 has been preconfigured correctly and loaded with a ready-to-use local certificate, CA certificate, and CRL.
- Ensure that the SSG5 device is configured to use the FQDN of ssg5.example.net (IKE ID).
- Ensure that PKI certificates with 1024-bit keys are used for the IKE negotiations on both sides.
- Ensure that the CA is a standalone CA at the domain example.com for both VPN peers.

Overview

Figure 279 shows the network topology used for this example to configure a policy-based IPsec VPN to allow data to be securely transferred between a corporate office and a remote office.

Copyright © 2016, Juniper Networks, Inc.

Figure 279: Network Topology Diagram

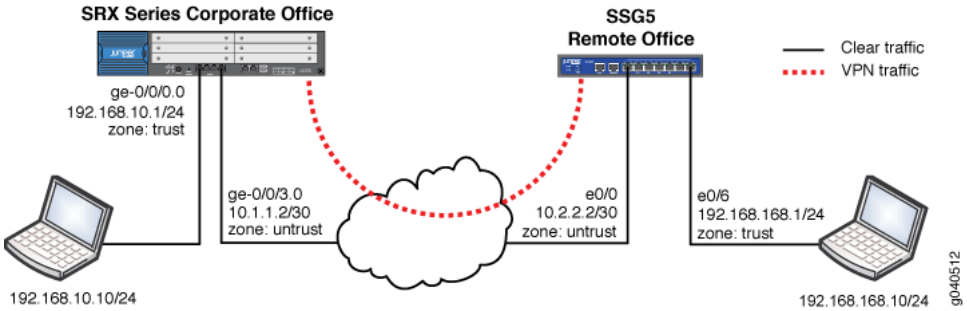

NOTE: The PKI administration is the same for both policy-based VPNs and route-based VPNs.

In this example, the VPN traffic is incoming on interface ge-0/0/0.0 with the next hop of 10.1.1.1. Thus the traffic is outgoing on interface ge-0/0/3.0. Any tunnel policy must consider incoming and outgoing interfaces.

NOTE: Optionally, you can use a dynamic routing protocol such as OSPF (not described in this document). When processing the first packet of a new session, the device running Junos OS first performs a route lookup. The static route, which is also the default route, dictates the zone for the outgoing VPN traffic.

Many CAs use hostnames (for example, FQDN) to specify various elements of the PKI. Because the CDP is usually specified using a URL containing an FQDN, you must configure a DNS resolver on the device running Junos OS.

The certificate request can be generated by the following methods:

- Creating a CA profile to specify the CA settings
- Generating the PKCS10 certificate request

The PKCS10 certificate request process involves generating a public or private key pair and then generating the certificate request itself, using the key pair.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Take note of the following information about the CA profile:

- The CA profile defines the attributes of a certificate authority.
- Each CA profile is associated with a CA certificate. If a new or renewed CA certificate needs to be loaded without removing the older CA certificate, a new profile must be created. This profile can also be used for online fetching of the CRL.
- There can be multiple such profiles present in the system created for different users.

NOTE: If you specify a CA administrator e-mail address to send the certificate request to, then the system composes an e-mail from the certificate request file and forwards it to the specified e-mail address. The e-mail status notification is sent to the administrator.

NOTE: The certificate request can be sent to the CA through an out-of-band method.

The following options are available to generate the PKCS10 certificate request:

- **certificate-id** — Name of the local digital certificate and the public/private key pair. This ensures that the proper key pair is used for the certificate request and ultimately for the local certificate.
- **subject** — Distinguished name format that contains the common name, department, company name, state, and country:
  - CN — Common name
  - OU — Department
  - O — Company name
  - L — Locality
  - ST — State
  - C — Country
  - CN — Phone
  - DC — Domain component

  NOTE: You are not required to enter all subject name components. Note also that you can enter multiple values of each type.

Copyright © 2016, Juniper Networks, Inc.

- **domain-name** — FQDN. The FQDN provides the identity of the certificate owner for IKE negotiations and provides an alternative to the subject name.
- **filename (path | terminal)** — (Optional) Location where the certificate request should be placed, or the login terminal.
- **ip-address** — (Optional) IP address of the device.
- **email** — (Optional) E-mail address of the CA administrator.

NOTE: You must use a domain-name, an ip-address, or an e-mail address.

The generated certificate request is stored in a specified file location. A local copy of the certificate request is saved in the local certificate storage. If the administrator reissues this command, the certificate request is generated again.

The PKCS10 certificate request is stored in a specified file and location, from which you can download it and send it to the CA for enrollment. If you have not specified the filename or location, you can get PKCS10 certificate request details by using the **show security pki certificate-request certificate-id <id-name>** command in the CLI. You can copy the command output and paste it into a Web front end for the CA server or into an e-mail.

The PKCS10 certificate request is generated and stored on the system as a pending certificate or certificate request. An e-mail notification is sent to the administrator of the CA (in this example, certadmin@example.com).

NOTE: A unique identity called certificate-ID is used to name the generated key pair. This ID is also used in certificate enrollment and request commands to get the right key pair. The generated key pair is saved in the certificate store in a file with the same name as the certificate-ID. The file size can be 512, 1024, or 2048 bits.

NOTE:

A default (fallback) profile can be created if intermediate CAs are not preinstalled in the device. The default profile values are used in the absence of a specifically configured CA profile.

In the case of a CDP, the following order is followed:

- Per CA profile
- CDP embedded in CA certificate
- Default CA profile

We recommend using a specific CA profile instead of a default profile.

Copyright © 2016, Juniper Networks, Inc.

The administrator submits the certificate request to the CA. The CA administrator verifies the certificate request and generates a new certificate for the device. The administrator for the Juniper Networks device retrieves it, along with the CA certificate and CRL.

The process of retrieving the CA certificate, the device's new local certificate, and the CRL from the CA depends on the CA configuration and software vendor in use.

NOTE:

Junos OS supports the following CA vendors:

- Entrust
- Verisign
- Microsoft

Although other CA software services such as OpenSSL can be used to generate certificates, these certificates are not verified by Junos OS.

Configuration

- PKI Basic Configuration on page 6384
- Configuring a CA Profile on page 6385
- Generating a Public-Private Key Pair on page 6386
- Enrolling a Local Certificate on page 6386
- Loading CA and Local Certificates on page 6387
- Configuring the IPsec VPN with the Certificates on page 6389

PKI Basic Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PKI:

1. Configure an IP address and protocol family on the Gigabit Ethernet interfaces.

   ```
   [edit interfaces]
   user@host# set ge-0/0/0 unit 0 family inet address 192.168.10.1/24
   user@host# set ge-0/0/3 unit 0 family inet address 10.1.1.2/30
   ```

2. Configure a default route to the Internet next hop.

   ```
   [edit]
   user@host# set routing-options static route 0.0.0.0/0 next-hop 10.1.1.1
   ```

3. Set the system time and date.

   ```
   [edit]
   user@host# set system time-zone PST8PDT
   ```


Copyright © 2016, Juniper Networks, Inc.

After the configuration is committed, verify the clock settings using the **show system uptime** command.

```
user@host> show system uptime
Current time: 2007-11-01 17:57:09 PDT
System booted: 2007-11-01 14:36:38 PDT (03:20:31 ago)
Protocols started: 2007-11-01 14:37:30 PDT (03:19:39 ago)
Last configured: 2007-11-01 17:52:32 PDT (00:04:37 ago) by root
5:57PM up 3:21, 4 users, load averages: 0.00, 0.00, 0.00
```

4. Set the NTP server address.

```
user@host> set date ntp 130.126.24.24
1 Nov 17:52:52 ntpdate[5204]: step time server 172.16.24.24 offset -0.220645
 sec
```

5. Set the DNS configuration.

```
[edit]
user@host# set system name-server 172.31.2.1
user@host# set system name-server 172.31.2.2
```

Configuring a CA Profile

Step-by-Step Procedure

1. Create a trusted CA profile.

```
[edit]
user@host# set security pki ca-profile ms-ca ca-identity example.com
```

2. Create a revocation check to specify a method for checking certificate revocation.

```
[edit]
user@host# set security pki ca-profile ms-ca revocation-check crl
```

NOTE: You can use the disable option to disable the revocation check or select the crl option to configure the CRL attributes.

You can select the disable on-download-failure option to allow the sessions matching the CA profile, when CRL download failed for a CA profile. The sessions will be allowed only if no old CRL is present in the same CA profile.

3. Set the refresh interval, in hours, to specify the frequency in which to update the CRL. The default values are next-update time in CRL, or 1 week, if no next-update time is specified.

```
[edit]
user@host# set security pki ca-profile ms-ca revocation-check crl refresh-interval
  48
```

4. Specify the location (URL) to retrieve the CRL (HTTP or LDAP). By default, the URL is empty and uses CDP information embedded in the CA certificate.

```
[edit]
user@host# set security pki ca-profile ms-ca revocation-check crl url
  http://srv1.example.com/CertEnroll/EXAMPLE.crl
```

Copyright © 2016, Juniper Networks, Inc.

NOTE: Currently you can configure only one URL. Support for backup URL configuration is not available.

5. Specify an e-mail address to send the certificate request directly to a CA administrator.

```
user@host# set security pki ca-profile ms-ca administrator email-address
  certadmin@example.com
```

6. Commit the configuration:

```
user@host# commit and-quit
commit complete
Exiting configuration mode
```

Generating a Public-Private Key Pair

Step-by-Step Procedure

When the CA profile is configured, the next step is to generate a key pair on the Juniper Networks device. To generate the private and public key pair:

1. Create a certificate key pair.

```
user@host> request security pki generate-key-pair certificate-id ms-cert size 1024
```

Results

After the public-private key pair is generated, the Juniper Networks device displays the following:

```
Generated key pair ms-cert, key size 1024 bits
```

Enrolling a Local Certificate

Step-by-Step Procedure

1. Generate a local digital certificate request in the PKCS-10 format.

```
user@host> request security pki generate-certificate-request certificate-id ms-cert subject
"CN=john doe,CN=10.1.1.2,OU=sales,O=example, L=Sunnyvale,ST=CA,C=US" email
user@example.net filename ms-cert-req
Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIB3DCCAUUCAQAwbDERMA8GA1UEAxMIam9obiBkb2UxDjAMBgNVBAsTBXNhbGVz
MRkwFwYDVQQKExBKdW5pcGVyIE5ldHdvcmtzMRIwEAYDVQQHEwlTdW5ueXZhbGUx
CzAJBgNVBAgTAkNBMQswCQYDVQQGEwJVUzCBnzANBgkqhkiG9w0BAQEFAAOBjQAw
gYkCgYEA5EG6sgG/CTFzX6KC/hz6CzalOBxakUxfGxF7UWYWHaWFFYLqo6vXNO8r
OS5Yak7rWANAsMob3E2X/1adlQIRi4QFTjkBqGI+MTEDGnqFsJBqrB6oyqGtdcSU
uOqUivMvgKQVCx8hpx99J3EBTurfWL1pCNlBmZggNogb6MbwESOCAwEAAaAwMC4G
CSqGSIb3DQEJDjEhMB8wHQYDVR0RBBYwFIESInVzZXJAanVuaXBlci5uZXQiMA0G
CSqGSIb3DQEBBQUAA4GBAI6GhBaCsXk6/1lE2e5AakFFDhY7oqzHhgd1yMjiSUMV
djmf9JbDz2gM2UKpI+yKgtUjyCK/lV2ui57hpZMvnhAW4AmgwkOJg6mpR5rsxdLr
4/HHSHuEGOF17RHO6x0YwJ+KE1rYDRWj3Dtz447ynaLxcDF7buwd4IrMcRJJI9ws
-----END CERTIFICATE REQUEST-----
Fingerprint:
47:b0:e1:4c:be:52:f7:90:c1:56:13:4e:35:52:d8:8a:50:06:e6:c8 (sha1)
a9:a1:cd:f3:0d:06:21:f5:31:b0:6b:a8:65:1b:a9:87 (md5)
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: In the sample of the PKCS10 certificate, the request starts with and includes the BEGIN CERTIFICATE REQUEST line and ends with and includes the END CERTIFICATE REQUEST line. This portion can be copied and pasted to your CA for enrollment. Optionally, you can also offload the ms-cert-req file and send that to your CA.

2. Generate the PKCS10 certificate request to be sent to the CA.

   user@host> **request security pki generate-certificate-request certificate-id *id-name* subject subject-name (domain-name *domain-name* | ip-address *device-ip* | email *email-id*) filename *filename***

3. Submit the certificate request to the CA, and retrieve the certificate.

Loading CA and Local Certificates

Step-by-Step Procedure

1. Load the local certificate, CA certificate, and CRL.

```
user@host> file copy ftp://192.168.10.10/certnew.cer certnew.cer
  /var/tmp//...transferring.file.....crYdEC/100% of 1459 B 5864 kBps
user@host> file copy ftp:// 192.168.10.10/CA-certnew.cer CA-certnew.cer
  /var/tmp//...transferring.file.....UKXUWu/100% of 1049 B 3607 kBps
user@host> file copy ftp:// 192.168.10.10/certcrl.crl certcrl.crl
  /var/tmp//...transferring.file.....wpqnpA/100% of 401 B 1611 kBps
```

NOTE: You can verify that all files have been uploaded by using the command **file list**.

2. Load the certificate into local storage from the specified external file.

   You must also specify the certificate ID to keep the proper linkage with the private or public key pair. This step loads the certificate into the RAM cache storage of the PKI module, checks the associated private key, and verifies the signing operation.

```
user@host> request security pki local-certificate load certificate-id ms-cert filename
certnew.cer
Local certificate loaded successfully
```

3. Load the CA certificate from the specified external file.

   You must specify the CA profile to associate the CA certificate to the configured profile.

```
user@host> request security pki ca-certificate load ca-profile ms-ca filename
CA-certnew.cer
Fingerprint:
1b:02:cc:cb:0f:d3:14:39:51:aa:0f:ff:52:d3:38:94:b7:11:86:30 (sha1)
90:60:53:c0:74:99:f5:da:53:d0:a0:f3:b0:23:ca:a3 (md5)
Do you want to load this CA certificate ? [yes,no] (no) yes
CA certificate for profile ms-ca loaded successfully
```

4. Load the CRL into the local storage.

Copyright © 2016, Juniper Networks, Inc.

The maximum size of the CRL is 5 MB. You must specify the associated CA profile in the command.

```
user@host> request security pki crl load ca-profile ms-ca filename certcrl.crl
CRL for CA profile ms-ca loaded successfully
```

Results Verify that all local certificates are loaded.

```
user@host> show security pki local-certificate certificate-id ms-cert detail Certificate
identifier: ms-cert
Certificate version: 3
Serial number: 3a01c5a0000000000011
Issuer:
Organization: Example, Organizational unit: example, Country: US, State:
CA, Locality: Sunnyvale,
Common name: Sales
Subject:
Organization: Example, Organizational unit: example, Country: US,
State: CA, Locality: Sunnyvale,
Common name: john doe
Alternate subject: "user@example.net", fqdn empty, ip empty
Validity:
Not before: 11- 2-2007 22:54
Not after: 11- 2-2008 23:04
Public key algorithm: rsaEncryption(1024 bits)
30:81:89:02:81:81:00:e4:41:ba:b2:01:bf:09:31:73:5f:a2:82:fe
1c:fa:0b:36:a5:d0:1c:5a:91:4c:5f:1b:11:7b:51:66:16:1d:a5:85
15:82:ea:a3:ab:d7:34:ef:2b:39:2e:58:6a:4e:eb:58:03:40:b0:ca
1b:dc:4d:97:ff:56:9d:95:02:11:8b:84:05:4e:39:01:a8:62:3e:31
31:03:1a:7a:85:b0:90:6a:ac:1e:a8:ca:a1:ad:75:c4:94:bb:4a:94
8a:f3:2f:80:a4:15:0b:1f:21:a7:1f:7d:27:71:01:4e:ea:df:58:bd
69:08:d9:41:99:98:20:36:88:1b:e8:c6:f0:11:2d:02:03:01:00:01
Signature algorithm: sha1WithRSAEncryption
Distribution CRL:
ldap:///CN=Sales,CN=SRV1,CN=CDP,CN=Public%20Key%20Services,CN=Services,
CN=Configuration,DC=domain,DC=com?certificateRevocationList?base?
objectclass=cRLDistributionPoint
http://srv1.domain.com/CertEnroll/example.crl
Fingerprint:
c9:6d:3d:3e:c9:3f:57:3c:92:e0:c4:31:fc:1c:93:61:b4:b1:2d:58 (sha1)
50:5d:16:89:c9:d3:ab:5a:f2:04:8b:94:5d:5f:65:bd (md5)
```

NOTE: You can display the individual certificate details by specifying certificate-ID in the command line.

Verify all CA certificates or the CA certificates of an individual CA profile (specified).

```
user@host> show security pki ca-certificate ca-profile ms-ca detail
Certificate identifier: ms-ca
Certificate version: 3
Serial number: 44b033d1e5e158b44597d143bbfa8a13
Issuer:
Organization: Example, Organizational unit: example, Country: US, State:
CA, Locality: Sunnyvale,
Common name: example
Subject:
Organization: Example, Organizational unit: example, Country: US, State:
```


Copyright © 2016, Juniper Networks, Inc.

```
CA, Locality: Sunnyvale,
Common name: example
Validity:
Not before: 09-25-2007 20:32
Not after: 09-25-2012 20:41
Public key algorithm: rsaEncryption(1024 bits)
30:81:89:02:81:81:00:d1:9e:6f:f4:49:c8:13:74:c3:0b:49:a0:56
11:90:df:3c:af:56:29:58:94:40:74:2b:f8:3c:61:09:4e:1a:33:d0
8d:53:34:a4:ec:5b:e6:81:f5:a5:1d:69:cd:ea:32:1e:b3:f7:41:8e
7b:ab:9c:ee:19:9f:d2:46:42:b4:87:27:49:85:45:d9:72:f4:ae:72
27:b7:b3:be:f2:a7:4c:af:7a:8d:3e:f7:5b:35:cf:72:a5:e7:96:8e
30:e1:ba:03:4e:a2:1a:f2:1f:8c:ec:e0:14:77:4e:6a:e1:3b:d9:03
ad:de:db:55:6f:b8:6a:0e:36:81:e3:e9:3b:e5:c9:02:03:01:00:01
Signature algorithm: sha1WithRSAEncryption
Distribution CRL:
ldap:///CN=Sales,CN=SRV1,CN=CDP,CN=Public%20Key%20Services,CN=Services,
CN=Configuration,DC=domain,DC=com?certificateRevocationList?base?
objectclass=cRLDistributionPoint
http://srv1.domain.com/CertEnroll/example.crl
Use for key: CRL signing, Certificate signing, Non repudiation
Fingerprint:
1b:02:cc:cb:0f:d3:14:39:51:aa:0f:ff:52:d3:38:94:b7:11:86:30 (sha1)
90:60:53:c0:74:99:f5:da:53:d0:a0:f3:b0:23:ca:a3 (md5)
```

Verify all loaded CRLs or the CRLs of the specified individual CA profile.

```
user@host> show security pki crl ca-profile ms-ca detail
CA profile: ms-ca
CRL version: V00000001
CRL issuer: emailAddress = certadmin@example.net, C = US, ST = CA,
L = Sunnyvale, O = Example, OU = example, CN = example
Effective date: 10-30-2007 20:32
Next update: 11- 7-2007 08:52
```

Verify the certificate path for the local certificate and the CA certificate.

```
user@host> request security pki local-certificate verify certificate-id ms-cert
Local certificate ms-cert verification success

user@host> request security pki ca-certificate verify ca-profile ms-ca
CA certificate ms-ca verified successfully
```

Configuring the IPsec VPN with the Certificates

Step-by-Step Procedure

To configure the IPsec VPN with the certificate, refer to the network diagram shown in Figure 279

1. Configure security zones and assign interfaces to the zones.

   In this example packets are incoming on **ge-0/0/0**, and the ingress zone is the trust zone.

   ```
   [edit security zones security-zone]
   user@host# set trust interfaces ge-0/0/0.0
   user@host# set untrust interfaces ge-0/0/3.0
   ```

2. Configure host-inbound services for each zone.

   Host-inbound services are for traffic destined for the Juniper Networks device. These settings include but are not limited to the FTP, HTTP, HTTPS, IKE, ping, rlogin, RSH, SNMP, SSH, Telnet, TFTP, and traceroute.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security zones security-zone]
user@host# set trust host-inbound-traffic system-services all
user@host# set untrust host-inbound-traffic system-services ike
```

3. Configure the address book entries for each zone.

```
[edit security zones security-zone]
user@host# set trust address-book address local-net 192.168.10.0/24
user@host# set untrust address-book address remote-net 192.168.168.0/24
```

4. Configure the IKE (Phase 1) proposal to use RSA encryption.

```
[edit security ike proposal rsa-prop1]
user@host# set authentication-method rsa-signatures
user@host# set encryption-algorithm 3des-cbc
user@host# set authentication-algorithm sha1
user@host# set dh-group group2
```

5. Configure an IKE policy.

   The phase 1 exchange can take place in either main mode or aggressive mode.

```
[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals rsa-prop1
user@host# set certificate local-certificate ms-cert
user@host# set certificate peer-certificate-type x509- signature
user@host# set certificate trusted-ca use-all
```

6. Configure an IKE gateway.

   In this example, the peer is identified by an FQDN (hostname). Therefore the gateway IKE ID should be the remote peer domain name. You must specify the correct external interface or peer ID to properly identify the IKE gateway during Phase 1 setup.

```
[edit security ike gateway ike-gate]
user@host# set external-interface ge-0/0/3.0
user@host# set ike-policy ike-policy1
user@host# set dynamic hostname ssg5.example.net
```

7. Configure the IPsec policy.

   This example uses the Standard proposal set, which includes **esp-group2-3des-sha1** and **esp-group2- aes128-sha1** proposals. However, a unique proposal can be created and then specified in the IPsec policy if needed.

```
[edit security ipsec policy vpn-policy1]
user@host# set proposal-set standard
user@host# set perfect-forward-secrecy keys group2
```

8. Configure the IPsec VPN with an IKE gateway and IPsec policy.

   In this example, the ike-vpn VPN name must be referenced in the tunnel policy to create a security association. Additionally, if required, an idle time and a proxy ID can be specified if they are different from the tunnel policy addresses.

```
[edit security ipsec vpn ike-vpn ike]
user@host# set gateway ike-gate
user@host# set ipsec-policy vpn-policy1
```

9. Configure bidirectional tunnel policies for VPN traffic.

 Copyright © 2016, Juniper Networks, Inc.

In this example, traffic from the host LAN to the remote office LAN requires a from-zone trust to-zone untrust tunnel policy. However, if a session needs to originate from the remote LAN to the host LAN, then a tunnel policy in the opposite direction from from-zone untrust to-zone trust is also required. When you specify the policy in the opposite direction as the pair-policy, the VPN becomes bidirectional. Note that in addition to the permit action, you also need to specify the IPsec profile to be used. Note that for tunnel policies, the action is always permit. In fact, if you are configuring a policy with the deny action, you will not see an option for specifying the tunnel.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy tunnel-policy-out match source-address local-net
user@host# set policy tunnel-policy-out match destination-address remote-net
user@host# set policy tunnel-policy-out match application any
user@host# set policy tunnel-policy-out then permit tunnel ipsec-vpn ike-vpn
  pair-policy tunnel-policy-in
user@host# top edit security policies from-zone untrust to-zone trust
user@host# set policy tunnel-policy-in match source-address remote-net
user@host# set policy tunnel-policy-in match destination-address local-net
user@host# set policy tunnel-policy-in match application any
user@host# set policy tunnel-policy-in then permit tunnel ipsec-vpn ike-vpn
  pair-policy tunnel-policy-out
```

10. Configure a source NAT rule and a security policy for Internet traffic.

    The device uses the specified source-nat interface, and translates the source IP address and port for outgoing traffic, using the IP address of the egress interface as the source IP address and a random higher port for the source port. If required, more granular policies can be created to permit or deny certain traffic.

    ```
    [edit security nat source rule-set nat-out]
    user@host#set from zone trust
    user@host#set to zone untrust
    user@host#set rule interface-nat match source-address 192.168.10.0/24
    user@host#set rule interface-nat match destination-address 0.0.0.0/0
    user@host#set rule interface-nat then source-nat interface

    [edit security policies from-zone trust to-zone untrust]
    user@host# set policy any-permit match source-address any
    user@host# set policy any-permit match destination-address any
    user@host# set policy any-permit match application any
    user@host# set policy any-permit then permit
    ```

11. Move the tunnel policy above the any-permit policy.

    ```
    [edit security policies from-zone trust to-zone untrust]
    user@host# insert policy tunnel-policy-out before policy any-permit
    ```

NOTE: The security policy should be below the tunnel policy in the hierarchy because the policy list is read from top to bottom. If this policy were above the tunnel policy, then the traffic would always match this policy and would not continue to the next policy. Thus no user traffic would be encrypted.

Copyright © 2016, Juniper Networks, Inc.

12. Configure the tcp-mss setting for TCP traffic across the tunnel.

    TCP-MSS is negotiated as part of the TCP 3-way handshake. It limits the maximum size of a TCP segment to accommodate the MTU limits on a network. This is very important for VPN traffic because the IPsec encapsulation overhead along with the IP and frame overhead can cause the resulting ESP packet to exceed the MTU of the physical interface, causing fragmentation. Because fragmentation increases the bandwidth and device resources usage, and in general it should be avoided.

    The recommended value to use for tcp-mss is 1350 for most Ethernet-based networks with an MTU of 1500 or higher. This value might need to be altered if any device in the path has a lower value of MTU or if there is any added overhead such as PPP, Frame Relay, and so on. As a general rule, you might need to experiment with different tcp-mss values to obtain optimal performance.

    user@host# **set security flow tcp-mss ipsec-vpn mss** ***mss-value***

    Example:
    [edit]
    user@host# **set security flow tcp-mss ipsec-vpn mss 1350**
    user@host# **commit and-quit**
    commit complete
    Exiting configuration mode

Verification

Confirm that the configuration is working properly.

- Confirming IKE Phase 1 Status on page 6392
- Getting Details on Individual Security Associations on page 6393
- Confirming IPsec Phase 2 Status on page 6394
- Displaying IPsec Security Association Details on page 6394
- Checking IPsec SA Statistics on page 6396
- Testing Traffic Flow Across the VPN on page 6397
- Confirming the Connectivity on page 6397

Confirming IKE Phase 1 Status

Purpose Confirm the VPN status by checking any IKE Phase 1 security associations status.

PKI related to IPsec tunnels is formed during Phase 1 setup. Completion of Phase 1 indicates that PKI was successful.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations

Index Remote Address State Initiator cookie Responder cookie Mode
 2010.2.2.2 UP af4f78bc135e4365 48a35f853ee95d21 Main
```


Copyright © 2016, Juniper Networks, Inc.

Meaning The output indicates that:

- The remote peer is 10.2.2.2 and the status is UP, which means the successful association of Phase 1 establishment.
- The remote peer IKE ID, IKE policy, and external interfaces are all correct.
- Index 20 is a unique value for each IKE security association. You can use this output details to get further details on each security association. See "Getting Details on Individual Security Associations" on page 6393.

Incorrect output would indicate that:

- The remote peer status is Down.
- There are no IKE security associations .
- There are IKE policy parameters, such as the wrong mode type (Aggr or Main), PKI issues, or Phase 1 proposals (all must match on both peers). For more information, see "Troubleshooting IKE, PKI, and IPsec Issues" on page 6398.
- External interface is invalid for receiving the IKE packets. Check the configurations for PKI-related issues, check the key management daemon (kmd) log for any other errors, or run traceoptions to find the mismatch. For more information, see "Troubleshooting IKE, PKI, and IPsec Issues" on page 6398.

Getting Details on Individual Security Associations

Purpose Get details on individual IKE.

Action From operational mode, enter the **show security ike security-associations index 20 detail** command.

```
user@host> show security ike security-associations index 20 detail
IKE peer 10.2.2.2, Index 20,
Role: Responder, State: UP
Initiator cookie: af4f78bc135e4365, Responder cookie: 48a35f853ee95d21
Exchange type: Main, Authentication method: RSA-signatures
Local: 10.1.1.2:500, Remote: 10.2.2.2:500
Lifetime: Expires in 23282 seconds
Algorithms:
Authentication : sha1
Encryption : 3des-cbc
Pseudo random function: hmac-sha1
Traffic statistics:
Input bytes : 10249
Output bytes : 4249
Input packets: 10
Output packets: 9
Flags: Caller notification sent
IPsec security associations: 2 created, 1 deleted
Phase 2 negotiations in progress: 0
```

Meaning The output displays the details of the individual IKE SAs such as role (initiator or responder), status, exchange type, authentication method, encryption algorithms, traffic statistics, Phase 2 negotiation status, and so on.

Copyright © 2016, Juniper Networks, Inc.

You can use the output data to:

- Know the role of the IKE SA. Troubleshooting is easier when the peer has the responder role.
- Get the traffic statistics to verify the traffic flow in both directions.
- Get the number of IPsec security associations created or in progress.
- Get the status of any completed Phase 2 negotiations.

Confirming IPsec Phase 2 Status

Purpose View IPsec (Phase 2) security associations.

When IKE Phase 1 is confirmed, view the IPsec (Phase 2) security associations.

Action From operational mode, enter the **show security ipsec security-associations** command.

```
user@host> show security ipsec security-associations

total configured sa: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<2 10.2.2.2 500 ESP:3des/sha1 bce1c6e0 1676/ unlim - 0
>2 10.2.2.2 500 ESP:3des/sha1 1a24eab9 1676/ unlim - 0
```

Meaning The output indicates that:

- There is a configured IPsec SA pair available . The port number 500 indicates that a standard IKE port is used. Otherwise, it is Network Address Translation-Traversal (NAT-T), 4500, or random high port.
- The security parameter index (SPI) is used for both directions. The lifetime or usage limits of the SA is expressed either in seconds or in kilobytes. In the output, 1676/ unlim indicates Phase 2 lifetime is set to expire in 1676 seconds and there is no specified lifetime size.
- The ID number shows the unique index value for each IPsec SA.
- A hyphen (-) in the Mon column indicates that VPN monitoring is not enabled for this SA.
- The virtual system (vsys) is zero, which is the default value.

NOTE: Phase 2 lifetime can be different from the Phase 1 lifetime because Phase 2 is not dependent on Phase 1 after the VPN is up.

Displaying IPsec Security Association Details

Purpose Display the individual IPsec SA details identified by the index number.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security ipsec security-associations index 2 detail** command.

```
user@host> show security ipsec security-associations index 2 detail
Virtual-system: Root
Local Gateway: 10.1.1.2, Remote Gateway: 10.2.2.2
Local Identity: ipv4_subnet(any:0,[0..7]=192.168.10.0/24)
Remote Identity: ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
DF-bit: clear
Policy-name: tunnel-policy-out
Direction: inbound, SPI: bce1c6e0, AUX-SPI: 0
Hard lifetime: Expires in 1667 seconds
Lifesize Remaining: Unlimited
Soft lifetime: Expires in 1093 seconds
Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: enabled, Replay window size: 32
Direction: outbound, SPI: 1a24eab9, AUX-SPI: 0
Hard lifetime: Expires in 1667 seconds
Lifesize Remaining: Unlimited
Soft lifetime: Expires in 1093 seconds
Mode: tunnel, Type: dynamic, State: installed, VPN Monitoring: -
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
Anti-replay service: enabled, Replay window size: 32
```

Copyright © 2016, Juniper Networks, Inc.

Meaning The output displays the local Identity and the remote Identity.

Note that a proxy ID mismatch may cause Phase 2 completion to fail. The proxy ID is derived from the tunnel policy (for policy-based VPNs). The local address and remote address are derived from the address book entries, and the service is derived from the application configured for the policy.

If Phase 2 fails due to a proxy ID mismatch, verify which address book entries are configured in the policy and ensure that the correct addresses are sent. Also ensure that the ports are matching. Double-check the service to ensure that the ports match for the remote and local servers.

NOTE: If multiple objects are configured in a tunnel policy for source address, destination address, or application, then the resulting proxy ID for that parameter is changed to zeroes.

For example, assume the following scenario for a tunnel policy:

- Local addresses of 192.168.10.0/24 and 10.10.20.0/24
- Remote address of 192.168.168.0/24
- Application as junos-http

The resulting proxy ID is local 0.0.0.0/0, remote 192.168.168.0/24, service 80.

The resulting proxy IDs can affect the interoperability if the remote peer is not configured for the second subnet. Also if you are employing a third-party vendor's application, you may have to manually enter the proxy ID to match.

If IPsec fails to complete, then check the kmd log or use the set traceoptions command. For more information, see "Troubleshooting IKE, PKI, and IPsec Issues" on page 6398.

Checking IPsec SA Statistics

Purpose Check statistics and errors for an IPsec SA.

For troubleshooting purpose, check the Encapsulating Security Payload/Authentication Header (ESP/AH) counters for any errors with a particular IPsec SA.

Action From operational mode, enter the **show security ipsec statistics index 2** command.

```
user@host> show security ipsec statistics index 2
ESP Statistics:
Encrypted bytes: 674784
Decrypted bytes: 309276
Encrypted packets: 7029
Decrypted packets: 7029
AH Statistics:
Input bytes: 0
```

 Copyright © 2016, Juniper Networks, Inc.

```
Output bytes: 0
Input packets: 0
Output packets: 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```

Meaning An error value of zero in the output indicates a normal condition.

We recommend running this command multiple times to observe any packet loss issues across a VPN. Output from this command also displays the statistics for encrypted and decrypted packet counters, error counters, and so on.

You must enable security flow traceoptions to investigate which ESP packets are experiencing errors and why. For more information, see "Troubleshooting IKE, PKI, and IPsec Issues" on page 6398.

Testing Traffic Flow Across the VPN

Purpose Test traffic flow across the VPN after Phase 1 and Phase 2 have completed successfully. You can test traffic flow by using the **ping** command. You can ping from local host to remote host. You can also initiate pings from the Juniper Networks device itself.

This example shows how to initiate a ping request from the Juniper Networks device to the remote host. Note that when pings are initiated from the Juniper Networks device, the source interface must be specified to ensure that the correct route lookup takes place and the appropriate zones are referenced in the policy lookup.

In this example, the ge-0/0/0.0 interface resides in the same security zone as the local host and must be specified in the ping request so that the policy lookup can be from zone trust to zone untrust.

Action From operational mode, enter the **ping 192.168.168.10 interface ge-0/0/0 count 5** command.

```
user@host> ping 192.168.168.10 interface ge-0/0/0 count 5
PING 192.168.168.10 (192.168.168.10): 56 data bytes
64 bytes from 192.168.168.10: icmp_seq=0 ttl=127 time=8.287 ms
64 bytes from 192.168.168.10: icmp_seq=1 ttl=127 time=4.119 ms
64 bytes from 192.168.168.10: icmp_seq=2 ttl=127 time=5.399 ms
64 bytes from 192.168.168.10: icmp_seq=3 ttl=127 time=4.361 ms
64 bytes from 192.168.168.10: icmp_seq=4 ttl=127 time=5.137 ms
--- 192.168.168.10 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.119/5.461/8.287/1.490 ms
```

Confirming the Connectivity

Purpose Confirm the connectivity between a remote host and a local host.

Action From operational mode, enter the **ping 192.168.10.10 from ethernet0/6** command.

```
ssg5-> ping 192.168.10.10 from ethernet0/6
```

```
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 192.168.10.10, timeout is 1 seconds from
ethernet0/6
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/4/5 ms
```

Meaning You can confirm end-to-end connectivity by using the **ping** command from the remote host to the local host. In this example, the command is initiated from the SSG5 device.

Failed end-to-end connectivity can indicate an issue with routing, policy, end host, or encryption/decryption of the ESP packets. To verify the exact causes of the failure:

- Check IPsec statistics for details on errors as described in "Checking IPsec SA Statistics" on page 6396.
- Confirm end host connectivity by using the **ping** command from a host on the same subnet as the end host. If the end host is reachable by other hosts, then you can assume that the issue is not with the end host.
- Enable security flow traceoptions for troubleshooting the routing-related and policy-related issues.

Troubleshooting IKE, PKI, and IPsec Issues

Troubleshoot IKE, PKI, and IPsec issues.

- Basic Troubleshooting Steps on page 6398
- Checking the Free Disk Space on Your Device on page 6399
- Checking the Log Files to Verify Different Scenarios and Uploading Log Files to an FTP on page 6400
- Enabling IKE Traceoptions to View Messages on IKE on page 6400
- Enabling PKI Traceoptions to View Messages on IPsec on page 6401
- Setting up IKE and PKI Traceoptions to Troubleshoot IKE Setup Issues with Certificates on page 6402
- Analyzing the Phase 1 Success Message on page 6402
- Analyzing the Phase 1 Failure Message (Proposal Mismatch) on page 6403
- Analyzing the Phase 1 Failure Message (Authentication Failure) on page 6403
- Analyzing the Phase 1 Failure Message (Timeout Error) on page 6404
- Analyzing the Phase 2 Failure Message on page 6404
- Analyzing the Phase 2 Failure Message on page 6405
- Troubleshooting Common Problems Related to IKE and PKI on page 6406

Basic Troubleshooting Steps

Problem The basic troubleshooting steps are as follows:

1. Identifying and isolating the problem.

Copyright © 2016, Juniper Networks, Inc.

2. Debugging the problem.

The common approach of starting troubleshooting is with the lowest layer of the OSI layers and working your way up the OSI stack to confirm the layer in which the failure occurs.

Solution Basic steps for troubleshooting IKE, PKI, and IPsec are as follows:

- Confirm the physical connectivity of the Internet link at the physical and data link levels.
- Confirm that the Juniper Networks device has connectivity to the Internet next hop and connectivity to the remote IKE peer.
- Confirm IKE Phase 1 completion.
- Confirm IKE Phase 2 completion if IKE Phase 1 completion is successful.
- Confirm the traffic flow across the VPN (if the VPN is up and active).

Junos OS includes the traceoptions feature. Using this feature, you can enable a traceoption flag to write the data from the traceoption to a log file, which may be predetermined or manually configured and stored in flash memory. These trace logs can be retained even after a system reboot. Check the available flash storage before implementing traceoptions.

You can enable the traceoptions feature in configuration mode and commit the configuration to use the traceoptions feature. Similarly to disable traceoptions, you must deactivate traceoptions in configuration mode and commit the configuration.

Checking the Free Disk Space on Your Device

Problem Check the statistics on the free disk space in your device file systems.

Solution From operational mode, enter the **show system storage** command.

```
user@host> show system storage
Filesystem Size Used Avail Capacity Mounted on
/dev/ad0s1a 213M 74M 137M 35% /
devfs 1.0K 1.0K 0B 100% /dev
devfs 1.0K 1.0K 0B 100% /dev/
/dev/md0 180M 180M 0B 100% /junos
/cf 213M 74M 137M 35% /junos/cf
devfs 1.0K 1.0K 0B 100% /junos/dev/
procfs 4.0K 4.0K 0B 100% /proc
/dev/bo0s1e 24M 13K 24M 0% /config
/dev/md1 168M 7.6M 147M 5% /mfs
/cf/var/jail 213M 74M 137M 35% /jail/var
```

The **/dev/ad0s1a** represents the onboard flash memory and is currently at 35 percent capacity.

Copyright © 2016, Juniper Networks, Inc.

Checking the Log Files to Verify Different Scenarios and Uploading Log Files to an FTP

Problem View the log files to check security IKE debug messages, security flow debugs, and the state of logging to the syslog.

Solution From operational mode, enter the **show log kmd**, **show log pkid**, **show log security-trace**, and **show log messages** commands.

```
user@host> show log kmd
user@host> show log pkid
user@host> show log security-trace
user@host> show log messages
```

NOTE: You can view a list of all logs in the /var/log directory by using the **show log** command.

Log files can also be uploaded to an FTP server by using the **file copy** command.

```
(operational mode):
user@host> file copy path/filename dest-path/filename
Example:
```

```
user@host> file copy /var/log/kmd ftp://192.168.10.10/kmd.log
ftp://192.168.10.10/kmd.log 100% of 35 kB 12 MBps
```

Enabling IKE Traceoptions to View Messages on IKE

Problem To view success or failure messages for IKE or IPsec, you can view the kmd log by using the **show log kmd** command. Because the kmd log displays some general messages, it can be useful to obtain additional details by enabling IKE and PKI traceoptions.

NOTE: Generally, it is best practice to troubleshoot the peer that has the responder role. You must obtain the trace output from the initiator and responder to understand the cause of a failure.

Configure IKE tracing options.

Solution

```
user@host> configure
Entering configuration mode

[edit]
user@host# edit security ike traceoptions
[edit security ike traceoptions]
```

```
user@host# set file ?
Possible completions:
<filename> Name of file in which to write trace information
```


Copyright © 2016, Juniper Networks, Inc.

```
files Maximum number of trace files (2..1000)
match Regular expression for lines to be logged
no-world-readable Don't allow any user to read the log file
size Maximum trace file size (10240..1073741824)
world-readable Allow any user to read the log file
```

[edit security ike traceoptions]

```
user@host# set flag ?
Possible completions:
all Trace everything
certificates Trace certificate events
database Trace security associations database events
general Trace general events
ike Trace IKE module processing
parse Trace configuration processing
policy-manager Trace policy manager processing
routing-socket Trace routing socket messages
timer Trace internal timer events
```

NOTE: If you do not specify file names for the <filename> field, then all IKE traceoptions are written to the kmd log.

You must specify at least one flag option to write trace data to the log. For example:

- **file size** — Maximum size of each trace file, in bytes. For example, 1 million (1,000,000) can generate a maximum file size of 1 MB.
- **files** — Maximum number of trace files to be generated and stored in a flash memory device.

NOTE: You must commit your configuration to start the trace.

Enabling PKI Traceoptions to View Messages on IPsec

Problem Enable PKI traceoptions to identify whether an IKE failure is related to the certificate or to a non-PKI issue.

Solution [edit security pki traceoptions]

```
user@host# set file ?
Possible completions:
<filename> Name of file in which to write trace information
files Maximum number of trace files (2..1000)
match Regular expression for lines to be logged
no-world-readable Don't allow any user to read the log file
size Maximum trace file size (10240..1073741824)
world-readable Allow any user to read the log file
```

Copyright © 2016, Juniper Networks, Inc.

[edit security pki traceoptions]

```
user@host# set flag ?
Possible completions:
all Trace with all flags enabled
certificate-verification PKI certificate verification tracing
online-crl-check PKI online crl tracing
```

Setting up IKE and PKI Traceoptions to Troubleshoot IKE Setup Issues with Certificates

Problem Configure the recommended settings for IKE and PKI traceoptions.

NOTE: The IKE and PKI traceoptions use the same parameters, but the default filename for all PKI-related traces is found in the pkid log.

Solution

user@host> **configure**
Entering configuration mode

[edit security ike traceoptions]
user@host# **set file size 1m**
user@host# **set flag ike**
user@host# **set flag policy-manager**
user@host# **set flag routing-socket**
user@host# **set flag certificates**

[edit security pki traceoptions]
user@host# **set file size 1m**
user@host# **set flag all**
user@host# **commit and-quit**
commit complete
Exiting configuration mode

Analyzing the Phase 1 Success Message

Problem Understand the output of the **show log kmd** command when the IKE Phase 1 and Phase 2 conditions are successful.

Copyright © 2016, Juniper Networks, Inc.

Solution

```
Nov 7 11:52:14 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=
10.1.1.2) remote=fqdn(udp:500,[0..15]=ssg5.example.net)
Nov 7 11:52:14 Phase-2 [responder] done for
p1_local=ipv4(udp:500,[0..3]=10.1.1.2)
p1_remote=fqdn(udp:500,[0..15]=ssg5.example.net)
p2_local=ipv4_subnet(any:0,[0..7]=192.168.10.0/24)
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
```

The sample output indicates:

- **10.1.1.2**—Local address.
- **ssg5.example.net** —Remote peer (hostname with FQDN).
- **udp: 500**—NAT-T was not negotiated.
- **Phase 1 [responder] done**—Phase 1 status, along with the role (initiator or responder).
- **Phase 2 [responder] done**—Phase 1 status, along with the proxy ID information.

  You can also confirm the IPsec SA status by using the verification commands mentioned in "Confirming IKE Phase 1 Status" on page 6392.

Analyzing the Phase 1 Failure Message (Proposal Mismatch)

Problem Understanding the output of the **show log kmd** command, where the IKE Phase 1 condition is a failure, helps in determining the reason for the VPN not establishing Phase 1.

Solution

```
Nov 7 11:52:14 Phase-1 [responder] failed with error(No proposal chosen) for
local=unknown(any:0,[0..0]=) remote=fqdn(udp:500,[0..15]=ssg5.example.net)
Nov 7 11:52:14 10.1.1.2:500 (Responder) <-> 10.2.2.2:500 { 011359c9 ddef501d -
2216ed2a bfc50f5f [-
1] / 0x00000000 } IP; Error = No proposal chosen (14)
```

The sample output indicates:

- **10.1.1.2**—Local address.
- **ssg5.example.net** —Remote peer (hostname with FQDN).
- **udp: 500**—NAT-T was not negotiated.
- **Phase-1 [responder] failed with error (No proposal chosen)**—Phase 1 failure because of proposal mismatch.

To resolve this issue, ensure that the parameters for the IKE gateway Phase 1 proposals on both the responder and the initiator match. Also confirm that a tunnel policy exists for the VPN.

Analyzing the Phase 1 Failure Message (Authentication Failure)

Problem Understand the output of the **show log kmd** command when the IKE Phase 1 condition is a failure. This helps in determining the reason for the VPN not establishing Phase 1.

Copyright © 2016, Juniper Networks, Inc.

Solution

```
Nov 7 12:06:36 Unable to find phase-1 policy as remote peer:10.2.2.2 is not
recognized.
Nov 7 12:06:36 Phase-1 [responder] failed with error(Authentication failed) for
local=ipv4(udp:500,[0..3]=10.1.1.2) remote=ipv4(any:0,[0..3]=10.2.2.2)
Nov 7 12:06:36 10.1.1.2:500 (Responder) <-> 10.2.2.2:500 { f725ca38 dad47583 -
dab1ba4c ae26674b [-
1] / 0x00000000 } IP; Error = Authentication failed (24)
```

The sample output indicates:

- **10.1.1.2**—Local address.
- **10.2.2.2**—Remote peer
- **Phase 1 [responder] failed with error (Authentication failed)**—Phase 1 failure due to the responder not recognizing the incoming request originating from a valid gateway peer. In the case of IKE with PKI certificates, this failure typically indicates that an incorrect IKE ID type was specified or entered.

To resolve this issue, confirm that the correct peer IKE ID type is specified on the local peer based on the following:

- How the remote peer certificate was generated
- Subject Alternative Name or DN information in the received remote peer certificate

Analyzing the Phase 1 Failure Message (Timeout Error)

Problem Understand the output of the **show log kmd** command when the IKE Phase 1 condition is a failure.

Solution

```
Nov 7 13:52:39 Phase-1 [responder] failed with error(Timeout) for
local=unknown(any:0,[0..0]=)
remote=ipv4(any:0,[0..3]=10.2.2.2)
```

The sample output indicates:

- **10.1.1.2**—Llocal address.
- **10.2.2.2**—Remote peer.
- **Phase 1 [responder] failed with error(Timeout)**—Phase 1 failure.

  This error indicates that either the IKE packet is lost enroute to the remote peer or there is a delay or no response from the remote peer.

Because this timeout error is the result of waiting on a response from the PKI daemon, you must review the PKI traceoptions output to see whether there is a problem with PKI.

Analyzing the Phase 2 Failure Message

Problem Understand the output of the **show log kmd** command when the IKE Phase 2 condition is a failure.

Copyright © 2016, Juniper Networks, Inc.

Solution

```
Nov 7 11:52:14 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=
10.1.1.2) remote=fqdn(udp:500,[0..15]=ssg5.example.net)
Nov 7 11:52:14 Failed to match the peer proxy ids
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
p2_local=ipv4_subnet(any:0,[0..7]=10.10.20.0/24) for the remote
peer:ipv4(udp:500,[0..3]=10.2.2.2)
Nov 7 11:52:14 KMD_PM_P2_POLICY_LOOKUP_FAILURE: Policy lookup for Phase-2
[responder] failed for
p1_local=ipv4(udp:500,[0..3]=10.1.1.2) p1_remote=ipv4(udp:500,[0..3]=10.2.2.2)
p2_local=ipv4_subnet(any:0,[0..7]=10.10.20.0/24)
p2_remote=ipv4_subnet(any:0,[0..7]=192.168.168.0/24)
Nov 7 11:52:14 10.1.1.2:500 (Responder) <-> 10.2.2.2:500 { 41f638eb cc22bbfe -
43fd0e85 b4f619d5 [0]
/ 0xc77fafcf } QM; Error = No proposal chosen (14)
```

The sample output indicates:

- **10.1.1.2**—Local address.
- **ssg5.example.net** —Remote peer (IKE ID type hostname with FQDN).
- **Phase 1 [responder] done**—Phase 1 success.
- **Failed to match the peer proxy ids**—The Incorrect proxy IDs are received. In the previous sample, the two proxy IDs received are 192.168.168.0/24 (remote) and 10.10.20.0/24 (local) (for service=any). Based on the configuration given in this example, the expected local address is 192.168.10.0/24. This shows that there is a mismatch of configurations on the local peer, resulting in the failure of proxy ID match.

  To resolve this issue, correct the address book entry or configure the proxy ID on either peer so that it matches the other peer.

  The output also indicates the reason for failure is **No proposal chosen**. However in this case you also see the message **Failed to match the peer proxy ids**.

Analyzing the Phase 2 Failure Message

Problem Understand the output of the **show log kmd** command when the IKE Phase 2 condition is a failure.

Copyright © 2016, Juniper Networks, Inc.

Solution

```
Nov 7 11:52:14 Phase-1 [responder] done for local=ipv4(udp:500,[0..3]=
10.1.1.2) remote=fqdn(udp:500,[0..15]=ssg5.example.net)
Nov 7 11:52:14 10.1.1.2:500 (Responder) <-> 10.2.2.2:500 { cd9dff36 4888d398 -
6b0d3933 f0bc8e26 [0]
/ 0x1747248b } QM; Error = No proposal chosen (14)
```

The sample output indicates:

- **10.1.1.2** —Local address.
- **fqdn(udp:500,[0..15]=ssg5.example.net**—Remote peer.
- **Phase 1 [responder] done**—Phase 1 success.
- **Error = No proposal chosen**—No proposal was chosen during Phase 2. This issue is due to proposal mismatch between the two peers.

  To resolve this issue, confirm that the Phase 2 proposals match on both peers.

Troubleshooting Common Problems Related to IKE and PKI

Problem Troubleshoot common problems related to IKE and PKI.

Enabling the traceoptions feature helps you to gather more information on the debugging issues than is obtainable from the normal log entries. You can use the traceoptions log to understand the reasons for IKE or PKI failures.

Solution Methods for troubleshooting the IKE -and-PKI-related issues:

- Ensure that the clock, date, time zone, and daylight savings settings are correct. Use NTP to keep the clock accurate.
- Ensure that you use a two-letter country code in the "C=" (country) field of the DN.

  For example: use “US” and not “USA” or “United States.” Some CAs require that the country field of the DN be populated, allowing you to enter the country code value only with a two-letter value.
- Ensure that if a peer certificate is using multiple OU=or CN= fields, you are using the distinguished name with container method (the sequence must be maintained and is case- sensitive).
- If the certificate is not valid yet, check the system clock and, if required, adjust the system time zone or just add a day in the clock for a quick test.
- Ensure that a matching IKE ID type and value are configured.
- PKI can fail due to a revocation check failure. To confirm this, temporarily disable revocation checking and see whether IKE Phase 1 is able to complete.

  To disable revocation checking, use the following command in configure mode:

  set security pki ca-profile <ca-profile> revocation-check disable

Related Documentation

- IPsec VPN Overview on page 6041

 Copyright © 2016, Juniper Networks, Inc.

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 269

Managing Certificate Revocation

- Understanding Online Certificate Status Protocol on page 6409
- Understanding Certificate Revocation Lists on page 6410
- Comparison of Online Certificate Status Protocol and Certificate Revocation List on page 6411
- Improving Security by Configuring OCSP for Certificate Revocation Status on page 6412
- Example: Manually Loading a CRL onto the Device on page 6427
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428
- Example: Verifying Certificate Validity on page 6429
- Deleting a Loaded CRL (CLI Procedure) on page 6431

Understanding Online Certificate Status Protocol

OCSP is used to check the revocation status of X509 certificates. OCSP provides revocation status on certificates in real time and is useful in time-sensitive situations such as bank transactions and stock trades.

The revocation status of a certificate is checked by sending a request to an OCSP server that resides outside of an SRX Series device. Based on the response from the server, the VPN connection is allowed or denied. OCSP responses are not cached on SRX Series devices.

The OCSP server can be the certificate authority (CA) that issues a certificate or a designated authorized responder. The location of the OCSP server can be configured manually or extracted from the certificate that is being verified. Requests are sent first to OCSP server locations that are manually configured in CA profiles with the **ocsp url** statement at the [**edit security pki ca-profile** *profile-name* **revocation-check**] hierarchy level; up to two locations can be configured for each CA profile. If the first configured OCSP server is not reachable, the request is sent to the second OCSP server. If the second OCSP server is not reachable, the request is then sent to the location in the certificate's AuthorityInfoAccess extension field. The **use-ocsp** option must also be configured, as certificate revocation list (CRL) is the default checking method.

Copyright © 2016, Juniper Networks, Inc.

SRX Series devices accept only signed OCSP responses from the CA or authorized responder. The response received is validated using trusted certificates. The response is validated as follows:

1. The CA certificate enrolled for the configured CA profile is used to validate the response.
2. The OCSP response might contain a certificate to validate the OCSP response. The received certificate must be signed by a CA certificate enrolled in the SRX Series device. After the received certificate is validated by the CA certificate, it is used to validate the OCSP response.

The response from the OCSP server can be signed by different CAs. The following scenarios are supported:

- The CA server that issues the end entity certificate for a device also signs the OCSP revocation status response. The SRX Series device verifies the OCSP response signature using the CA certificate enrolled in the SRX Series device. After the OCSP response is validated, the certificate revocation status is checked.
- An authorized responder signs the OCSP revocation status response. The certificate for the authorized responder and the end entity certificate being verified must be issued by the same CA. The authorized responder is first verified using the CA certificate enrolled in the SRX Series device. The OCSP response is validated using the responder's CA certificate. The SRX Series device then uses the OCSP response to check the revocation status of the end entity certificate.
- There are different CA signers for the end entity certificate being verified and the OCSP response. The OCSP response is signed by a CA in the certificate chain for the end entity certificate being verified. (All peers participating in an IKE negotiation need to have at least one common trusted CA in their respective certificate chains.) The OCSP responder's CA is verified using a CA in the certificate chain. After validating the responder CA certificate, the OCSP response is validated using the responder's CA certificate.

To prevent replay attacks, a nonce payload can be sent in an OCSP request. Nonce payloads are sent by default unless it is explicitly disabled. If enabled, the SRX Series device expects the OCSP response to contain a nonce payload, otherwise the revocation check fails. If OCSP responders are not capable of responding with a nonce payload, then the nonce payload must be disabled on the SRX Series device.

Related Documentation

- Comparison of Online Certificate Status Protocol and Certificate Revocation List on page 6411
- Improving Security by Configuring OCSP for Certificate Revocation Status on page 6412

Understanding Certificate Revocation Lists

In the normal course of business, certificates are revoked for various reasons. You might wish to revoke a certificate if you suspect that it has been compromised, for example, or when a certificate holder leaves the company.

Copyright © 2016, Juniper Networks, Inc.

You can manage certificate revocations and validations in two ways:

- Locally— This is a limited solution.
- By referencing a Certificate Authority (CA) certificate revocation list (CRL)— You can automatically access the CRL online at intervals you specify or at the default interval set by the CA.

In Phase 1 negotiations, participants check the CRL list to see if certificates received during an IKE exchange are still valid. If a CRL did not accompany a CA certificate and is not loaded on the device, the device tries to download it automatically from the CRL distribution point of the local certificate. If the device fails to connect to the URL in the certificate distribution point (CDP), it tries to retrieve the CRL from the URL configured in the CA profile.

If the certificate does not contain a certificate distribution point extension, and you cannot automatically retrieve the CRL through Lightweight Directory Access Protocol (LDAP) or Hypertext Transfer Protocol (HTTP), you can retrieve a CRL manually and load that in the device.

Related Documentation

- Understanding Certificates and PKI on page 6347
- Example: Manually Loading a CRL onto the Device on page 6427
- Example: Verifying Certificate Validity on page 6429
- Deleting a Loaded CRL (CLI Procedure) on page 6431
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428

Comparison of Online Certificate Status Protocol and Certificate Revocation List

Online Certificate Status Protocol (OCSP) and certificate revocation list (CRL) can both be used to check the revocation status of a certificate. There are advantages and disadvantages to each method.

- OCSP provides certificate status in real time, while CRL uses cached data. For time-sensitive applications, OCSP is the preferred approach.
- CRL checking is faster because lookup for certificate status is done on information cached on the VPN device. OCSP requires time to obtain the revocation status from an external server.
- CRL requires additional memory to store the revocation list received from a CRL server. OCSP does not require additional memory to save the revocation status of certificates.
- OCSP requires that the OCSP server be available at all times. CRL can use cached data to check the revocation status of certificates when the server is unreachable.

NOTE: On SRX Series devices, CRL is the default method used to check the revocation status of a certificate.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Online Certificate Status Protocol on page 6409
- Understanding Certificate Revocation Lists on page 6410

Improving Security by Configuring OCSP for Certificate Revocation Status

This example shows how to improve security by configuring two peers using the Online Certificate Status Protocol (OCSP) to check the revocation status of the certificates used in Phase 1 negotiations for the IPsec VPN tunnel.

- Requirements on page 6412
- Overview on page 6412
- Configuration on page 6414
- Verification on page 6422

Requirements

On each device:

- Obtain and enroll a local certificate. This can be done either manually or by using the Simple Certificate Enrollment Protocol (SCEP).
- Optionally, enable automatic renewal of the local certificate.
- Configure security policies to permit traffic to and from the peer device.

Overview

On both peers, a certificate authority (CA) profile OCSP-ROOT is configured with the following options:

- CA name is OCSP-ROOT.
- Enrollment URL is http://10.1.1.1:8080/scep/OCSP-ROOT/. This is the URL where SCEP requests to the CA are sent.
- The URL for the OCSP server is http://10.157.88.56:8210/OCSP-ROOT/.
- OCSP is used first to check the certificate revocation status. If there is no response from the OCSP server, then the certificate revocation list (CRL) is used to check the status. The CRL URL is http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45.
- The CA certificate received in an OCSP response is not checked for certificate revocation. Certificates received in an OCSP response generally have shorter lifetimes and a revocation check is not required.

Table 551 shows the Phase 1 options used in this example.

Table 551: Phase 1 Options for OCSP Configuration Example

Option	Peer A	Peer B
IKE proposal	ike_prop	ike_prop

Copyright © 2016, Juniper Networks, Inc.

Table 551: Phase 1 Options for OCSP Configuration Example *(continued)*

Option	Peer A	Peer B
Authentication method	RSA signatures	RSA signatures
DH group	group2	group2
Authentication algorithm	SHA 1	SHA 1
Encryption algorithm	3DES CBC	3DES CBC
IKE policy	ike_policy	ike_policy
Mode	aggressive	aggressive
Proposal	ike_prop	ike_prop
Certificate	local-certificate localcert1	local-certificate localcert1
IKE gateway	jsr_gateway	jsr_gateway
Policy	ike_policy	ike_policy
Gateway address	198.51.100.50	192.0.2.50
Remote identity	localcert11.example.net	-
Local identity	-	localcert11.example.net
External interface	reth1	ge-0/0/2.0
Version	v2	v2

Table 552 shows the Phase 2 options used in this example.

Table 552: Phase 2 Options for OCSP Configuration Example

Option	Peer A	Peer B
IPsec proposal	ipsec_prop	ipsec_prop
Protocol	ESP	ESP
Authentication algorithm	HMAC SHA1-96	HMAC SHA1-96
Encryption algorithm	3DES CBC	3DES CBC
Lifetime seconds	1200	1200
Lifetime kilobytes	150,000	150,000

Copyright © 2016, Juniper Networks, Inc.

Table 552: Phase 2 Options for OCSP Configuration Example *(continued)*

Option	Peer A	Peer B
IPsec policy	ipsec_policy	ipsec_policy
PFC keys	group2	group2
Proposal	ipsec_prop	ipsec_prop
VPN	test_vpn	test_vpn
Bind interface	st0.1	st0.1
IKE gateway	jsr_gateway	jsr_gateway
Policy	ipsec_policy	ipsec_policy
Establish tunnels	-	immediately

Topology

Figure 280 shows the peer devices that are configured in this example.

Figure 280: OCSP Configuration Example

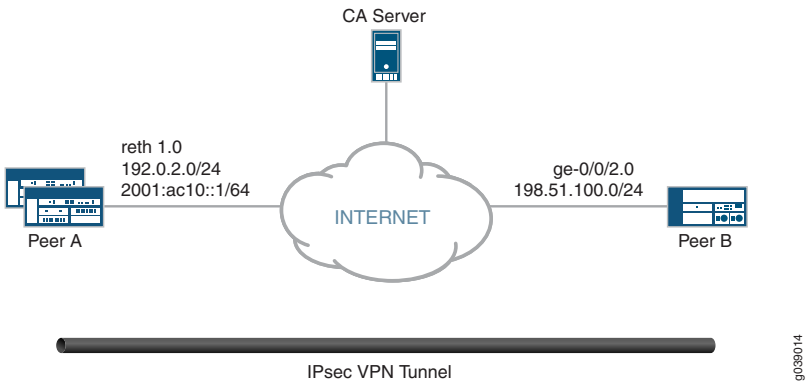

Configuration

- Configuring Peer A on page 6414
- Configuring Peer B on page 6418

Configuring Peer A

CLI Quick Configuration

To quickly configure VPN peer A to use OCSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/3 gigether-options redundant-parent reth1
set interfaces ge-9/0/3 gigether-options redundant-parent reth1
set interfaces lo0 unit 0 family inet address 172.16.1.100/24
set interfaces lo0 redundant-pseudo-interface-options redundancy-group 1
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 192.0.2.50/24
set interfaces st0 unit 1 family inet address 172.18.1.100/24
set security pki ca-profile OCSP-ROOT ca-identity OCSP-ROOT
set security pki ca-profile OCSP-ROOT enrollment url
  http://10.1.1.1:8080/scep/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check ocsp url
  http://10.157.88.56:8210/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check use-ocsp
set security pki ca-profile OCSP-ROOT revocation-check ocsp
  disable-responder-revocation-check
set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure
  fallback-crl
set security pki ca-profile OCSP-ROOT revocation-check crl url
  http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
set security ike proposal ike_prop authentication-method rsa-signatures
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_policy mode aggressive
set security ike policy ike_policy proposals ike_prop
set security ike policy ike_policy certificate local-certificate localcert1
set security ike gateway jsr_gateway ike-policy ike_policy
set security ike gateway jsr_gateway address 198.51.100.50
set security ike gateway jsr_gateway remote-identity hostname localcert11.example.net
set security ike gateway jsr_gateway external-interface reth1
set security ike gateway jsr_gateway version v2-only
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec proposal ipsec_prop lifetime-seconds 1200
set security ipsec proposal ipsec_prop lifetime-kilobytes 150000
set security ipsec policy ipsec_policy perfect-forward-secrecy keys group2
set security ipsec policy ipsec_policy proposals ipsec_prop
set security ipsec vpn test_vpn bind-interface st0.1
set security ipsec vpn test_vpn ike gateway jsr_gateway
set security ipsec vpn test_vpn ike ipsec-policy ipsec_policy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure VPN peer A to use OCSP:

1. Configure interfaces.

```
[edit interfaces]
set ge-0/0/3 gigether-options redundant-parent reth1
set ge-9/0/3 gigether-options redundant-parent reth1
set lo0 unit 0 family inet address 172.16.1.100/24
set lo0 redundant-pseudo-interface-options redundancy-group 1
set reth1 redundant-ether-options redundancy-group 1
```

Copyright © 2016, Juniper Networks, Inc.

```
set reth1 unit 0 family inet address 192.0.2.0/24
set st0 unit 1 family inet address 172.18.1.100/24
```

2. Configure the CA profile.

```
[edit security pki ca-profile OCSP-ROOT]
set ca-identity OCSP-ROOT
set enrollment url http://10.1.1.1:8080/scep/OCSP-ROOT/
set revocation-check ocsp url http://10.157.88.56:8210/OCSP-ROOT/
set revocation-check use-ocsp
set revocation-check ocsp disable-responder-revocation-check
set revocation-check ocsp connection-failure fallback-crl
set revocation-check crl url http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike_prop]
set authentication-method rsa-signatures
set dh-group group2
set authentication-algorithm sha1
set encryption-algorithm 3des-cbc

[edit security ike policy ike_policy]
set mode aggressive
set proposals ike_prop
set certificate local-certificate localcert1

[edit security ike gateway jsr_gateway]
set ike-policy ike_policy
set address 198.51.100.50
set remote-identity hostname localcert11.example.net
set external-interface reth1
set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_prop]
set protocol esp
set authentication-algorithm hmac-sha1-96
set encryption-algorithm 3des-cbc
set lifetime-seconds 1200
set lifetime-kilobytes 150000

[edit security ipsec policy ipsec_policy]
set perfect-forward-secrecy keys group2
set proposals ipsec_prop

[edit security ipsec vpn test_vpn]
set bind-interface st0.1
set ike gateway jsr_gateway
set ike ipsec-policy ipsec_policy
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security pki ca-profile OCSP-ROOT**, **show security ike**, and **show security ipsec**

Copyright © 2016, Juniper Networks, Inc.

commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/3 {
    gigether-options {
        redundant-parent reth1;
    }
}
ge-9/0/3 {
    gigether-options {
        redundant-parent reth1;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.1.100/24;
        }
    }
    redundant-pseudo-interface-options {
        redundancy-group 1;
    }
}
reth1 {
    redundant-ether-options {
        redundancy-group 1;
    }
    unit 0 {
        family inet {
            address 192.0.2.0/24;
        }
    }
}
st0 {
    unit 1 {
        family inet {
            address 172.18.1.100/24;
        }
    }
}
[edit]
user@host# show security pki ca-profile OCSP-ROOT
ca-identity OCSP-ROOT;
enrollment {
    url http://10.1.1.1:8080/scep/OCSP-ROOT/;
}
revocation-check {
    crl {
        url http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45;
    }
    ocsp {
        disable-responder-revocation-check;
        url http://10.157.88.56:8210/OCSP-ROOT/;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    use-ocsp;
}
[edit]
user@host# show security ike
proposal ike_prop {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
}
policy ike_policy {
    mode aggressive;
    proposals ike_prop;
    certificate {
        local-certificate localcert1;
    }
}
gateway jsr_gateway {
    ike-policy ike_policy;
    address 10.10.2.50;
    remote-identity hostname localcert11.example.net;
    external-interface reth1;
    version v2-only;
}
[edit]
user@host# show security ipsec
proposal ipsec_prop {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 1200;
    lifetime-kilobytes 150000;
}
policy ipsec_policy {
    perfect-forward-secrecy {
        keys group2;
    }
    proposals ipsec_prop;
}
vpn test_vpn {
    bind-interface st0.1;
    ike {
        gateway jsr_gateway;
        ipsec-policy ipsec_policy;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Peer B

CLI Quick Configuration

To quickly configure VPN peer B to use OCSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [**edit**] hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/2 unit 0 family inet address 198.51.100.0/24
set interfaces lo0 unit 0 family inet address 172.17.1.100/24
set interfaces st0 unit 1 family inet address 172.18.1.1/24
set security pki ca-profile OCSP-ROOT ca-identity OCSP-ROOT
set security pki ca-profile OCSP-ROOT enrollment url
  http://10.1.1.1:8080/scep/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check ocsp url
  http://10.157.88.56:8210/OCSP-ROOT/
set security pki ca-profile OCSP-ROOT revocation-check use-ocsp
set security pki ca-profile OCSP-ROOT revocation-check ocsp
  disable-responder-revocation-check
set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure
  fallback-crl
set security pki ca-profile OCSP-ROOT revocation-check crl url
  http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
set security ike proposal ike_prop authentication-method rsa-signatures
set security ike proposal ike_prop dh-group group2
set security ike proposal ike_prop authentication-algorithm sha1
set security ike proposal ike_prop encryption-algorithm 3des-cbc
set security ike policy ike_policy mode aggressive
set security ike policy ike_policy proposals ike_prop
set security ike policy ike_policy certificate local-certificate localcert11
set security ike gateway jsr_gateway ike-policy ike_policy
set security ike gateway jsr_gateway address 192.0.2.50
set security ike gateway jsr_gateway local-identity hostname localcert11.example.net
set security ike gateway jsr_gateway external-interface ge-0/0/2.0
set security ike gateway jsr_gateway version v2-only
set security ipsec proposal ipsec_prop protocol esp
set security ipsec proposal ipsec_prop authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop encryption-algorithm 3des-cbc
set security ipsec proposal ipsec_prop lifetime-seconds 1200
set security ipsec proposal ipsec_prop lifetime-kilobytes 150000
set security ipsec policy ipsec_policy perfect-forward-secrecy keys group2
set security ipsec policy ipsec_policy proposals ipsec_prop
set security ipsec vpn test_vpn bind-interface st0.1
set security ipsec vpn test_vpn ike gateway jsr_gateway
set security ipsec vpn test_vpn ike ipsec-policy ipsec_policy
set security ipsec vpn test_vpn establish-tunnels immediately
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure VPN peer B to use OCSP:

1. Configure interfaces.

   ```
   [edit interfaces]
   set ge-0/0/2 unit 0 family inet address 198.51.100.0/24
   set lo0 unit 0 family inet address 172.17.1.100/24
   set st0 unit 1 family inet address 172.18.1.1/24
   ```

2. Configure the CA profile.

   ```
   [edit security pki ca-profile OCSP-ROOT]
   set ca-identity OCSP-ROOT
   ```

Copyright © 2016, Juniper Networks, Inc.

```
set enrollment url http://10.1.1.1:8080/scep/OCSP-ROOT/
set revocation-check ocsp url http://10.157.88.56:8210/OCSP-ROOT/
set revocation-check use-ocsp
set revocation-check ocsp disable-responder-revocation-check
set revocation-check ocsp connection-failure fallback-crl
set revocation-check crl url http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike_prop]
set authentication-method rsa-signatures
set dh-group group2
set authentication-algorithm sha1
set encryption-algorithm 3des-cbc

[edit security ike policy ike_policy]
set mode aggressive
set proposals ike_prop
set certificate local-certificate localcert1

[edit security ike gateway jsr_gateway]
set ike-policy ike_policy
set address 192.0.2.50
set local-identity hostname localcert11.example.net
set external-interface ge-0/0/2.0
set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_prop]
set protocol esp
set authentication-algorithm hmac-sha1-96
set encryption-algorithm 3des-cbc
set lifetime-seconds 1200
set lifetime-kilobytes 150000

[edit security ipsec policy ipsec_policy]
set perfect-forward-secrecy keys group2
set proposals ipsec_prop

[edit security ipsec vpn test_vpn]
set bind-interface st0.1
set ike gateway jsr_gateway
set ike ipsec-policy ipsec_policy
set establish-tunnels immediately
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security pki ca-profile OCSP-ROOT**, **show security ike**, and **show security ipsec** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
   unit 0 {
```

 Copyright © 2016, Juniper Networks, Inc.

```
      family inet {
        address 198.51.100.0/24;
      }
    }
}
lo0 {
    unit 0 {
      family inet {
        address 172.17.1.100/24;
      }
    }
}
st0 {
    unit 1 {
      family inet {
        address 172.18.1.1/24;
      }
    }
}
[edit]
user@host# show security pki ca-profile OCSP-ROOT
ca-identity OCSP-ROOT;
enrollment {
    url http://10.1.1.1:8080/scep/OCSP-ROOT/;
}
revocation-check {
    crl {
      url http://10.1.1.1:8080/crl-as-der/currentcrl-45.crlid=45;
    }
    ocsp {
      disable-responder-revocation-check;
      url http://10.157.88.56:8210/OCSP-ROOT/;
    }
    use-ocsp;
}
[edit]
user@host# show security ike
proposal ike_prop {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm 3des-cbc;
}
policy ike_policy {
    mode aggressive;
    proposals ike_prop;
    certificate {
      local-certificate localcert11;
    }
}
gateway jsr_gateway {
    ike-policy ike_policy;
    address 192.0.2.50;
    local-identity hostname localcert11.example.net;
    external-interface ge-0/0/2.0;
    version v2-only;
```

Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show security ipsec
proposal ipsec_prop {
   protocol esp;
   authentication-algorithm hmac-sha1-96;
   encryption-algorithm 3des-cbc;
   lifetime-seconds 1200;
   lifetime-kilobytes 150000;
}
policy ipsec_policy {
   perfect-forward-secrecy {
      keys group2;
   }
   proposals ipsec_prop;
}
vpn test_vpn {
   bind-interface st0.1;
   ike {
      gateway jsr_gateway;
      ipsec-policy ipsec_policy;
   }
   establish-tunnels immediately;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying CA Certificates on page 6422
- Verifying Local Certificates on page 6424
- Verifying IKE Phase 1 Status on page 6425
- Verifying IPsec Phase 2 Status on page 6426

Verifying CA Certificates

Purpose Verify the validity of a CA certificate on each peer device.

Action From operational mode, enter the **show security pki ca-certificate ca-profile OCSP-ROOT** or **show security pki ca-certificate ca-profile OCSP-ROOT detail** command.

```
user@host> show security pki ca-certificate ca-profile OCSP-ROOT
  Certificate identifier: OCSP-ROOT
    Issued to: OCSP-ROOT, Issued by: C = US, O = example, CN = OCSP-ROOT
    Validity:
      Not before: 11-15-2013 22:26 UTC
      Not after: 11-14-2016 22:26 UTC
    Public key algorithm: rsaEncryption(2048 bits)
```

```
user@host> show security pki ca-certificate ca-profile OCSP-ROOT detail
```


Copyright © 2016, Juniper Networks, Inc.

```
Certificate identifier: OCSP-ROOT
  Certificate version: 3
  Serial number: 0000a17f
  Issuer:
    Organization: example, Country: US, Common name: OCSP-ROOT
  Subject:
    Organization: example, Country: US, Common name: OCSP-ROOT
  Subject string:
    C=US, O=example, CN=OCSP-ROOT
  Validity:
    Not before: 11-15-2013 22:26 UTC
    Not after: 11-14-2016 22:26 UTC
  Public key algorithm: rsaEncryption(2048 bits)
    30:82:01:0a:02:82:01:01:00:c6:38:e9:03:69:5e:45:d8:a3:ea:3d
    2e:e3:b8:3f:f0:5b:39:f0:b7:35:64:ed:60:a0:ba:89:28:63:29:e7
    27:82:47:c4:f6:41:53:c8:97:d7:1e:3c:ca:f0:a0:b9:09:0e:3d:f8
    76:5b:10:6f:b5:f8:ef:c5:e8:48:b9:fe:46:a3:c6:ba:b5:05:de:2d
    91:ce:20:12:8f:55:3c:a6:a4:99:bb:91:cf:05:5c:89:d3:a7:dc:a4
    d1:46:f2:dc:36:f3:f0:b5:fd:1d:18:f2:e6:33:d3:38:bb:44:8a:19
    ad:e0:b1:1a:15:c3:56:07:f9:2d:f6:19:f7:cd:80:cf:61:de:58:b8
    a3:f5:e0:d1:a3:3a:19:99:80:b0:63:03:1f:25:05:cc:b2:0c:cd:18
    ef:37:37:46:91:20:04:bc:a3:4a:44:a9:85:3b:50:33:76:45:d9:ba
    26:3a:3b:0d:ff:82:40:36:64:4e:ea:6a:d8:9b:06:ff:3f:e2:c4:a6
    76:ee:8b:58:56:a6:09:d3:4e:08:b0:64:60:75:f3:e2:06:91:64:73
    d2:78:e9:7a:cb:8c:57:0e:d1:9a:6d:3a:4a:9e:5b:d9:e4:a2:ef:31
    5d:2b:2b:53:ab:a1:ad:45:49:fd:a5:e0:8b:4e:0b:71:52:ca:6b:fa
    8b:0e:2c:7c:7b:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://10.1.1.1:8080/crl-as-der/currentcrl-45.crl?id=45
  Authority Information Access OCSP:
    http://10.1.1.1:8090/OCSP-ROOT/
  Use for key: CRL signing, Certificate signing, Key encipherment, Digital signature
  Fingerprint:
    ed:ce:ec:13:1a:d2:ab:0a:76:e5:26:6d:2c:29:5d:49:90:57:f9:41 (sha1)
    af:87:07:69:f0:3e:f7:c6:b8:2c:f8:df:0b:ae:b0:28 (md5)
```

NOTE: In this example, IP addresses are used in the URLs in the CA profile configuration. If IP addresses are not used with CA-issued certificates or CA certificates, DNS must be configured in the device's configuration. DNS must be able to resolve the host in the distribution CRL and in the CA URL in the CA profile configuration. Additionally, you must have network reachability to the same host to receive revocation checks.

Meaning The output shows the details and validity of CA certificate on each peer as follows:

- **C**—Country.
- **O**—Organization.
- **CN**—Common name.

Copyright © 2016, Juniper Networks, Inc.

- **Not before**—Begin date of validity.
- **Not after**—End date of validity.

Verifying Local Certificates

Purpose Verify the validity of a local certificate on each peer device.

Action From operational mode, enter the **show security pki local-certificate certificate-id localcert1 detail** command.

```
user@host> show security pki local-certificate certificate-id localcert1 detail
  Certificate identifier: localcert1
    Certificate version: 3
    Serial number: 013e3f1d
    Issuer:
      Organization: example, Country: US, Common name: OCSP-ROOT
    Subject:
      Organization: example, Organizational unit: example, State: california1, Locality:
    sunnyvale1, Common name: localcert1, Domain component: domain_component1
    Subject string:
      DC=domain_component1, CN=localcert1, OU=example, O=example, L=sunnyvale1,
    ST=california1, C=us1
    Alternate subject: "localcert1@example.net", localcert1.example.net, 10.10.1.50
    Validity:
      Not before: 01-28-2014 22:23 UTC
      Not after: 03-29-2014 22:53 UTC
    Public key algorithm: rsaEncryption(1024 bits)
      30:81:89:02:81:81:00:a6:df:c1:57:59:f8:4d:0f:c4:a8:96:25:97
      03:c4:a0:fb:df:d5:f3:d5:56:b6:5a:26:65:b8:1a:ec:be:f6:c6:5f
      b3:d7:d3:59:39:48:52:4a:e3:1b:e4:e0:6d:24:c3:c1:50:8c:55:3b
      c0:c1:29:a0:45:29:8e:ec:3e:52:2f:84:b3:e8:89:9a:0f:8b:7d:e8
      90:4b:c1:28:48:95:b3:aa:11:ab:b4:8c:a8:80:ce:90:07:2a:13:a2
      2f:84:44:92:3b:be:7d:39:5b:2f:9a:4c:7a:2f:2d:31:8b:12:6d:52
      34:7d:6b:e4:69:7e:f3:86:55:e2:89:31:98:c9:15:02:03:01:00:01
    Signature algorithm: sha1WithRSAEncryption
    Distribution CRL:
      http://10.1.1.1:8080/crl-as-der/currentcrl-45.crl?id=45
    Authority Information Access OCSP:
      http://10.1.1.1/:8090/OCSP-ROOT/
    Fingerprint:
      00:c6:56:64:ad:e3:ce:8e:26:6b:df:17:1e:de:fc:14:a4:bb:8c:e4 (sha1)
      7f:43:c6:ed:e4:b3:7a:4f:9a:8c:0b:61:95:01:c9:52 (md5)
    Auto-re-enrollment:
      Status: Disabled
      Next trigger time: Timer not started
```

Meaning The output shows the details and validity of a local certificate on each peer as follows:

- **DC**—Domain component.
- **CN**—Common name.
- **OU**—Organizational unit.

Copyright © 2016, Juniper Networks, Inc.

- **O**—Organization.
- **L**—Locality
- **ST**—State.
- **C**—Country.
- **Not before**—Begin date of validity.
- **Not after**—End date of validity.

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status on each peer device.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
  Index   State  Initiator cookie  Responder cookie  Mode        Remote Address
  6534660 UP     3e62e05abd6a703f  c552b238e8a26668  IKEv2       198.51.100.50
```

From operational mode, enter the **show security ike security-associations detail** command.

```
user@host> show security ike security-associations detail
  IKE peer 198.51.100.50, Index 6534660, Gateway Name: jsr_gateway
    Role: Responder, State: UP
    Initiator cookie: 3e62e05abd6a703f, Responder cookie: c552b238e8a26668
    Exchange type: IKEv2, Authentication method: RSA-signatures
    Local: 192.0.2.50:500, Remote: 198.51.100.50:500
    Lifetime: Expires in 26906 seconds
    Peer ike-id: localcert11.example.net
    Xauth assigned IP: 0.0.0.0
    Algorithms:
     Authentication      : hmac-sha1-96
     Encryption          : 3des-cbc
     Pseudo random function: hmac-sha1
     Diffie-Hellman group  : DH-group-2
    Traffic statistics:
     Input  bytes  :              2152
     Output bytes  :              2097
     Input  packets:                 4
     Output packets:                 4
    Flags: IKE SA is created
    IPSec security associations: 4 created, 0 deleted
    Phase 2 negotiations in progress: 0

      Negotiation type: Quick mode, Role: Responder, Message ID: 0
      Local: 192.0.2.50:500, Remote: 198.51.100.50:500
      Local identity: 192.0.2.50
      Remote identity: localcert11.example.net
      Flags: IKE SA is created
```

Meaning The **flags** field in the output shows that, IKE security association is created.

Copyright © 2016, Juniper Networks, Inc.

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status on each peer device.

Action From operational mode, enter the **show security ipsec security-associations** command.

```
user@host> show security ipsec security-associations
    Total active tunnels: 1
    ID   Algorithm     SPI    Life:sec/kb  Mon lsys Port  Gateway
    <131073 ESP:3des/sha1 9d1066e2 252/  150000 -  root 500  198.51.100.50
    >131073 ESP:3des/sha1 82079c2c 252/  150000 -  root 500  198.51.100.50
```

From operational mode, enter the **show security ipsec security-associations detail** command.

```
user@host> show security ipsec security-associations detail
    ID: 131073 Virtual-system: root, VPN Name: test_vpn
    Local Gateway: 192.0.2.50, Remote Gateway: 198.51.100.50
    Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    Version: IKEv2
     DF-bit: clear
     Bind-interface: st0.1

    Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
    Last Tunnel Down Reason: Delete payload received
     Direction: inbound, SPI: 9d1066e2, AUX-SPI: 0
                        , VPN Monitoring: -
     Hard lifetime: Expires in 249 seconds
     Lifesize Remaining:  150000 kilobytes
     Soft lifetime: Expires in 10 seconds
     Mode: Tunnel(0 0), Type: dynamic, State: installed
     Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
     Anti-replay service: counter-based enabled, Replay window size: 64

     Direction: outbound, SPI: 82079c2c, AUX-SPI: 0
                        , VPN Monitoring: -
     Hard lifetime: Expires in 249 seconds
     Lifesize Remaining:  150000 kilobytes
     Soft lifetime: Expires in 10 seconds
     Mode: Tunnel(0 0), Type: dynamic, State: installed
     Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
     Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The output shows the ipsec security associations details.

Related Documentation

- Understanding Online Certificate Status Protocol on page 6409
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

Example: Manually Loading a CRL onto the Device

This example shows how to load a CRL manually onto the device.

- Requirements on page 6427
- Overview on page 6427
- Configuration on page 6427
- Verification on page 6428

Requirements

Before you begin:

1. Generate a public and private key pair. See "Example: Generating a Public-Private Key Pair" on page 6368.
2. Generate a certificate request. See "Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server" on page 6376.
3. Configure a certificate authority (CA) profile. See "Example: Configuring a CA Profile" on page 6369.
4. Load your certificate onto the device. See "Example: Loading CA and Local Certificates Manually" on page 6378.

Overview

You can load a CRL manually, or you can have the device load it automatically, when you verify certificate validity. To load a CRL manually, you obtain the CRL from a CA and transfer it to the device (for example, using FTP).

In this example, you load a CRL certificate called **revoke.crl** from the /var/tmp directory on the device. The CA profile is called **ca-profile-ipsec**. (Maximum file size is 5 MB.)

NOTE: If a CRL is already loaded into the ca-profile the command **clear security pki crl ca-profile ca-profile-ipsec** must be run first to clear the old CRL.

Configuration

Step-by-Step Procedure

To load a CRL certificate manually:

1. Load a CRL certificate.

```
[edit]
user@host> request security pki crl load ca-profile ca-profile-ipsec filename
  /var/tmp/revoke.crl
```

NOTE: Junos OS supports loading of CA certificates in X509, PKCS #7, DER, or PEM formats.

Copyright © 2016, Juniper Networks, Inc.

Verification

To verify the configuration is working properly, enter the **show security pki crl** operational mode command.

Related Documentation

- Understanding Certificate Revocation Lists on page 6410
- Digital Certificates Configuration Overview on page 6355
- Example: Verifying Certificate Validity on page 6429
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428
- Deleting a Loaded CRL (CLI Procedure) on page 6431

Example: Configuring a Certificate Authority Profile with CRL Locations

This example shows how to configure a certificate authority profile with CRL locations.

- Requirements on page 6428
- Overview on page 6428
- Configuration on page 6429
- Verification on page 6429

Requirements

Before you begin:

1. Generate a key pair in the device. See "Example: Generating a Public-Private Key Pair" on page 6368.
2. Create a CA profile or profiles containing information specific to a CA. See "Example: Configuring a CA Profile" on page 6369.
3. Obtain a personal certificate from the CA. See "Example: Manually Generating a CSR for the Local Certificate and Sending It to the CA Server" on page 6376.
4. Load the certificate onto the device. See "Example: Loading CA and Local Certificates Manually" on page 6378.
5. Configure automatic reenrollment. See "Example: Configuring SecurID User Authentication" on page 5241.
6. If necessary, load the certificate's CRL on the device. See "Example: Manually Loading a CRL onto the Device" on page 6427.

Overview

In Phase 1 negotiations, you check the CRL list to see if the certificate that you received during an IKE exchange is still valid. If a CRL did not accompany a CA certificate and is not loaded on the device, Junos OS tries to retrieve the CRL through the LDAP or HTTP CRL location defined within the CA certificate itself. If no URL address is defined in the CA certificate, the device uses the URL of the server that you define for that CA certificate.

Copyright © 2016, Juniper Networks, Inc.

If you do not define a CRL URL for a particular CA certificate, the device gets the CRL from the URL in the CA profile configuration.

NOTE: **The CRL distribution point extension (.cdp) in an X509 certificate can be added to either an HTTP URL or an LDAP URL.**

In this example, you direct the device to check the validity of the CA profile called **my_profile** and, if a CRL did not accompany a CA certificate and is not loaded on the device, to retrieve the CRL from the URL **http://abc/abc-crl.crl**.

Configuration

Step-by-Step Procedure

To configure certificate using CRL:

1. Specify the CA profile and URL.

   ```
   [edit]
   user@host# set security pki ca-profile my_profile revocation-check crl url
     http://abc/abc-crl.crl
   ```

2. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security pki** operational mode command.

Related Documentation

- Understanding Certificate Revocation Lists on page 6410
- Example: Manually Loading a CRL onto the Device on page 6427
- Example: Verifying Certificate Validity on page 6429
- Deleting a Loaded CRL (CLI Procedure) on page 6431
- Deleting Certificates (CLI Procedure) on page 6379

Example: Verifying Certificate Validity

This example shows how to verify the validity of a certificate.

- Requirements on page 6430
- Overview on page 6430
- Configuration on page 6430
- Verification on page 6430

Copyright © 2016, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you verify certificates manually to find out whether a certificate has been revoked or whether the CA certificate used to create a local certificate is no longer present on the device.

When you verify certificates manually, the device uses the CA certificate (**ca-cert**) to verify the local certificate (**local.cert**). If the local certificate is valid, and if **revocation-check** is enabled in the CA profile, the device verifies that the CRL is loaded and valid. If the CRL is not loaded and valid, the device downloads the new CRL.

For CA-issued certificates or CA certificates, a DNS must be configured in the device's configuration. The DNS must be able to resolve the host in the distribution CRL and in the CA cert/revocation list url in the ca-profile configuration. Additionally, you must have network reachability to the same host in order for the checks to receive.

Configuration

Step-by-Step Procedure

To manually verify the validity of a certificate:

1. Verify the validity of a local certificate.

   ```
   [edit]
   user@host> request security pki local-certificate verify certificate-id local.cert
   ```

2. Verify the validity of a CA certificate.

   ```
   [edit]
   user@host> request security pki ca-certificate verify ca-profile ca-profile-ipsec
   ```

NOTE: The associated private key and the signature are also verified.

Verification

To verify the configuration is working properly, enter the **show security pki ca-profile** command.

NOTE: If an error is returned instead of a positive verification the failure is logged in pkid.

Related Documentation

- Understanding Certificate Revocation Lists on page 6410
- Example: Manually Loading a CRL onto the Device on page 6427

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428
- Deleting a Loaded CRL (CLI Procedure) on page 6431

Deleting a Loaded CRL (CLI Procedure)

You can choose to delete a loaded CRL if you no longer need to use it to manage certificate revocations and validation.

Use the following command to delete a loaded certificate revocation list:

user@host> clear security pki crl ca-profile (ca-profile **all**)

Specify a CA profile to delete a CRL associated with the CA identified by the profile, or use **all** to delete all CRLs.

Related Documentation

- Understanding Certificate Revocation Lists on page 6410
- Example: Manually Loading a CRL onto the Device on page 6427
- Example: Verifying Certificate Validity on page 6429
- Example: Configuring a Certificate Authority Profile with CRL Locations on page 6428

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 270

Generating Self-Signed Certificates

- Understanding Self-Signed Certificates on page 6433
- Example: Manually Generating Self-Signed Certificates on page 6434
- Using Automatically Generated Self-Signed Certificates (CLI Procedure) on page 6435

Understanding Self-Signed Certificates

A self-signed certificate is a certificate that is signed by its creator rather than by a Certificate Authority (CA).

Self-signed certificates allow for use of SSL-based (Secure Sockets Layer) services without requiring that the user or administrator to undertake the considerable task of obtaining an identity certificate signed by a CA.

NOTE: **Self-signed certificates do not provide additional security as do those generated by CAs. This is because a client cannot verify that the server he or she has connected to is the one advertised in the certificate.**

This topic includes the following sections:

- Generating Self-Signed Certificates on page 6433
- Automatically Generating Self-Signed Certificates on page 6434
- Manually Generating Self-Signed Certificates on page 6434

Generating Self-Signed Certificates

Junos OS provides two methods for generating a self-signed certificate:

- Automatic generation

  In this case, the creator of the certificate is the Juniper Networks device. An automatically generated self-signed certificate is configured on the device by default.

  After the device is initialized, it checks for the presence of an automatically generated self-signed certificate. If it does not find one, the device generates one and saves it in the file system.

- Manual generation

Copyright © 2016, Juniper Networks, Inc.

In this case, you create the self-signed certificate for the device.

At any time, you can use the CLI to generate a self-signed certificate. These certificates are also used to gain access to SSL services.

Self-signed certificates are valid for five years from the time they were generated.

Automatically Generating Self-Signed Certificates

An automatically generated self-signed certificate allows for use of SSL-based services without requiring that the administrator obtain an identity certificate signed by a CA.

A self-signed certificate that is automatically generated by the device is similar to a Secure Shell (SSH) host key. It is stored in the file system, not as part of the configuration. It persists when the device is rebooted, and it is preserved when a **request system snapshot** command is issued.

Manually Generating Self-Signed Certificates

A self-signed certificate that you manually generate allows for use of SSL-based services without requiring that you obtain an identity certificate signed by a CA. A manually generated self-signed certificate is one example of a public key infrastructure (PKI) local certificate. As is true of all PKI local certificates, manually generated self-signed certificates are stored in the file system.

Related Documentation

- Understanding Certificates and PKI on page 6347
- Using Automatically Generated Self-Signed Certificates (CLI Procedure) on page 6435
- Example: Manually Generating Self-Signed Certificates on page 6434

Example: Manually Generating Self-Signed Certificates

This example shows how to generate self-signed certificates manually.

- Requirements on page 6434
- Overview on page 6434
- Configuration on page 6435
- Verification on page 6435

Requirements

Before you begin, generate a public private key pair. See "Example: Generating a Public-Private Key Pair" on page 6368.

Overview

For a manually generated self-signed certificate, you specify the DN when you create it. For an automatically generated self-signed certificate, the system supplies the DN, identifying itself as the creator.

Copyright © 2016, Juniper Networks, Inc.

In this example, you generate a self-signed certificate with the e-mail address as **mholmes@example.net**. You specify a certificate-id of **self-cert** to be referenced by web management, which refers a key-pair of the same certificate-id.

Configuration

Step-by-Step Procedure

To generate the self-signed certificate manually:

- Create the self-signed certificate.

  user@host> **request security pki local-certificate generate-self-signed certificate-id self-cert subject CN=abc domain-name Example.net ip-address 1.2.3.4 email mholmes@example.net**

Verification

To verify the certificate was properly generated and loaded, enter the **show security pki local-certificate** operational mode command.

Related Documentation

- Understanding Self-Signed Certificates on page 6433
- Digital Certificates Configuration Overview on page 6355
- Using Automatically Generated Self-Signed Certificates (CLI Procedure) on page 6435

Using Automatically Generated Self-Signed Certificates (CLI Procedure)

After the device is initialized, it checks for the presence of a self-signed certificate. If a self-signed certificate is not present, the device automatically generates one.

You can add the following statement to your configuration if you want to use the automatically generated self-signed certificate to provide access to HTTPS services:

```
system {
 services {
  web-management {
   http {
    interface [ ... ];
   } https {
    system-generated-certificate;
    interface [ ... ];
   }
  }
 }
}
```

The device uses the following distinguished name for the automatically generated certificate:

" CN=<device serial number>, CN=system generated, CN=self-signed"

Use the following command to specify that the automatically generated self-signed certificate is to be used for Web management HTTPS services:

user@host# **set system services web-management https system-generated-certificate**

Copyright © 2016, Juniper Networks, Inc.

Use the following operational command to delete the automatically generated self-signed certificate:

```
user@host# clear security pki local-certificate system-generated
```

After you delete the system-generated self-signed certificate, the device automatically generates a new one and saves it in the file system.

Related Documentation

- Understanding Self-Signed Certificates on page 6433
- Digital Certificates Configuration Overview on page 6355
- Example: Manually Generating Self-Signed Certificates on page 6434

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 271

Configuring a Device for Certificate Chains

- Understanding Certificate Chains on page 6437
- Example: Configuring a Device for Peer Certificate Chain Validation on page 6440

Understanding Certificate Chains

- Multilevel Hierarchy for Certificate Authentication on page 6437
- Dynamic CRL Download and Checking on page 6439

Multilevel Hierarchy for Certificate Authentication

Certificate-based authentication is an authentication method supported on SRX Series devices during IKE negotiation. In large networks, multiple certificate authorities (CAs) can issue end entity (EE) certificates to their respective end devices. It is common to have separate CAs for individual locations, departments, or organizations.

When a single-level hierarchy for certificate-based authentication is employed, all EE certificates in the network must be signed by the same CA. All firewall devices must have the same CA certificate enrolled for peer certificate validation. The certificate payload sent during IKE negotiation only contains EE certificates.

Alternatively, the certificate payload sent during IKE negotiation can contain a chain of EE and CA certificates. A certificate chain is the list of certificates required to validate a peer's EE certificate. The certificate chain includes the EE certificate and any CA certificates that are not present in the local peer.

The network administrator needs to ensure that all peers participating in an IKE negotiation have at least one common trusted CA in their respective certificate chains. The common trusted CA does not have to be the root CA. The number of certificates in the chain, including certificates for EEs and the topmost CA in the chain, cannot exceed 10.

In the example CA hierarchy shown in Figure 281, Root-CA is the common trusted CA for all devices in the network. Root-CA issues CA certificates to the engineering and sales CAs, which are identified as Eng-CA and Sales-CA, respectively. Eng-CA issues CA certificates to the development and quality assurance CAs, which are identified as Dev-CA and Qa-CA, respectively. Host-A receives its EE certificate from Dev-CA while Host-B receives its EE certificate from Sales-CA.

Copyright © 2016, Juniper Networks, Inc.

Figure 281: Multilevel Hierarchy for Certificate-Based Authentication

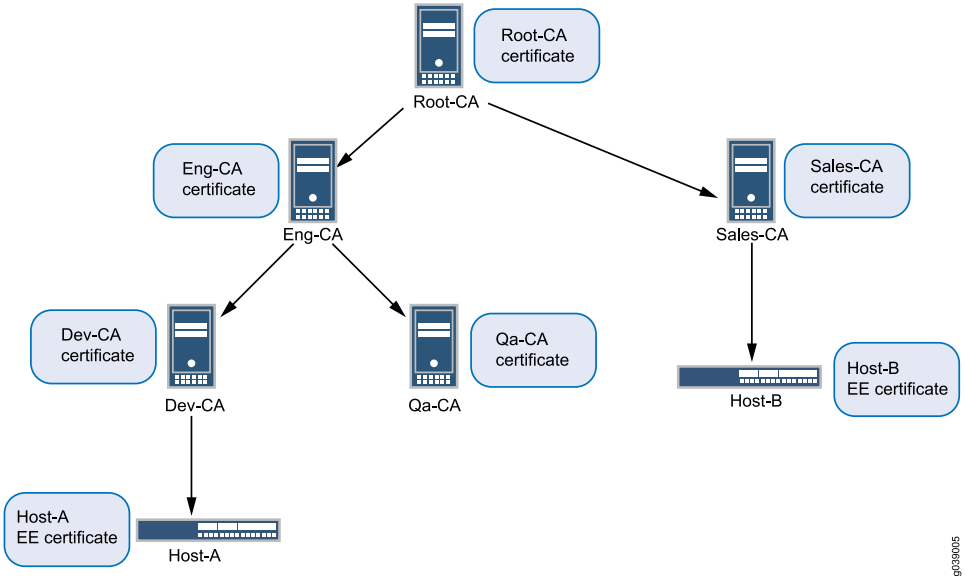

Each end device needs to be loaded with the CA certificates in its hierarchy. Host-A must have Root-CA, Eng-CA, and Dev-CA certificates; Sales-CA and Qa-CA certificates are not necessary. Host-B must have Root-CA and Sales-CA certificates. Certificates can be loaded manually in a device or enrolled using the Simple Certificate Enrollment Process (SCEP).

Each end device must be configured with a CA profile for each CA in the certificate chain. The following output shows the CA profiles configured on Host-A:

```
admin@host-A# show security
pki {
    ca-profile Root-CA {
        ca-identity Root-CA;
        enrollment {
            url “www.example.net/scep/Root/”;
        }
    }
    ca-profile Eng-CA {
        ca-identity Eng-CA;
        enrollment {
            url “www.example.net/scep/Eng/”;
        }
    }
    ca-profile Dev-CA {
        ca-identity Dev-CA;
        enrollment {
            url “www.example.net/scep/Dev/”;
        }
    }
}
```

The following output shows the CA profiles configured on Host-B:

```
admin@host-B# show security
pki {
    ca-profile Root-CA {
```


Copyright © 2016, Juniper Networks, Inc.

```
        ca-identity Root-CA;
        enrollment {
            url “www.example.net/scep/Root/”;
        }
    }
    ca-profile Sales-CA {
        ca-identity Sales-CA;
        enrollment {
            url “www.example.net/scep/Sales/”;
        }
    }
}
```

Dynamic CRL Download and Checking

Digital certificates are issued for a set period of time and are invalid after the specified expiration date. A CA can revoke an issued certificate by listing it in a certificate revocation list (CRL). During peer certificate validation, the revocation status of a peer certificate is checked by downloading the CRL from a CA server to the local device.

A VPN device must be able to check a peer’s certificate for its revocation status. A device can use the CA certificate received from its peer to extract the URL to dynamically download the CA’s CRL and check the revocation status of the peer’s certificate. A dynamic CA profile is automatically created on the local device with the format **dynamic-*nnn***. A dynamic CA profile allows the local device to download the CRL from the peer’s CA and check the revocation status of the peer’s certificate. In Figure 281, Host-A can use the Sales-CA and EE certificates received from Host-B to dynamically download the CRL for Sales-CA and check the revocation status of Host-B’s certificate.

To enable dynamic CA profiles, the **revocation-check crl** option must be configured on a parent CA profile at the [**edit security pki ca-profile *profile-name***] hierarchy level.

The properties of a parent CA profile are inherited for dynamic CA profiles. In Figure 281, the CA profile configuration on Host-A for Root-CA enables dynamic CA profiles as shown in the following output:

```
admin@host-A# show security
pki {
    ca-profile Root-CA {
        ca-identity Root-CA;
        enrollment {
            url “www.example.net/scep/Root/”;
        }
        revocation-check {
            crl;
        }
    }
}
```

A dynamic CA profile is created on Host-A for Sales-CA. Revocation checking is inherited for the Sales-CA dynamic CA profile from Root-CA.

If the **revocation-check disable** statement is configured in a parent CA profile, dynamic CA profiles are not created and dynamic CRL download and checking is not performed.

Copyright © 2016, Juniper Networks, Inc.

The data for CRLs downloaded from dynamic CA profiles are displayed with the **show security pki crl** command in the same way as CRLs downloaded by configured CA profiles. The CRL from a dynamic CA profile is updated periodically as are those for CA profiles that are configured in the device.

NOTE: The CA certificate is required to validate the CRL received from a CA server; therefore, the CA certificate received from a peer is stored on the local device. Because the CA certificate is not enrolled by an administrator, it is used only for validating the CRL received from the CA server and not for validating the peer certificate.

Related Documentation

- Example: Configuring a Device for Peer Certificate Chain Validation on page 6440
- Understanding Certificates and PKI on page 6347
- Understanding Certificate Authority Profiles on page 6369
- Understanding Certificate Revocation Lists on page 6410

Example: Configuring a Device for Peer Certificate Chain Validation

This example shows how to configure a device for certificate chains used to validate peer devices during IKE negotiation.

- Requirements on page 6440
- Overview on page 6440
- Configuration on page 6441
- Verification on page 6446
- IKE and IPsec SA Failure for a Revoked Certificate on page 6447

Requirements

Before you begin, obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates.

Overview

This example shows how to configure a local device for certificate chains, enroll CA and local certificates, check the validity of enrolled certificates, and check the revocation status of the peer device.

This example shows the configuration and operational commands on Host-A, as shown in Figure 282. A dynamic CA profile is automatically created on Host-A to allow Host-A to download the CRL from Sales-CA and check the revocation status of Host-B's certificate.

Copyright © 2016, Juniper Networks, Inc.

Figure 282: Certificate Chain Example

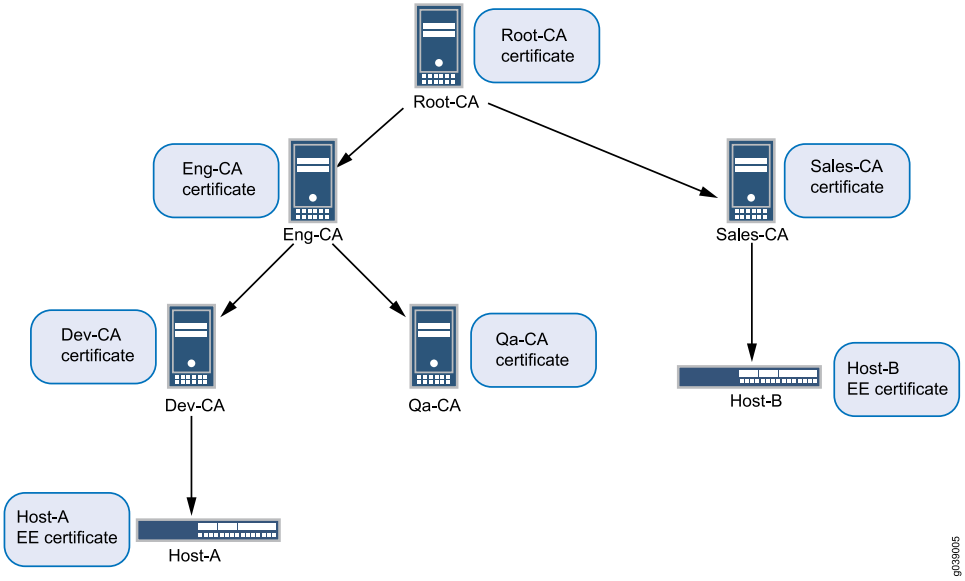

NOTE: The IPsec VPN configuration for Phase 1 and Phase 2 negotiation is shown for Host-A in this example. The peer device (Host-B) must be properly configured so that Phase 1 and Phase 2 options are successfully negotiated and security associations (SAs) are established. See "Configuring Remote IKE IDs for Site-to-Site VPNs" on page 6062 for examples of configuring peer devices for VPNs.

Configuration

To configure a device for certificate chains:

- Configure CA Profiles on page 6441
- Enroll Certificates on page 6443
- Configure IPsec VPN Options on page 6444

Configure CA Profiles

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security pki ca-profile Root-CA ca-identity CA-Root
set security pki ca-profile Root-CA enrollment url http://10.157.88.230:8080/scep/Root/
set security pki ca-profile Root-CA revocation-check crl
set security pki ca-profile Eng-CA ca-identity Eng-CA
set security pki ca-profile Eng-CA enrollment url http://10.157.88.230:8080/scep/Eng/
set security pki ca-profile Eng-CA revocation-check crl
set security pki ca-profile Dev-CA ca-identity Dev-CA
set security pki ca-profile Dev-CA enrollment url http://10.157.88.230:8080/scep/Dev/
```

Copyright © 2016, Juniper Networks, Inc.

```
set security pki ca-profile Dev-CA revocation-check crl
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure CA profiles:

1. Create the CA profile for Root-CA.

```
[edit security pki]
user@host# set ca-profile Root-CA ca-identity CA-Root
user@host# set ca-profile Root-CA enrollment url
  http://10.157.88.230:8080/scep/Root/
user@host# set ca-profile Root-CA revocation-check crl
```

2. Create the CA profile for Eng-CA.

```
[edit security pki]
user@host# set ca-profile Eng-CA ca-identity Eng-CA
user@host# set ca-profile Eng-CA enrollment url
  http://10.157.88.230:8080/scep/Eng/
user@host# set ca-profile Eng-CA revocation-check crl
```

3. Create the CA profile for Dev-CA.

```
[edit security pki]
user@host# set ca-profile Dev-CA ca-identity Dev-CA
user@host# set ca-profile Dev-CA enrollment url
  http://10.157.88.230:8080/scep/Dev/
user@host# set ca-profile Dev-CA revocation-check crl
```

Results

From configuration mode, confirm your configuration by entering the **show security pki** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security pki
ca-profile Root-CA {
  ca-identity Root-CA;
  enrollment {
    url "http:/;/10.157.88.230:8080/scep/Root/";
  }
  revocation-check {
    crl ;
  }
}
ca-profile Eng-CA {
  ca-identity Eng-CA;
  enrollment {
    url "http:/;/10.157.88.230:8080/scep/Eng/";
  }
  revocation-check {
    crl ;
  }
}
```


Copyright © 2016, Juniper Networks, Inc.

```
ca-profile Dev-CA {
  ca-identity Dev-CA;
  enrollment {
    url "http:/;/10.157.88.230:8080/scep/Dev/";
  }
  revocation-check {
    crl ;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Enroll Certificates

Step-by-Step Procedure

To enroll certificates:

1. Enroll the CA certificates.

   user@host> **request security pki ca-certificate enroll ca-profile Root-CA**

   user@host> **request security pki ca-certificate enroll ca-profile Eng-CA**

   user@host> **request security pki ca-certificate enroll ca-profile Dev-CA**

   Type **yes** at the prompts to load the CA certificate.

2. Verify that the CA certificates are enrolled in the device.

```
user@host> show security pki ca-certificate ca-profile Root-CA
Certificate identifier: Root-CA
        Issued to: Root-CA, Issued by: C = us, O = example, CN = Root-CA
        Validity:
          Not before: 08-14-2012 22:19
          Not after: 08-13-2017 22:19
        Public key algorithm: rsaEncryption(2048 bits)

user@host> show security pki ca-certificate ca-profile Eng-CA
Certificate identifier: Eng-CA
        Issued to: Eng-CA, Issued by: C = us, O = example, CN = Root-CA
        Validity:
          Not before: 08-15-2012 01:02
          Not after: 08-13-2017 22:19
        Public key algorithm: rsaEncryption(2048 bits)

user@host> show security pki ca-certificate ca-profile Dev-CA
        Certificate identifier: Dev-CA
        Issued to: Dev-CA, Issued by: C = us, O = example, CN = Eng-CA
        Validity:
          Not before: 08-15-2012 17:41
          Not after: 08-13-2017 22:19
        Public key algorithm: rsaEncryption(2048 bits)
```

3. Verify the validity of the enrolled CA certificates.

```
user@host> request security pki ca-certificate verify ca-profile Root-CA
CA certificate Root-CA verified successfully

user@host> request security pki ca-certificate verify ca-profile Eng-CA
CA certificate Eng-CA verified successfully

user@host> request security pki ca-certificate verify ca-profile Dev-CA
CA certificate Dev-CA verified successfully
```

Copyright © 2016, Juniper Networks, Inc.

4. Enroll the local certificate.

```
user@host> request security pki local-certificate enroll certificate-id Host-A
  ca-profile Dev-CA challenge-password example domain-name host-a.example.net
  email host-a@example.net subject DC=example,CN=Host-A,
  OU=DEV,O=PKI,L=Sunnyvale,ST=CA,C=US
```

5. Verify that the local certificate is enrolled in the device.

```
user@host> show security pki local-certificate
Issued to: Host-A, Issued by: C = us, O = example, CN = Dev-CA
        Validity:
          Not before: 09-17-2012 22:22
          Not after: 08-13-2017 22:19
        Public key algorithm: rsaEncryption(1024 bits)
```

6. Verify the validity of the enrolled local certificate.

```
user@host> request security pki local-certificate verify certificate-id Host-A
Local certificate Host-A verification success
```

7. Check the CRL download for configured CA profiles.

```
user@host> show security pki crl
     CA profile: Root-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Root-CA
        Effective date: 09- 9-2012 13:08
        Next update: 09-21-2012 02:55

      CA profile: Eng-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Eng-CA
        Effective date: 08-22-2012 17:46
        Next update: 10-24-2015 03:33

      CA profile: Dev-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Dev-CA
        Effective date: 09-14-2012 21:15
        Next update: 09-26-2012 11:02
```

Configure IPsec VPN Options

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security ike proposal ike_cert_prop_01 authentication-method rsa-signatures
set security ike proposal ike_cert_prop_01 dh-group group5
set security ike proposal ike_cert_prop_01 authentication-algorithm sha1
set security ike proposal ike_cert_prop_01 encryption-algorithm aes-256-cbc
set security ike policy ike_cert_pol_01 mode main
set security ike policy ike_cert_pol_01 proposals ike_cert_prop_01
set security ike policy ike_cert_pol_01 certificate local-certificate Host-A
set security ike gateway ike_cert_gw_01 ike-policy ike_cert_pol_01
set security ike gateway ike_cert_gw_01 address 30.1.1.51
set security ike gateway ike_cert_gw_01 external-interface ge-0/0/1.0
set security ike gateway ike_cert_gw_01 local-identity 30.1.1.31
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec proposal ipsec_prop_01 protocol esp
set security ipsec proposal ipsec_prop_01 authentication-algorithm hmac-sha1-96
set security ipsec proposal ipsec_prop_01 encryption-algorithm 3des-cbc
set security ipsec proposal ipsec_prop_01 lifetime-seconds 300
set security ipsec policy ipsec_pol_01 proposals ipsec_prop_01
set security ipsec vpn ipsec_cert_vpn_01 bind-interface st0.1
set security ipsec vpn ipsec_cert_vpn_01 ike gateway ike_cert_gw_01
set security ipsec vpn ipsec_cert_vpn_01 ike ipsec-policy ipsec_pol_01
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IPsec VPN options:

1. Configure Phase 1 options.

```
[edit security ike proposal ike_cert_prop_01]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy ike_cert_pol_01]
user@host# set mode main
user@host# set proposals ike_cert_prop_01
user@host# set certificate local-certificate Host-A

[edit security ike gateway ike_cert_gw_01]
user@host# set ike-policy ike_cert_pol_01
user@host# set address 30.1.1.51
user@host# set external-interface ge-0/0/1.0
user@host# set local-identity 30.1.1.31
```

2. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec_prop_01]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm 3des-cbc
user@host# set lifetime-seconds 300

[edit security ipsec policy ipsec_pol_01]
user@host# set proposals ipsec_prop_01

[edit security ipsec vpn ipsec_cert_vpn_01]
user@host# set bind-interface st0.1
user@host# set ike gateway ike_cert_gw_01
user@host# set ike ipsec-policy ipsec_pol_01
```

Results

From configuration mode, confirm your configuration by entering the **show security ike** and **show security ipsec** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security ike
proposal ike_cert_prop_01 {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
    policy ike_cert_pol_01 {
        mode main;
        proposals ike_cert_prop_01;
        certificate {
            local-certificate Host-A;
        }
    }
    gateway ike_cert_gw_01 {
        ike-policy ike_cert_pol_01;
        address 30.1.1.51;
        external-interface ge-0/0/1.0;
    }
[edit]
user@host# show security ipsec
proposal ipsec_prop_01 {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm 3des-cbc;
    lifetime-seconds 300;
}
    policy ipsec_pol_01 {
        proposals ipsec_prop_01;
    }
    vpn ipsec_cert_vpn_01 {
        bind-interface st0.1;
        ike {
            gateway ike_cert_gw_01;
            ipsec-policy ipsec_pol_01;
        }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

If certificate validation is successful during IKE negotiation between peer devices, both IKE and IPsec security associations (SAs) are established.

- Verifying IKE Phase 1 Status on page 6446
- Verifying IPsec Phase 2 Status on page 6447

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Copyright © 2016, Juniper Networks, Inc.

Action Enter the **show security ike security-associations** command from operational mode.

```
user@host> show security ike security-associations
    Index   State  Initiator cookie  Responder cookie  Mode           Remote
Address
    2090205 UP     285feacb50824495  59fca3f72b64da10  Main           30.1.1.51
```

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action Enter the **show security ipsec security-associations** command from operational mode.

```
user@host> show security ipsec security-associations
      Total active tunnels: 1
      ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
      <131073 ESP:3des/sha1 a4756de9 207/  unlim   -   root 500   30.1.1.51

      >131073 ESP:3des/sha1 353bacd3 207/  unlim   -   root 500   30.1.1.51
```

IKE and IPsec SA Failure for a Revoked Certificate

- Checking for Revoked Certificates on page 6447

Checking for Revoked Certificates

Problem If certificate validation fails during IKE negotiation between peer devices, check to make sure that the peer's certificate has not been revoked. A dynamic CA profile allows the local device to download the CRL from the peer's CA and check the revocation status of the peer's certificate. To enable dynamic CA profiles, the **revocation-check crl** option must be configured on a parent CA profile.

Solution To check the revocation status of a peer's certificate:

1. Identify the dynamic CA profile that will show the CRL for the peer device by entering the **show security pki crl** command from operational mode.

```
user@host> show security pki crl
     CA profile: Root-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Root-CA
        Effective date: 09- 9-2012 13:08
        Next update: 09-21-2012 02:55

      CA profile: Eng-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Eng-CA
        Effective date: 08-22-2012 17:46
        Next update: 10-24-2015 03:33

      CA profile: Dev-CA
        CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Dev-CA
```

Copyright © 2016, Juniper Networks, Inc.

```
    Effective date: 09-14-2012 21:15
    Next update: 09-26-2012 11:02

  CA profile: dynamic-001
    CRL version: V00000001
    CRL issuer: C = us, O = example, CN = Sales-CA
    Effective date: 09-14-2012 21:15
    Next update: 09-26-2012 11:02
```

The CA profile **dynamic-001** is automatically created on Host-A so that Host-A can download the CRL from Host-B's CA (Sales-CA) and check the revocation status of the peer's certificate.

2. Display CRL information for the dynamic CA profile by entering the **show security pki crl ca-profile dynamic-001 detail** command from operational mode.

   Enter

```
user@host> show security pki crl ca-profile dynamic-001 detail
    CA profile: dynamic-001
      CRL version: V00000001
        CRL issuer: C = us, O = example, CN = Sub11
        Effective date: 09-19-2012 17:29
        Next update: 09-20-2012 01:49
        Revocation List:
          Serial number              Revocation date
          10647C84                   09-19-2012 17:29 UTC
```

   Host-B's certificate (serial number 10647084) has been revoked.

Related Documentation

- Understanding Certificate Chains on page 6437
- Understanding Certificates and PKI on page 6347
- Understanding Certificate Authority Profiles on page 6369
- Understanding Certificate Revocation Lists on page 6410

Copyright © 2016, Juniper Networks, Inc.

PART 86

Configuring AutoVPN

- Configuring AutoVPN on Hub-and-Spoke Devices on page 6451
- Configuring Auto Discovery VPNs on page 6559
- Configuring AutoVPN and Traffic Selectors on page 6607

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 272

Configuring AutoVPN on Hub-and-Spoke Devices

- Understanding AutoVPN on page 6451
- Understanding AutoVPN Limitations on page 6453
- Understanding Spoke Authentication in AutoVPN Deployments on page 6454
- AutoVPN Configuration Overview on page 6456
- Example: Configuring Basic AutoVPN with iBGP on page 6457
- Example: Configuring Basic AutoVPN with OSPF on page 6482
- Example: Configuring AutoVPN with iBGP and ECMP on page 6506
- Example: Configuring AutoVPN with iBGP and Active-Backup Tunnels on page 6531

Understanding AutoVPN

AutoVPN supports an IPsec VPN aggregator (known as a *hub*) that serves as a single termination point for multiple tunnels to remote sites (known as *spokes*). AutoVPN allows network administrators to configure a hub for current and future spokes. No configuration changes are required on the hub when spoke devices are added or deleted, thus allowing administrators flexibility in managing large-scale network deployments.

- Secure Tunnel Modes on page 6451
- Authentication on page 6452
- Configuration and Management on page 6452

Secure Tunnel Modes

AutoVPN is supported on route-based IPsec VPNs. For route-based VPNs, you configure a secure tunnel (st0) interface and bind it to an IPsec VPN tunnel. st0 interfaces in AutoVPN networks can be configured in one of two modes:

- Point-to-point mode—By default, an st0 interface configured at the [**edit interfaces st0 unit** *x*] hierarchy level is in point-to-point mode.
- Point-to-multipoint mode—In this mode, the **multipoint** option is configured at the [**edit interfaces st0 unit** *x*] hierarchy level on both AutoVPN hub and spokes. st0

Copyright © 2016, Juniper Networks, Inc.

interfaces on the hub and spokes must be numbered and the IP address configured on a spoke must exist in the hub's st0 interface subnetwork.

Table 553 compares AutoVPN point-to-point and point-to-multipoint secure tunnel interface modes.

Table 553: Comparison Between AutoVPN Point-to-Point and Point-to-Multipoint Secure Tunnel Modes

Point-to-Point Mode	Point-to-Multipoint Mode
Uses traffic selectors to forward packets through VPN tunnels. Traffic selectors must be configured on each spoke. Administrator needs to be aware of the types of traffic that need to be permitted through the VPN tunnel.	Uses dynamic routing protocol to forward packets through VPN tunnels. The dynamic routing protocol must run in point-to-multipoint mode.
Does not support dynamic routing protocols on the st0 interface when traffic selectors are configured.	Cannot configure an st0 interface in point-to-multipoint mode with traffic selectors.
Supports IPv4 traffic only.	Supports IPv4 traffic only.
Allows spoke devices to be non-SRX Series devices.	Requires that hub and spoke devices are SRX Series devices.
Supports IKEv1 or IKEv2.	Supports IKEv1 only.
Supports dead peer detection only.	Supports dead peer detection and VPN monitoring.
Supports larger numbers of tunnels and spokes.	

Authentication

The supported authentication for AutoVPN hubs and spokes is X.509 public key infrastructure (PKI) certificates. The group IKE user type configured on the hub allows strings to be specified to match the alternate subject field in spoke certificates. Partial matches for the subject fields in spoke certificates can also be specified. See "Understanding Spoke Authentication in AutoVPN Deployments" on page 6454.

Configuration and Management

AutoVPN is configured and managed on SRX Series devices using the CLI. Multiple AutoVPN hubs can be configured on a single SRX Series device. The maximum number of spokes supported by a configured hub is specific to the model of the SRX Series device.

Related Documentation

- Understanding AutoVPN Limitations on page 6453
- Understanding Spoke Authentication in AutoVPN Deployments on page 6454
- AutoVPN Configuration Overview on page 6456

Copyright © 2016, Juniper Networks, Inc.

Understanding AutoVPN Limitations

The following features are not supported for AutoVPN:

- AutoVPN does not support IPv6 traffic.
- AutoVPN tunnels are only supported on SRX Series devices for specific releases of Junos OS. AutoVPN tunnels cannot interoperate with any other Juniper Networks devices or other vendors' devices. However, host-protected third-party devices like LTE eNodeBs and dial-up clients are supported with point-to-point secure tunnel mode.
- Policy-based VPNs are not supported.
- The RIP dynamic routing protocol is not supported with AutoVPN tunnels. We recommend using OSPF and iBGP for dynamic routing when using point-to-multipoint VPN tunnels.
- Manual keys and Autokey IKE with preshared keys are not supported.
- Configuring static next-hop tunnel binding (NHTB) on the hub for spokes is not supported.
- AutoVPN does not support multicast traffic.
- When IKE main mode is used with PKI authentication, all gateway configurations that use the same external interface on a device must use the same IKE policy.
- The IKE gateway connections-limit configuration is not supported for high-end SRX Series devices.
- The group IKE ID user type is not supported with an IP address as the IKE ID.
- The IKE ID should not overlap with other IKE gateways when the group IKE ID user type is used.
- VPNs on SRX Series devices support only one IKE connection between two peers. The IKE connection is identified by a set of local IP addresses and ports and remote IP addresses and ports.
- ASN1 distinguished names that are longer than 2047 characters are not supported for PKI authentication.
- A secure tunnel (st0) interface supports only one IPv4 address and one IPv6 address at the same time. This applies to all route-based VPNs, including AutoVPNs.
- Configuring XAuth with AutoVPN st0 interfaces in point-to-multipoint mode and dynamic IKE gateways is not supported.

Related Documentation

- Understanding AutoVPN on page 6451

Copyright © 2016, Juniper Networks, Inc.

Understanding Spoke Authentication in AutoVPN Deployments

In AutoVPN deployments, the hub and spoke devices must have valid X.509 PKI certificates loaded. You can use the **show security pki local-certificate detail** command to display information about the certificates loaded in a device.

This topic covers the configuration on the hub that allows spokes to authenticate and connect to the hub:

- Group IKE ID Configuration on the Hub on page 6454
- Excluding a Spoke Connection on page 6456

Group IKE ID Configuration on the Hub

The group IKE ID feature allows a number of spoke devices to share an IKE configuration on the hub. The certificate holder's identification, in the subject or alternate subject fields in each spoke's X.509 certificate, must contain a part that is common to all spokes; the common part of the certificate identification is specified for the IKE configuration on the hub.

For example, the IKE ID **example.net** can be configured on the hub to identify spokes with the hostnames **host1.example.net**, **host2.example.net**, and **host3.example.net**. The certificate on each spoke must contain a hostname identity in the alternate subject field with **example.net** in the right-most part of the field; for example, **host1.example.net**. In this example, all spokes use this hostname identity in their IKE ID payload. During IKE negotiation, the IKE ID from a spoke is used to match the common part of the peer IKE identity configured on the hub. A valid certificate authenticates the spoke.

The common part of the certificate identification can be one of the following:

- A partial hostname in the right-most part of the alternate subject field of the certificate, for example **example.net**.
- A partial e-mail address in the right-most part of the alternate subject field of the certificate, for example **@example.net**.
- A container string, a set of wildcards, or both to match the subject fields of the certificate. The subject fields contain details of the digital certificate holder in Abstract Syntax Notation One (ASN.1) distinguished name (DN) format. Fields can include organization, organizational unit, country, locality, or common name.

  To configure a group IKE ID to match subject fields in certificates, you can specify the following types of identity matches:

  - Container—The hub authenticates the spoke's IKE ID if the subject fields of the spoke's certificate exactly match the values configured on the hub. Multiple entries can be specified for each subject field (for example, **ou=eng,ou=sw**). The order of values in the fields must match.
  - Wildcard—The hub authenticates the spoke's IKE ID if the subject fields of the spoke's certificate match the values configured on the hub. The wildcard match supports

Copyright © 2016, Juniper Networks, Inc.

only one value per field (for example, **ou=eng** or **ou=sw** but not **ou=eng,ou=sw**). The order of the fields is inconsequential.

The following example configures a group IKE ID with the partial hostname **example.net** in the alternate subject field of the certificate.

```
[edit]
security {
  ike {
    policy common-cert-policy {
      proposals common-ike-proposal;
      certificate {
        local-certificate hub-local-certificate;
      }
    }
    gateway common-gateway-to-all-spoke-peer {
      ike-policy common-cert-policy;
      dynamic {
        hostname example.net;
        ike-user-type group-ike-id;
      }
      external-interface fe-0/0/2;
    }
  }
}
```

In this example, **example.net** is the common part of the hostname identification used for all spokes. All X.509 certificates on the spokes must contain a hostname identity in the alternate subject field with **example.net** in the right-most part. All spokes must use the hostname identity in their IKE ID payload.

The following example configures a group IKE ID with wildcards to match the values **sales** in the organizational unit and **example** in the organization subject fields of the certificate.

```
[edit]
security {
  ike {
    policy common-cert-policy {
      proposals common-ike-proposal;
      certificate {
        local-certificate hub-local-certificate;
      }
    }
    gateway common-gateway-to-all-spoke-peer {
      ike-policy common-cert-policy;
      dynamic {
        distinguished-name {
          wildcard ou=sales,o=example;
        }
        ike-user-type group-ike-id;
      }
      external-interface fe-0/0/2;
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
```

In this example, the fields **ou=sales,o=example** are the common part of the subject field in the certificates expected from the spokes. During IKE negotiation, if a spoke presents a certificate with the subject fields **cn=alice,ou=sales,o=example** in its certificate, authentication succeeds and the tunnel is established. If a spoke presents a certificate with the subject fields **cn=thomas,ou=engineer,o=example** in its certificate, the certificate is rejected by the hub as the organization unit should be **sales**.

Excluding a Spoke Connection

To exclude a particular spoke from connecting to the hub, the certificate for that spoke must be revoked. The hub needs to retrieve the latest certificate revocation list (CRL) from the CA that contains the serial number of the revoked certificate. The hub will then refuse a VPN connection from the revoked spoke. Until the latest CRL is available in the hub, the hub might continue to establish a tunnel from the revoked spoke. For more information, see "Understanding Certificate Revocation Lists" on page 6410 and "Understanding Certificate Authority Profiles" on page 6369.

Related Documentation

- Understanding AutoVPN on page 6451
- AutoVPN Configuration Overview on page 6456

AutoVPN Configuration Overview

The following steps describe the basic tasks for configuring AutoVPN on hub and spoke devices. The AutoVPN hub is configured *once* for all current and new spokes.

To configure the AutoVPN hub:

1. Enroll a CA certificate and the local certificate in the device.
2. Create a secure tunnel (st0) interface and configure it in point-to-multipoint mode.
3. Configure a single IKE policy.
4. Configure an IKE gateway with a group IKE ID that is common to all spokes.
5. Configure a single IPsec policy and VPN.
6. Configure a dynamic routing protocol.

To configure an SRX Series AutoVPN spoke device:

1. Enroll a CA certificate and the local certificate in the device.
2. Create an st0 interface and configure it in point-to-multipoint mode.
3. Configure an IKE policy to match the IKE policy configured on the hub.
4. Configure an IKE gateway with an ID to match the group IKE ID configured on the hub.

NOTE: Only IKEv1 is supported on an SRX Series spoke with st0 interfaces in point-to-point mode.

Copyright © 2016, Juniper Networks, Inc.

5. Configure an IPsec policy to match the IPsec policy configured on the hub.
6. Configure a dynamic routing protocol.

Related Documentation

- Understanding Traffic Selectors in Route-Based VPNs on page 6195
- Understanding AutoVPN on page 6451

Example: Configuring Basic AutoVPN with iBGP

This example shows how to configure an AutoVPN hub to act as a single termination point, and then configure two spokes to act as tunnels to remote sites. This example configures iBGP to forward packets through the VPN tunnels.

- Requirements on page 6457
- Overview on page 6457
- Configuration on page 6460
- Verification on page 6480

Requirements

This example uses the following hardware and software components:

- Three supported SRX Series devices as AutoVPN hub and spokes
- Junos OS Release 12.1X44-D10 and later that support AutoVPN

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates.

NOTE: You should be familiar with the dynamic routing protocol that is used to forward packets through the VPN tunnels. For more information about specific requirements for a dynamic routing protocol, see the *Routing Protocols Overview for Security Devices*.

Overview

This example shows the configuration of an AutoVPN hub and the subsequent configurations of two spokes.

In this example, the first step is to enroll digital certificates in each device using the Simple Certificate Enrollment Protocol (SCEP). The certificates for the spokes contain the organizational unit (OU) value “SLT” in the subject field; the hub is configured with a group IKE ID to match the value “SLT” in the OU field.

The spokes establish IPsec VPN connections to the hub, which allows them to communicate with each other as well as access resources on the hub. Phase 1 and Phase

Copyright © 2016, Juniper Networks, Inc.

2 IKE tunnel options configured on the AutoVPN hub and all spokes must have the same values. Table 554 shows the options used in this example.

Table 554: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke Configurations

Option	Value
IKE proposal:	
Authentication method	RSA digital certificates
Diffie-Hellman (DH) group	2
Authentication algorithm	SHA-1
Encryption algorithm	AES 128 CBC
IKE policy:	
Mode	Main
IPsec proposal:	
Protocol	ESP
Authentication algorithm	HMAC MD5 96
Encryption algorithm	DES CBC
IPsec policy:	
Perfect Forward Secrecy (PFS) group	14

The same certificate authority (CA) is configured on all devices.

NOTE: **Junos OS only supports a single level of certificate hierarchy.**

Table 555 shows the options configured on the hub and on all spokes.

Table 555: AutoVPN Configuration for Hub and All Spokes

Option	Hub	All Spokes
IKE gateway:		
Remote IP address	Dynamic	1.1.1.1
Remote IKE ID	Distinguished name (DN) on the spoke's certificate with the string **SLT** in the organizational unit (OU) field	DN on the hub's certificate

Copyright © 2016, Juniper Networks, Inc.

Table 555: AutoVPN Configuration for Hub and All Spokes *(continued)*

Option	Hub	All Spokes
Local IKE ID	DN on the hub's certificate	DN on the spoke's certificate
External interface	ge-0/0/1.0	Spoke 1: fe-0/0/1.0 Spoke 2: ge-0/0/1.0
VPN:		
Bind interface	st0.0	st0.0
Establish tunnels	(not configured)	Immediately on configuration commit

Table 556 shows the configuration options that are different on each spoke.

Table 556: Comparison Between the Spoke Configurations

Option	Spoke 1	Spoke 2
st0.0 interface	10.10.10.2/24	10.10.10.3/24
Interface to internal network	(fe-0.0/4.0) 60.60.60.1/24	(fe-0.0/4.0) 70.70.70.1/24
Interface to Internet	(fe-0/0/1.0) 2.2.2.1/30	(ge-0/0/1.0) 3.3.3.1/30

Routing information for all devices is exchanged through the VPN tunnels.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Topology

Figure 283 shows the SRX Series devices to be configured for AutoVPN in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 283: Basic AutoVPN Deployment with iBGP

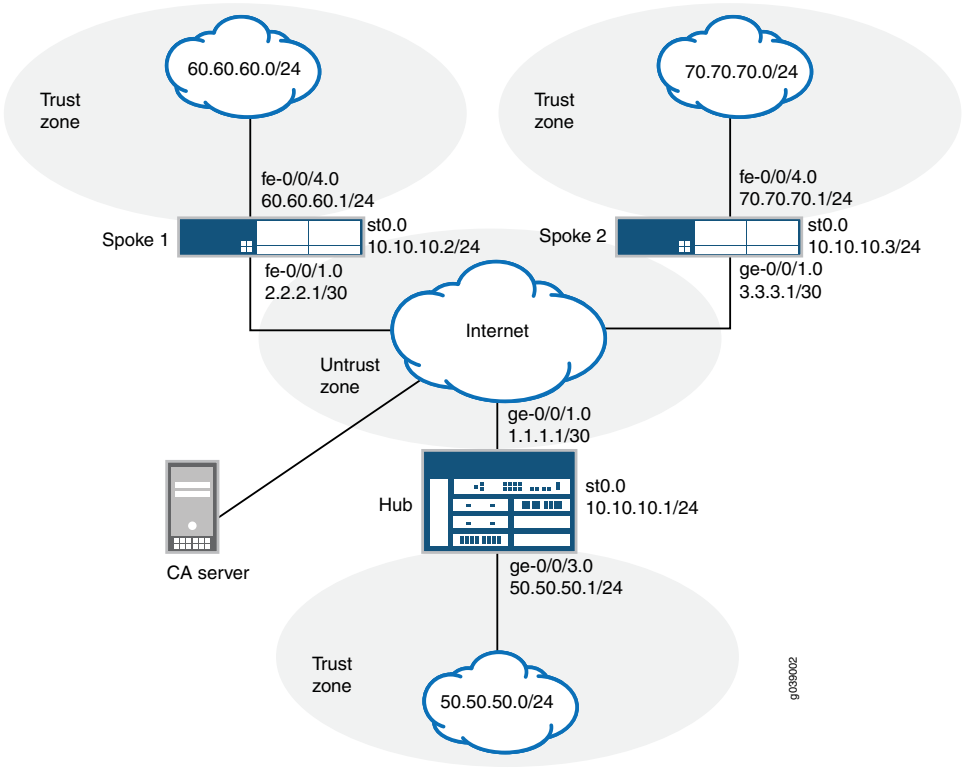

Configuration

To configure AutoVPN, perform these tasks:

NOTE: The first section describes how to obtain CA and local certificates online using the Simple Certificate Enrollment Protocol (SCEP) on the hub and spoke devices.

- Enroll Device Certificates with SCEP on page 6460
- Configuring the Hub on page 6464
- Configuring Spoke 1 on page 6469
- Configuring Spoke 2 on page 6474

Enroll Device Certificates with SCEP

Step-by-Step Procedure

To enroll digital certificates with SCEP on the hub:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

```
user@host> request security pki generate-key-pair certificate-id Local1
```

4. Enroll the local certificate.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email hub@example.net ip-address
  1.1.1.1 subject
  DC=example.net,CN=hub,OU=SLT,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificate.

```
user@host> show security pki local-certificate detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a6d5f300000000258d
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Bangalore, Common name: hub, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SLT, CN=hub
  Alternate subject: "hub@example.net", example.net, 1.1.1.1
  Validity:
    Not before: 11- 6-2012 09:39
    Not after: 11- 6-2013 09:49
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:c9:c9:cc:30:b6:7a:86:12:89:b5:18:b3:76
    01:2d:cc:65:a8:a8:42:78:cd:d0:9a:a2:c0:aa:c4:bd:da:af:88:f3
    2a:78:1f:0a:58:e6:11:2c:81:8f:0e:7c:de:86:fc:48:4c:28:5b:8b
    34:91:ff:2e:91:e7:b5:bd:79:12:de:39:46:d9:fb:5c:91:41:d1:da
    90:f5:09:00:9b:90:07:9d:50:92:7d:ff:fb:3f:3c:bc:34:e7:e3:c8
    ea:cb:99:18:b4:b6:1d:a8:99:d3:36:b9:1b:36:ef:3e:a1:fd:48:82
    6a:da:22:07:da:e0:d2:55:ef:57:be:09:7a:0e:17:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    e1:f7:a1:a6:1e:c3:97:69:a5:07:9b:09:14:1a:c7:ae:09:f1:f6:35 (sha1)
    a0:02:fa:8d:5c:63:e5:6d:f7:f4:78:56:ac:4e:b2:c4 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 1:

1. Configure the CA.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

```
user@host> request security pki generate-key-pair certificate-id Local1
```

4. Enroll the local certificate.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email spoke1@example.net
  ip-address 2.2.2.1 subject
  DC=example.net,CN=spoke1,OU=SLT,O=example,L=Mysore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificate.

```
user@host> show security pki local-certificate detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a7975f00000000258e
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Mysore, Common name: spoke1, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
  Alternate subject: "spoke1@example.net", example.net, 2.2.2.1
  Validity:
    Not before: 11- 6-2012 09:40
    Not after: 11- 6-2013 09:50
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:d8:45:09:77:cd:36:9a:6f:58:44:18:91:db
    b0:c7:8a:ee:c8:d7:a6:d2:e2:e7:20:46:2b:26:1a:92:e2:4e:8a:ce
    c9:25:d9:74:a2:81:ad:ea:e0:38:a0:2f:2d:ab:a6:58:ac:88:35:f4
    90:01:08:33:33:75:2c:44:26:f8:25:18:97:96:e4:28:de:3b:35:f2
    4a:f5:92:b7:57:ae:73:4f:8e:56:71:ab:81:54:1d:75:88:77:13:64
    1b:6b:01:96:15:0a:1c:54:e3:db:f8:ec:ec:27:5b:86:39:c1:09:a1
    e4:24:1a:19:0d:14:2c:4b:94:a4:04:91:3f:cb:ef:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    b6:24:2a:0e:96:5d:8c:4a:11:f3:5a:24:89:7c:df:ea:d5:c0:80:56 (sha1)
    31:58:7f:15:bb:d4:66:b8:76:1a:42:4a:8a:16:b3:a9 (md5)
  Auto-re-enrollment:
```


Copyright © 2016, Juniper Networks, Inc.

```
  Status: Disabled
  Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT. The IKE configuration on the hub includes ou=SLT to identify the spoke.

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 2:

1. Configure the CA.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
   user@host# set security pki ca-profile ca-profile1 enrollment url
     http://pc4/certsrv/mscep/mscep.dll
   user@host# set security pki ca-profile ca-profile1 revocation-check disable
   user@host# commit
   ```

2. Enroll the CA certificate.

   ```
   user@host> request security pki ca-certificate enroll ca-profile ca-profile1
   ```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

   ```
   user@host> request security pki generate-key-pair certificate-id Local1
   ```

4. Enroll the local certificate.

   ```
   user@host> request security pki local-certificate enroll ca-profile ca-profile1
     certificate-id Local1 domain-name example.net email spoke2@example.net
     ip-address 3.3.3.1 subject
     DC=example.net,CN=spoke2,OU=SLT,O=example,L=Tumkur,ST=KA,C=IN
     challenge-password <password>
   ```

5. Verify the local certificate.

   ```
   user@host> show security pki local-certificate detail

   Certificate identifier: Local1
     Certificate version: 3
     Serial number: 40bb71d400000000258f
     Issuer:
       Common name: CASERVER1, Domain component: net, Domain component: internal

     Subject:
       Organization: example, Organizational unit: SLT, Country: IN, State: KA,

       Locality: Tumkur, Common name: spoke2, Domain component: example.net
     Subject string:
       C=IN, DC=example.net, ST=KA, L=Tumkur, O=example, OU=SLT, CN=spoke2
     Alternate subject: "spoke2@example.net", example.net, 3.3.3.1
     Validity:
       Not before: 11- 6-2012 10:02
       Not after: 11- 6-2013 10:12
     Public key algorithm: rsaEncryption(1024 bits)
       30:81:89:02:81:81:00:b6:2e:e2:da:e6:ac:57:e4:5d:ff:de:f6:89
       27:d6:3e:1b:4a:3f:b2:2d:b3:d3:61:ed:ed:6a:07:d9:8a:d2:24:03
   ```

Copyright © 2016, Juniper Networks, Inc.

```
    77:1a:fe:84:e1:12:8a:2d:63:6e:bf:02:6b:15:96:5a:4f:37:a0:46
    44:09:96:c0:fd:bb:ab:79:2c:5d:92:bd:31:f0:3b:29:51:ce:89:8e
    7c:2b:02:d0:14:5b:0a:a9:02:93:21:ea:f9:fc:4a:e7:08:bc:b1:6d
    7c:f8:3e:53:58:8e:f1:86:13:fe:78:b5:df:0b:8e:53:00:4a:46:11
    58:4a:38:e9:82:43:d8:25:47:7d:ef:18:f0:ef:a7:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    1a:6d:77:ac:fd:94:68:ce:cf:8a:85:f0:39:fc:e0:6b:fd:fe:b8:66 (sha1)
    00:b1:32:5f:7b:24:9c:e5:02:e6:72:75:9e:a5:f4:77 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT. The IKE configuration on the hub includes ou=SLT to identify the spoke.

Configuring the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/3 unit 0 family inet address 50.50.50.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.1/24
set policy-options policy-statement lan_nw from interface ge-0/0/3.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.10.10.1
set protocols bgp group ibgp export lan_nw
set protocols bgp group ibgp cluster 1.2.3.4
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp allow 10.10.10.0/24
set routing-options static route 2.2.2.0/30 next-hop 1.1.1.2
set routing-options static route 3.3.3.0/30 next-hop 1.1.1.2
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
set security ike gateway hub-to-spoke-gw ike-policy ike-policy1
set security ike gateway hub-to-spoke-gw dynamic distinguished-name wildcard OU=SLT
set security ike gateway hub-to-spoke-gw dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw local-identity distinguished-name
set security ike gateway hub-to-spoke-gw external-interface ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn hub-to-spoke-vpn bind-interface st0.0
set security ipsec vpn hub-to-spoke-vpn ike gateway hub-to-spoke-gw
set security ipsec vpn hub-to-spoke-vpn ike ipsec-policy vpn-policy1
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the hub:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 1.1.1.1/30
user@host# set ge-0/0/3 unit 0 family inet address 50.50.50.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.1/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface ge-0/0/3.0
user@host# set policy-statement lan_nw then accept

[edit protocols bgp]
user@host# set group ibgp type internal
user@host# set group ibgp local-address 10.10.10.1
user@host# set group ibgp export lan_nw
user@host# set group ibgp cluster 1.2.3.4
user@host# set group ibgp peer-as 10
user@host# set group ibgp allow 10.10.10.0/24

[edit routing-options]
user@host# set static route 2.2.2.0/30 next-hop 1.1.1.2
user@host# set static route 3.3.3.0/30 next-hop 1.1.1.2
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike gateway hub-to-spoke-gw]
user@host# set ike-policy ike-policy1
user@host# set dynamic distinguished-name wildcard OU=SLT
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn hub-to-spoke-vpn]
user@host# set bind-interface st0.0
user@host# set ike gateway hub-to-spoke-gw
user@host# set ike ipsec-policy vpn-policy1
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/3.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
   unit 0 {
      family inet {
         address 1.1.1.1/30;
      }
   }
}
   ge-0/0/3 {
      unit 0 {
         family inet {
            address 50.50.50.1/24;
         }
      }
   }
   st0 {
      unit 0 {
         multipoint;
         family inet {
            address 10.10.10.1/24;
         }
      }
   }
[edit]
user@host# show policy-options
policy-statement lan_nw {
   from interface ge-0/0/3.0;
   then accept;
}
[edit]
user@host# show protocols
bgp {
   group ibgp {
      type internal;
      local-address 10.10.10.1;
      export lan_nw;
      cluster 1.2.3.4;
      peer-as 10;
      allow 10.10.10.0/24;
   }
}
[edit]
user@host# show routing-options
static {
   route 2.2.2.0/30 next-hop 1.1.1.2;
   route 3.3.3.0/30 next-hop 1.1.1.2;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy1 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local1;
        }
    }
    gateway hub-to-spoke-gw {
        ike-policy ike-policy1;
        dynamic {
            distinguished-name {
                wildcard OU=SLT;
            }
            ike-user-type group-ike-id;
        }
        local-identity distinguished-name;
        external-interface ge-0/0/1.0;
    }
[edit]
user@host# show security ipsec
    proposal ipsec-proposal {
        protocol esp;
        authentication-algorithm hmac-md5-96;
        encryption-algorithm des-cbc;
    }
    policy vpn-policy1 {
        perfect-forward-secrecy {
            keys group14;
        }
        proposals ipsec-proposal;
    }
    vpn hub-to-spoke-vpn {
        bind-interface st0.0;
        ike {
            gateway hub-to-spoke-gw;
            ipsec-policy vpn-policy1;
        }
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
```


Copyright © 2016, Juniper Networks, Inc.

```
        all;
      }
    }
    interfaces {
      st0.0;
      ge-0/0/1.0;
    }
}
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        ge-0/0/3.0;
      }
    }
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/1 unit 0 family inet address 2.2.2.1/30
set interfaces fe-0/0/4 unit 0 family inet address 60.60.60.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.2/24
set policy-options policy-statement lan_nw from interface fe-0/0/4.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.10.10.2
```

Copyright © 2016, Juniper Networks, Inc.

```
set protocols bgp group ibgp export lan_nw
set protocols bgp group ibgp neighbor 10.10.10.1
set routing-options static route 1.1.1.0/30 next-hop 2.2.2.2
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
set security ike gateway spoke-to-hub-gw ike-policy ike-policy1
set security ike gateway spoke-to-hub-gw address 1.1.1.1
set security ike gateway spoke-to-hub-gw local-identity distinguished-name
set security ike gateway spoke-to-hub-gw remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw external-interface fe-0/0/1.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn spoke-to-hub bind-interface st0.0
set security ipsec vpn spoke-to-hub ike gateway spoke-to-hub-gw
set security ipsec vpn spoke-to-hub ike ipsec-policy vpn-policy1
set security ipsec vpn spoke-to-hub establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 1:

1. Configure interfaces.

```
[edit interfaces]
user@host# set fe-0/0/1 unit 0 family inet address 2.2.2.1/30
user@host# set fe-0/0/4 unit 0 family inet address 60.60.60.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.2/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface fe-0/0/4.0
user@host# set policy-statement lan_nw then accept
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols bgp]
user@host# set group ibgp type internal
user@host# set group ibgp local-address 10.10.10.2
user@host# set group ibgp export lan_nw
user@host# set group ibgp neighbor 10.10.10.1

[edit routing-options]
user@host# set static route 1.1.1.0/30 next-hop 2.2.2.2
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike gateway spoke-to-hub-gw]
user@host# set ike-policy ike-policy1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub]
user@host# set bind-interface st0.0
user@host# set ike gateway spoke-to-hub-gw
user@host# set ike ipsec-policy vpn-policy1
user@host# set establish-tunnels immediately
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 2.2.2.1/30;
    }
  }
}
  fe-0/0/4 {
    unit 0 {
      family inet {
        address 60.60.60.1/24;
      }
    }
  }
  st0 {
    unit 0 {
      multipoint;
      family inet {
        address 10.10.10.2/24;
      }
    }
  }
[edit]
user@host# show policy-options
policy-statement lan_nw {
  from interface fe-0/0/4.0;
  then accept;
}
[edit]
user@host# show protocols
bgp {
```


Copyright © 2016, Juniper Networks, Inc.

```
    group ibgp {
      type internal;
      local-address 10.10.10.2;
      export lan_nw;
      neighbor 10.10.10.1;
    }
}
[edit]
user@host# show routing-options
static {
    route 1.1.1.0/30 next-hop 2.2.2.2;
    }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy1 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local1;
      }
    }
    gateway spoke-to-hub-gw {
      ike-policy ike-policy1;
      address 1.1.1.1;
      local-identity distinguished-name;
      remote-identity distinguished-name;
      external-interface fe-0/0/1.0;
    }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
    protocol esp;
    authentication-algorithm hmac-md5-96;
    encryption-algorithm des-cbc;
}
    policy vpn-policy1 {
      perfect-forward-secrecy {
        keys group14;
      }
      proposals ipsec-proposal;
    }
    vpn spoke-to-hub {
      bind-interface st0.0;
      ike {
        gateway spoke-to-hub-gw;
        ipsec-policy vpn-policy1;
      }
      establish-tunnels immediately;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security zones
security-zone untrust {
   host-inbound-traffic {
      system-services {
         all;
      }
      protocols {
         all;
      }
   }
   interfaces {
      fe-0/0/1.0;
      st0.0;
   }
}
   security-zone trust {
      host-inbound-traffic {
         system-services {
            all;
         }
         protocols {
            all;
         }
      }
      interfaces {
         fe-0/0/4.0;
      }
   }
[edit]
user@host# show security policies
default-policy {
   permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
   ca-identity ca-profile1;
   enrollment {
      url http://pc4/certsrv/mscep/mscep.dll;
   }
   revocation-check {
      disable;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 2

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/1 unit 0 family inet address 3.3.3.1/30
set interfaces fe-0/0/4 unit 0 family inet address 70.70.70.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.3/24
set policy-options policy-statement lan_nw from interface fe-0/0/4.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.10.10.3
set protocols bgp group ibgp export lan_nw
set protocols bgp group ibgp neighbor 10.10.10.1
set routing-options static route 1.1.1.0/30 next-hop 3.3.3.2
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
set security ike gateway spoke-to-hub-gw ike-policy ike-policy1
set security ike gateway spoke-to-hub-gw address 1.1.1.1
set security ike gateway spoke-to-hub-gw local-identity distinguished-name
set security ike gateway spoke-to-hub-gw remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw external-interface ge-0/0/1.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn spoke-to-hub bind-interface st0.0
set security ipsec vpn spoke-to-hub ike gateway spoke-to-hub-gw
set security ipsec vpn spoke-to-hub ike ipsec-policy vpn-policy1
set security ipsec vpn spoke-to-hub establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 2:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 3.3.3.1/30
user@host# set fe-0/0/4 unit 0 family inet address 70.70.70.1/24
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.3/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface fe-0/0/4.0
user@host# set policy-statement lan_nw then accept

[edit protocols bgp]
user@host# set group ibgp type internal
user@host# set group ibgp local-address 10.10.10.3
user@host# set group ibgp export lan_nw
user@host# set group ibgp neighbor 10.10.10.1

[edit routing-options]
user@host# set static route 1.1.1.0/30 next-hop 3.3.3.2
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike gateway spoke-to-hub-gw]
user@host# set ike-policy ike-policy1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface ge-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub]
user@host# set bind-interface st0.0
user@host# set ike gateway spoke-to-hub-gw
user@host# set ike ipsec-policy vpn-policy1
user@host# set establish-tunnels immediately
```


Copyright © 2016, Juniper Networks, Inc.

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 3.3.3.1/30;
    }
  }
}
  fe-0/0/4 {
    unit 0 {
      family inet {
        address 70.70.70.1/24;
      }
    }
  }
  st0 {
    unit 0 {
      multipoint;
      family inet {
        address 10.10.10.3/24;
      }
    }
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show policy-options
policy-statement lan_nw {
   from interface fe-0/0/4.0;
   then accept;
}
[edit]
user@host# show protocols
bgp {
   group ibgp {
      type internal;
      local-address 10.10.10.3;
      export lan_nw;
      neighbor 10.10.10.1;
   }
}
[edit]
user@host# show routing-options
static {
   route 1.1.1.0/30 next-hop 3.3.3.2;
   }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
   authentication-method rsa-signatures;
   dh-group group2;
   authentication-algorithm sha1;
   encryption-algorithm aes-128-cbc;
}
   policy ike-policy1 {
      mode main;
      proposals ike-proposal;
      certificate {
         local-certificate Local1;
      }
   }
   gateway spoke-to-hub-gw {
      ike-policy ike-policy1;
      address 1.1.1.1;
      local-identity distinguished-name;
      remote-identity distinguished-name;
      external-interface ge-0/0/1.0;
   }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
   protocol esp;
   authentication-algorithm hmac-md5-96;
   encryption-algorithm des-cbc;
}
   policy vpn-policy1 {
      perfect-forward-secrecy {
         keys group14;
      }
      proposals ipsec-proposal;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    vpn spoke-to-hub {
      bind-interface st0.0;
      ike {
        gateway spoke-to-hub-gw;
        ipsec-policy vpn-policy1;
      }
      establish-tunnels immediately;
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/1.0;
      st0.0;
    }
}
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        fe-0/0/4.0;
      }
    }
[edit]
user@host# show security policies
default-policy {
    permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
    ca-identity ca-profile1;
    enrollment {
      url http://pc4/certsrv/mscep/mscep.dll;
    }
    revocation-check {
      disable;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying IKE Phase 1 Status on page 6480
- Verifying IPsec Phase 2 Status on page 6480
- Verifying IPsec Next-Hop Tunnels on page 6481
- Verifying BGP on page 6481
- Verifying Learned Routes on page 6481

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5480163 UP     a558717f387074ab  6d0135c5ecaed61d  Main           3.3.3.1

5480162 UP     7a63d16a5a723df1  c471f7ae166d3a34  Main           2.2.2.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spokes.

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **security ipsec security-associations** command.

```
user@host> security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <268173400 ESP:des/ md5 9bf33bc7 3567/ unlim -   root 500   2.2.2.1
  >268173400 ESP:des/ md5 aae5196b 3567/ unlim -   root 500   2.2.2.1
  <268173401 ESP:des/ md5 69c24d81 622/ unlim  -   root 500   3.3.3.1
  >268173401 ESP:des/ md5 e3fe0231 622/ unlim  -   root 500   3.3.3.1
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spokes.

Copyright © 2016, Juniper Networks, Inc.

Verifying IPsec Next-Hop Tunnels

Purpose Verify the IPsec next-hop tunnels.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPSec VPN name                    Flag     IKE-ID
                              XAUTH username
10.10.10.2        st0.0       hub-to-spoke-vpn                  Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
10.10.10.3        st0.0       hub-to-spoke-vpn                  Auto     C=IN,
DC=example.net, ST=KA, L=Tumkur, O=example, OU=SLT, CN=spoke2
```

Meaning The next-hop gateways are the IP addresses for the **st0** interfaces of the spokes. The next hop should be associated with the correct IPsec VPN name.

Verifying BGP

Purpose Verify that BGP references the IP addresses for the **st0** interfaces of the spokes.

Action From operational mode, enter the **show bgp summary** command.

```
user@host> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Unconfigured peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2               10        116        119       0       0       50:25
1/1/1/0              0/0/0/0
10.10.10.3               10        114        114       0       0       50:04
1/1/1/0              0/0/0/0
```

Verifying Learned Routes

Purpose Verify that routes to the spokes have been learned.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route 60.60.60.0** command.

```
user@host> show route 60.60.60.0
inet.0: 45 destinations, 45 routes (44 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

60.60.60.0/24      *[BGP/170] 00:50:57, localpref 100
                      AS path: I
                    > to 10.10.10.2 via st0.0
```

From operational mode, enter the **show route 70.70.70.0** command.

```
user@host> show route 70.70.70.0
inet.0: 45 destinations, 45 routes (44 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

70.70.70.0/24      *[BGP/170] 00:50:42, localpref 100
                      AS path: I
                    > to 10.10.10.3 via st0.0
```

Related Documentation

- Example: Configuring a Route-Based VPN on page 6074
- *Routing Protocols Overview for Security Devices*

Example: Configuring Basic AutoVPN with OSPF

This example shows how to configure an AutoVPN hub to act as a single termination point, and then configure two spokes to act as tunnels to remote sites. This example configures OSPF to forward packets through the VPN tunnels.

- Requirements on page 6482
- Overview on page 6483
- Configuration on page 6485
- Verification on page 6504

Requirements

This example uses the following hardware and software components:

- Three supported SRX Series devices as AutoVPN hub and spokes
- Junos OS Release 12.1X44-D10 and later that support AutoVPN

Copyright © 2016, Juniper Networks, Inc.

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates.

NOTE: You should be familiar with the dynamic routing protocol that is used to forward packets through the VPN tunnels. For more information about specific requirements for a dynamic routing protocol, see the *Routing Protocols Overview for Security Devices*.

Overview

This example shows the configuration of an AutoVPN hub and the subsequent configurations of two spokes.

In this example, the first step is to enroll digital certificates in each device using the Simple Certificate Enrollment Protocol (SCEP). The certificates for the spokes contain the organizational unit (OU) value “SLT” in the subject field; the hub is configured with a group IKE ID to match the value “SLT” in the OU field.

The spokes establish IPsec VPN connections to the hub, which allows them to communicate with each other as well as access resources on the hub. Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hub and all spokes must have the same values. Table 557 shows the options used in this example.

Table 557: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke Basic OSPF Configurations

Option	Value
IKE proposal:	
Authentication method	RSA digital certificates
Diffie-Hellman (DH) group	2
Authentication algorithm	SHA-1
Encryption algorithm	AES 128 CBC
IKE policy:	
Mode	Main
IPsec proposal:	
Protocol	ESP
Authentication algorithm	HMAC MD5 96
Encryption algorithm	DES CBC

Copyright © 2016, Juniper Networks, Inc.

Table 557: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke Basic OSPF Configurations *(continued)*

Option	Value
IPsec policy:	
Perfect Forward Secrecy (PFS) group	14

The same certificate authority (CA) is configured on all devices.

NOTE: Junos OS only supports a single level of certificate hierarchy.

Table 558 shows the options configured on the hub and on all spokes.

Table 558: AutoVPN Basic OSPF Configuration for Hub and All Spokes

Option	Hub	All Spokes
IKE gateway:		
Remote IP address	Dynamic	1.1.1.1
Remote IKE ID	Distinguished name (DN) on the spoke's certificate with the string **SLT** in the organizational unit (OU) field	DN on the hub's certificate
Local IKE ID	DN on the hub's certificate	DN on the spoke's certificate
External interface	ge-0/0/1.0	Spoke 1: fe-0/0/1.0 Spoke 2: ge-0/0/1.0
VPN:		
Bind interface	st0.0	st0.0
Establish tunnels	(not configured)	Immediately on configuration commit

Table 559 shows the configuration options that are different on each spoke.

Table 559: Comparison Between the Basic OSPF Spoke Configurations

Option	Spoke 1	Spoke 2
st0.0 interface	10.10.10.2/24	10.10.10.3/24
Interface to internal network	fe-0.0/4.0: 60.60.60.1/24	fe-0.0/4.0: 70.70.70.1/24
Interface to Internet	fe-0/0/1.0: 2.2.2.1/30	ge-0/0/1.0: 3.3.3.1/30

Copyright © 2016, Juniper Networks, Inc.

Routing information for all devices is exchanged through the VPN tunnels.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Topology

Figure 284 shows the SRX Series devices to be configured for AutoVPN in this example.

Figure 284: Basic AutoVPN Deployment with OSPF

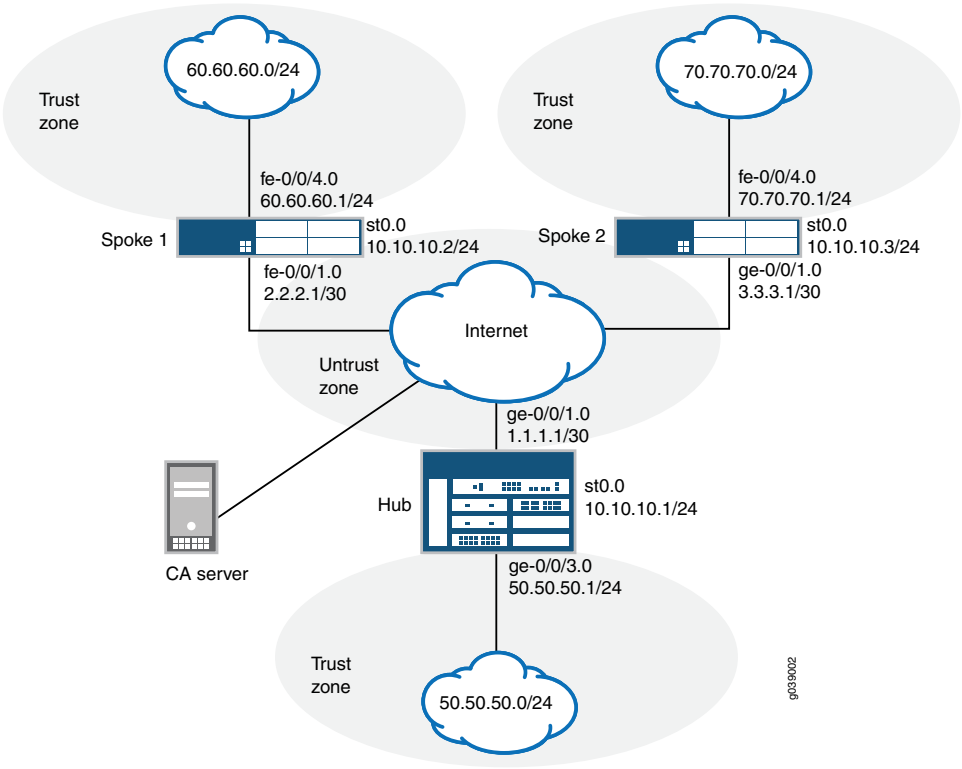

Configuration

To configure AutoVPN, perform these tasks:

NOTE: The first section describes how to obtain CA and local certificates online using the Simple Certificate Enrollment Protocol (SCEP) on the hub and spoke devices.

- Enroll Device Certificates with SCEP on page 6486
- Configuring the Hub on page 6489

Copyright © 2016, Juniper Networks, Inc.

- Configuring Spoke 1 on page 6494
- Configuring Spoke 2 on page 6499

Enroll Device Certificates with SCEP

Step-by-Step Procedure

To enroll digital certificates with SCEP on the hub:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

```
user@host> request security pki generate-key-pair certificate-id Local1
```

4. Enroll the local certificate.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email hub@example.net ip-address
  1.1.1.1 subject
  DC=example.net,CN=hub,OU=SLT,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificate.

```
user@host> show security pki local-certificate detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a6d5f300000000258d
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Bangalore, Common name: hub, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SLT, CN=hub
  Alternate subject: "hub@example.net", example.net, 1.1.1.1
  Validity:
    Not before: 11- 6-2012 09:39
    Not after: 11- 6-2013 09:49
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:c9:c9:cc:30:b6:7a:86:12:89:b5:18:b3:76
    01:2d:cc:65:a8:a8:42:78:cd:d0:9a:a2:c0:aa:c4:bd:da:af:88:f3
    2a:78:1f:0a:58:e6:11:2c:81:8f:0e:7c:de:86:fc:48:4c:28:5b:8b
    34:91:ff:2e:91:e7:b5:bd:79:12:de:39:46:d9:fb:5c:91:41:d1:da
    90:f5:09:00:9b:90:07:9d:50:92:7d:ff:fb:3f:3c:bc:34:e7:e3:c8
    ea:cb:99:18:b4:b6:1d:a8:99:d3:36:b9:1b:36:ef:3e:a1:fd:48:82
```


Copyright © 2016, Juniper Networks, Inc.

```
    6a:da:22:07:da:e0:d2:55:ef:57:be:09:7a:0e:17:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    e1:f7:a1:a6:1e:c3:97:69:a5:07:9b:09:14:1a:c7:ae:09:f1:f6:35 (sha1)
    a0:02:fa:8d:5c:63:e5:6d:f7:f4:78:56:ac:4e:b2:c4 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 1:

1. Configure the CA.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
   user@host# set security pki ca-profile ca-profile1 enrollment url
     http://pc4/certsrv/mscep/mscep.dll
   user@host# set security pki ca-profile ca-profile1 revocation-check disable
   user@host# commit
   ```

2. Enroll the CA certificate.

   ```
   user@host> request security pki ca-certificate enroll ca-profile ca-profile1
   ```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

   ```
   user@host> request security pki generate-key-pair certificate-id Local1
   ```

4. Enroll the local certificate.

   ```
   user@host> request security pki local-certificate enroll ca-profile ca-profile1
     certificate-id Local1 domain-name example.net email spoke1@example.net
     ip-address 2.2.2.1 subject
     DC=example.net,CN=spoke1,OU=SLT,O=example,L=Mysore,ST=KA,C=IN
     challenge-password <password>
   ```

5. Verify the local certificate.

   ```
   user@host> show security pki local-certificate detail

   Certificate identifier: Local1
     Certificate version: 3
     Serial number: 40a7975f00000000258e
     Issuer:
       Common name: CASERVER1, Domain component: net, Domain component: internal

     Subject:
       Organization: example, Organizational unit: SLT, Country: IN, State: KA,

       Locality: Mysore, Common name: spoke1, Domain component: example.net
     Subject string:
       C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
     Alternate subject: "spoke1@example.net", example.net, 2.2.2.1
     Validity:
       Not before: 11- 6-2012 09:40
       Not after: 11- 6-2013 09:50
     Public key algorithm: rsaEncryption(1024 bits)
   ```

Copyright © 2016, Juniper Networks, Inc.

```
    30:81:89:02:81:81:00:d8:45:09:77:cd:36:9a:6f:58:44:18:91:db
    b0:c7:8a:ee:c8:d7:a6:d2:e2:e7:20:46:2b:26:1a:92:e2:4e:8a:ce
    c9:25:d9:74:a2:81:ad:ea:e0:38:a0:2f:2d:ab:a6:58:ac:88:35:f4
    90:01:08:33:33:75:2c:44:26:f8:25:18:97:96:e4:28:de:3b:35:f2
    4a:f5:92:b7:57:ae:73:4f:8e:56:71:ab:81:54:1d:75:88:77:13:64
    1b:6b:01:96:15:0a:1c:54:e3:db:f8:ec:ec:27:5b:86:39:c1:09:a1
    e4:24:1a:19:0d:14:2c:4b:94:a4:04:91:3f:cb:ef:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    b6:24:2a:0e:96:5d:8c:4a:11:f3:5a:24:89:7c:df:ea:d5:c0:80:56 (sha1)
    31:58:7f:15:bb:d4:66:b8:76:1a:42:4a:8a:16:b3:a9 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT. The IKE configuration on the hub includes ou=SLT to identify the spoke.

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 2:

1. Configure the CA.

   ```
   [edit]
   user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
   user@host# set security pki ca-profile ca-profile1 enrollment url
     http://pc4/certsrv/mscep/mscep.dll
   user@host# set security pki ca-profile ca-profile1 revocation-check disable
   user@host# commit
   ```

2. Enroll the CA certificate.

   ```
   user@host> request security pki ca-certificate enroll ca-profile ca-profile1
   ```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair.

   ```
   user@host> request security pki generate-key-pair certificate-id Local1
   ```

4. Enroll the local certificate.

   ```
   user@host> request security pki local-certificate enroll ca-profile ca-profile1
     certificate-id Local1 domain-name example.net email spoke2@example.net
     ip-address 3.3.3.1 subject
     DC=example.net,CN=spoke2,OU=SLT,O=example,L=Tumkur,ST=KA,C=IN
     challenge-password <password>
   ```

5. Verify the local certificate.

   ```
   user@host> show security pki local-certificate detail

   Certificate identifier: Local1
     Certificate version: 3
     Serial number: 40bb71d400000000258f
     Issuer:
   ```


Copyright © 2016, Juniper Networks, Inc.

```
  Common name: CASERVER1, Domain component: net, Domain component: internal

Subject:
  Organization: example, Organizational unit: SLT, Country: IN, State: KA,

  Locality: Tumkur, Common name: spoke2, Domain component: example.net
Subject string:
  C=IN, DC=example.net, ST=KA, L=Tumkur, O=example, OU=SLT, CN=spoke2
Alternate subject: "spoke2@example.net", example.net, 3.3.3.1
Validity:
  Not before: 11- 6-2012 10:02
  Not after: 11- 6-2013 10:12
Public key algorithm: rsaEncryption(1024 bits)
  30:81:89:02:81:81:00:b6:2e:e2:da:e6:ac:57:e4:5d:ff:de:f6:89
  27:d6:3e:1b:4a:3f:b2:2d:b3:d3:61:ed:ed:6a:07:d9:8a:d2:24:03
  77:1a:fe:84:e1:12:8a:2d:63:6e:bf:02:6b:15:96:5a:4f:37:a0:46
  44:09:96:c0:fd:bb:ab:79:2c:5d:92:bd:31:f0:3b:29:51:ce:89:8e
  7c:2b:02:d0:14:5b:0a:a9:02:93:21:ea:f9:fc:4a:e7:08:bc:b1:6d
  7c:f8:3e:53:58:8e:f1:86:13:fe:78:b5:df:0b:8e:53:00:4a:46:11
  58:4a:38:e9:82:43:d8:25:47:7d:ef:18:f0:ef:a7:02:03:01:00:01
Signature algorithm: sha1WithRSAEncryption
Distribution CRL:
  http://ca-server1/CertEnroll/CASERVER1.crl
  file://\\ca-server1\CertEnroll\CASERVER1.crl
Fingerprint:
  1a:6d:77:ac:fd:94:68:ce:cf:8a:85:f0:39:fc:e0:6b:fd:fe:b8:66 (sha1)
  00:b1:32:5f:7b:24:9c:e5:02:e6:72:75:9e:a5:f4:77 (md5)
Auto-re-enrollment:
  Status: Disabled
  Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT. The IKE configuration on the hub includes ou=SLT to identify the spoke.

Configuring the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/3 unit 0 family inet address 50.50.50.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.1/24
set protocols ospf area 0.0.0.0 interface st0.0 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.0 dynamic-neighbors
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set routing-options static route 2.2.2.0/30 next-hop 1.1.1.2
set routing-options static route 3.3.3.0/30 next-hop 1.1.1.2
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
set security ike gateway hub-to-spoke-gw ike-policy ike-policy1
set security ike gateway hub-to-spoke-gw dynamic distinguished-name wildcard OU=SLT
set security ike gateway hub-to-spoke-gw dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw local-identity distinguished-name
set security ike gateway hub-to-spoke-gw external-interface ge-0/0/1.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn hub-to-spoke-vpn bind-interface st0.0
set security ipsec vpn hub-to-spoke-vpn ike gateway hub-to-spoke-gw
set security ipsec vpn hub-to-spoke-vpn ike ipsec-policy vpn-policy1
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the hub:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 1.1.1.1/30
user@host# set ge-0/0/3 unit 0 family inet address 50.50.50.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.1/24
```

2. Configure the routing protocol.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface st0.0 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.0 dynamic-neighbors
user@host# set area 0.0.0.0 interface ge-0/0/3.0


[edit routing-options]
user@host# set static route 2.2.2.0/30 next-hop 1.1.1.2
user@host# set static route 3.3.3.0/30 next-hop 1.1.1.2
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike gateway hub-to-spoke-gw]
user@host# set ike-policy ike-policy1
user@host# set dynamic distinguished-name wildcard OU=SLT
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn hub-to-spoke-vpn]
user@host# set bind-interface st0.0
user@host# set ike gateway hub-to-spoke-gw
user@host# set ike ipsec-policy vpn-policy1
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/3.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 1.1.1.1/30;
    }
  }
}
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 50.50.50.1/24;
      }
    }
  }
  st0 {
    unit 0 {
      multipoint;
      family inet {
        address 10.10.10.1/24;
      }
    }
  }
[edit]
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface st0.0 {
      interface-type p2mp;
      dynamic-neighbors;
    }
    interface ge-0/0/3.0;
  }
}
[edit]
user@host# show routing-options
static {
  route 2.2.2.0/30 next-hop 1.1.1.2;
  route 3.3.3.0/30 next-hop 1.1.1.2;
}
[edit]
user@host# show security ike
proposal ike-proposal {
  authentication-method rsa-signatures;
  dh-group group2;
  authentication-algorithm sha1;
  encryption-algorithm aes-128-cbc;
}
```


Copyright © 2016, Juniper Networks, Inc.

```
    policy ike-policy1 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local1;
      }
    }
    gateway hub-to-spoke-gw {
      ike-policy ike-policy1;
      dynamic {
        distinguished-name {
          wildcard OU=SLT;
        }
        ike-user-type group-ike-id;
      }
      local-identity distinguished-name;
      external-interface ge-0/0/1.0;
    }
[edit]
user@host# show security ipsec
traceoptions {
    flag all;
}
    proposal ipsec-proposal {
      protocol esp;
      authentication-algorithm hmac-md5-96;
      encryption-algorithm des-cbc;
    }
    policy vpn-policy1 {
      perfect-forward-secrecy {
        keys group14;
      }
      proposals ipsec-proposal;
    }
    vpn hub-to-spoke-vpn {
      bind-interface st0.0;
      ike {
        gateway hub-to-spoke-gw;
        ipsec-policy vpn-policy1;
      }
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      st0.0;
      ge-0/0/1.0;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/3.0;
    }
  }
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/1 unit 0 family inet address 2.2.2.1/30
set interfaces fe-0/0/4 unit 0 family inet address 60.60.60.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.2/24
set protocols ospf area 0.0.0.0 interface st0.0 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.0 neighbor 10.10.10.1
set protocols ospf area 0.0.0.0 interface fe-0/0/4.0
set routing-options static route 1.1.1.0/30 next-hop 2.2.2.2
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ike gateway spoke-to-hub-gw ike-policy ike-policy1
set security ike gateway spoke-to-hub-gw address 1.1.1.1
set security ike gateway spoke-to-hub-gw local-identity distinguished-name
set security ike gateway spoke-to-hub-gw remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw external-interface fe-0/0/1.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn spoke-to-hub bind-interface st0.0
set security ipsec vpn spoke-to-hub ike gateway spoke-to-hub-gw
set security ipsec vpn spoke-to-hub ike ipsec-policy vpn-policy1
set security ipsec vpn spoke-to-hub establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 1:

1. Configure interfaces.

```
[edit interfaces]
user@host# set fe-0/0/1 unit 0 family inet address 2.2.2.1/30
user@host# set fe-0/0/4 unit 0 family inet address 60.60.60.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.2/24
```

2. Configure the routing protocol.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface st0.0 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.0 neighbor 10.10.10.1
user@host# set area 0.0.0.0 interface fe-0/0/4.0

[edit routing-options]
user@host# set static route 1.1.1.0/30 next-hop 2.2.2.2
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike gateway spoke-to-hub-gw]
user@host# set ike-policy ike-policy1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub]
user@host# set bind-interface st0.0
user@host# set ike gateway spoke-to-hub-gw
user@host# set ike ipsec-policy vpn-policy1
user@host# set establish-tunnels immediately
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 2.2.2.1/30;
    }
  }
}
  fe-0/0/4 {
    unit 0 {
      family inet {
        address 60.60.60.1/24;
      }
    }
  }
  st0 {
    unit 0 {
      multipoint;
      family inet {
        address 10.10.10.2/24;
      }
    }
  }
[edit]
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface st0.0 {
      interface-type p2mp;
      neighbor 10.10.10.1;
    }
    interface fe-0/0/4.0;
  }
}
[edit]
user@host# show routing-options
static {
  route 1.1.1.0/30 next-hop 2.2.2.2;
}
[edit]
user@host# show security ike
proposal ike-proposal {
  authentication-method rsa-signatures;
  dh-group group2;
  authentication-algorithm sha1;
  encryption-algorithm aes-128-cbc;
}
  policy ike-policy1 {
```

Copyright © 2016, Juniper Networks, Inc.

```
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local1;
      }
    }
    gateway spoke-to-hub-gw {
      ike-policy ike-policy1;
      address 1.1.1.1;
      local-identity distinguished-name;
      remote-identity distinguished-name;
      external-interface fe-0/0/1.0;
    }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
  protocol esp;
  authentication-algorithm hmac-md5-96;
  encryption-algorithm des-cbc;
}
  policy vpn-policy1 {
    perfect-forward-secrecy {
      keys group14;
    }
    proposals ipsec-proposal;
  }
  vpn spoke-to-hub {
    bind-interface st0.0;
    ike {
      gateway spoke-to-hub-gw;
      ipsec-policy vpn-policy1;
    }
    establish-tunnels immediately;
  }
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
      all;
    }
  }
  interfaces {
    fe-0/0/1.0;
    st0.0;
  }
}
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
```


Copyright © 2016, Juniper Networks, Inc.

```
            all;
          }
        }
        interfaces {
          fe-0/0/4.0;
        }
      }
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 2

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 3.3.3.1/30
set interfaces fe-0/0/4 unit 0 family inet address 70.70.70.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.3/24
set protocols ospf area 0.0.0.0 interface st0.0 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.0 neighbor 10.10.10.1
set protocols ospf area 0.0.0.0 interface fe-0/0/4.0
set routing-options static route 1.1.1.1/32 next-hop 3.3.3.2
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy1 mode main
set security ike policy ike-policy1 proposals ike-proposal
set security ike policy ike-policy1 certificate local-certificate Local1
set security ike gateway spoke-to-hub-gw ike-policy ike-policy1
set security ike gateway spoke-to-hub-gw address 1.1.1.1
set security ike gateway spoke-to-hub-gw local-identity distinguished-name
set security ike gateway spoke-to-hub-gw remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw external-interface ge-0/0/1.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy1 proposals ipsec-proposal
set security ipsec vpn spoke-to-hub bind-interface st0.0
set security ipsec vpn spoke-to-hub ike gateway spoke-to-hub-gw
set security ipsec vpn spoke-to-hub ike ipsec-policy vpn-policy1
set security ipsec vpn spoke-to-hub establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 2:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 3.3.3.1/30
user@host# set fe-0/0/4 unit 0 family inet address 70.70.70.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.3/24
```

2. Configure the routing protocol.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface st0.0 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.0 neighbor 10.10.10.1
user@host# set area 0.0.0.0 interface fe-0/0/4.0

[edit routing-options]
user@host# set static route 1.1.1.1/32 next-hop 3.3.3.2
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike gateway spoke-to-hub-gw]
user@host# set ike-policy ike-policy1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface ge-0/0/1.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy1]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub]
user@host# set bind-interface st0.0
user@host# set ike gateway spoke-to-hub-gw
user@host# set ike ipsec-policy vpn-policy1
user@host# set establish-tunnels immediately
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces st0.0

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            address 3.3.3.1/30;
        }
    }
}
    fe-0/0/4 {
        unit 0 {
            family inet {
                address 70.70.70.1/24;
            }
        }
    }
    st0 {
        unit 0 {
            multipoint;
            family inet {
                address 10.10.10.3/24;
            }
        }
    }
[edit]
user@host# show protocols
ospf {
    area 0.0.0.0 {
        interface st0.0 {
            interface-type p2mp;
            neighbor 10.10.10.1;
        }
        interface fe-0/0/4.0;
    }
}
[edit]
user@host# show routing-options
static {
    route 1.1.1.1/32 next-hop 3.3.3.2;
}
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy1 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local1;
        }
    }
    gateway spoke-to-hub-gw {
        ike-policy ike-policy1;
```


Copyright © 2016, Juniper Networks, Inc.

```
    address 1.1.1.1;
    local-identity distinguished-name;
    remote-identity distinguished-name;
    external-interface ge-0/0/1.0;
  }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
  protocol esp;
  authentication-algorithm hmac-md5-96;
  encryption-algorithm des-cbc;
}
  policy vpn-policy1 {
    perfect-forward-secrecy {
      keys group14;
    }
    proposals ipsec-proposal;
  }
  vpn spoke-to-hub {
    bind-interface st0.0;
    ike {
      gateway spoke-to-hub-gw;
      ipsec-policy vpn-policy1;
    }
    establish-tunnels immediately;
  }
[edit]
user@host# show security zones
security-zone untrust {
  host-inbound-traffic {
    system-services {
      all;
    }
    protocols {
      all;
    }
  }
  interfaces {
    ge-0/0/1.0;
    st0.0;
  }
}
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      fe-0/0/4.0;
    }
  }
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security policies
default-policy {
    permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
    ca-identity ca-profile1;
    enrollment {
        url http://pc4/certsrv/mscep/mscep.dll;
    }
    revocation-check {
        disable;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying IKE Phase 1 Status on page 6504
- Verifying IPsec Phase 2 Status on page 6504
- Verifying IPsec Next-Hop Tunnels on page 6505
- Verifying OSPF on page 6505
- Verifying Learned Routes on page 6505

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5480159 UP     22432fb6f7fbc389  412b751f79b45099  Main           2.2.2.1

5480161 UP     d455050707bc3eaf  b3dde111232270d2  Main           3.3.3.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spokes.

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **security ipsec security-associations** command.

```
user@host> security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <268173400 ESP:des/ md5 f38eea12 2954/ unlim -   root 500   2.2.2.1
  >268173400 ESP:des/ md5 bb48d228 2954/ unlim -   root 500   2.2.2.1
  <268173401 ESP:des/ md5 bcd1390b 3530/ unlim -   root 500   3.3.3.1
  >268173401 ESP:des/ md5 77fcf6e2 3530/ unlim -   root 500   3.3.3.1
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spokes.

Verifying IPsec Next-Hop Tunnels

Purpose Verify the IPsec next-hop tunnels.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPSec VPN name                    Flag     IKE-ID
                            XAUTH username
10.10.10.2        st0.0       hub-to-spoke-vpn                  Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
10.10.10.3        st0.0       hub-to-spoke-vpn                  Auto     C=IN,
DC=example.net, ST=KA, L=Tumkur, O=example, OU=SLT, CN=spoke2
```

Meaning The next-hop gateways are the IP addresses for the **st0** interfaces of the spokes. The next hop should be associated with the correct IPsec VPN name.

Verifying OSPF

Purpose Verify that OSPF references the IP addresses for the **st0** interfaces of the spokes.

Action From operational mode, enter the **show ospf neighbor** command.

```
user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.10.10.3       st0.0                  Full      10.255.226.179   128    32
10.10.10.2       st0.0                  Full      10.207.36.182    128    38
```

Verifying Learned Routes

Purpose Verify that routes to the spokes have been learned.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route 60.60.60.0** command.

```
user@host> show route 60.60.60.0
 inet.0: 48 destinations, 48 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

60.60.60.0/24      *[OSPF/10] 00:51:13, metric 2
                    > to 10.10.10.2 via st0.0
```

From operational mode, enter the **show route 70.70.70.0** command.

```
user@host> show route 70.70.70.0
inet.0: 48 destinations, 48 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

70.70.70.0/24      *[OSPF/10] 00:51:48, metric 2
                    > to 10.10.10.3 via st0.0
```

Related Documentation

- Example: Configuring a Route-Based VPN on page 6074
- *Routing Protocols Overview for Security Devices*

Example: Configuring AutoVPN with iBGP and ECMP

This example shows how to configure two IPsec VPN tunnels between an AutoVPN hub and spoke. This example configures iBGP with equal-cost multipath (ECMP) to forward packets through the VPN tunnels.

- Requirements on page 6506
- Overview on page 6507
- Configuration on page 6509
- Verification on page 6528

Requirements

This example uses the following hardware and software components:

- Two supported SRX Series devices as AutoVPN hub and spoke
- Junos OS Release 12.1X44-D10 and later that support AutoVPN

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates.

NOTE: You should be familiar with the dynamic routing protocol that is used to forward packets through the VPN tunnels. For more information about specific requirements for a dynamic routing protocol, see the *Routing Protocols Overview for Security Devices*.

Copyright © 2016, Juniper Networks, Inc.

Overview

This example shows the configuration of an AutoVPN hub and a spoke with two IPsec VPN tunnels.

In this example, the first step is to enroll digital certificates in each device using the Simple Certificate Enrollment Protocol (SCEP). Certificates are enrolled in the hub and in the spoke for each IPsec VPN tunnel. One of the certificates for the spoke contains the organizational unit (OU) value "SLT" in the distinguished name (DN); the hub is configured with a group IKE ID to match the value "SLT" in the OU field. The other certificate for the spoke contains the OU value "SBU" in the DN; the hub is configured with a group IKE ID to match the value "SBU" in the OU field.

The spoke establishes IPsec VPN connections to the hub, which allows it to access resources on the hub. Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hub and the spoke must have the same values. Table 560 shows the options used in this example.

Table 560: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke iBGP ECMP Configurations

Option	Value
IKE proposal:	
Authentication method	RSA digital certificates
Diffie-Hellman (DH) group	2
Authentication algorithm	SHA-1
Encryption algorithm	AES 128 CBC
IKE policy:	
Mode	Main
IPsec proposal:	
Protocol	ESP
Authentication algorithm	HMAC MD5 96
Encryption algorithm	DES CBC
IPsec policy:	
Perfect Forward Secrecy (PFS) group	14

The same certificate authority (CA) is configured on all devices.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Junos OS only supports a single level of certificate hierarchy.

Table 561 shows the options configured on the hub and on the spoke.

Table 561: AutoVPN iBGP ECMP Configuration for Hub and Spoke 1

Option	Hub	Spoke 1
IKE gateway:		
Remote IP address	hub-to-spoke-gw-1: Dynamic hub-to-spoke-gw-2: Dynamic	spoke-to-hub-gw-1: 1.1.1.1 spoke-to-hub-gw-2: 1.1.2.1
Remote IKE ID	hub-to-spoke-gw-1: DN on the spoke's certificate with the string **SLT** in the OU field hub-to-spoke-gw-2: DN on the spoke's certificate with the string **SBU** in the OU field	spoke-to-hub-gw-1: DN on the hub's certificate spoke-to-hub-gw-2: DN on the hub's certificate
Local IKE ID	DN on the hub's certificate	DN on the spoke's certificate
External interface	hub-to-spoke-gw-1: ge-0/0/1.0 hub-to-spoke-gw-2: ge-0/0/2.0	spoke-to-hub-gw-1: fe-0/0/1.0 spoke-to-hub-gw-2: fe-0/0/2.0
VPN:		
Bind interface	hub-to-spoke-vpn-1: st0.0 hub-to-spoke-vpn-2: st0.1	spoke-to-hub-1: st0.0 spoke-to-hub-2: st0.1
Establish tunnels	(not configured)	Immediately on configuration commit

Routing information for all devices is exchanged through the VPN tunnels.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Topology

Figure 285 shows the SRX Series devices to be configured for AutoVPN in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 285: AutoVPN Deployment with iBGP and ECMP

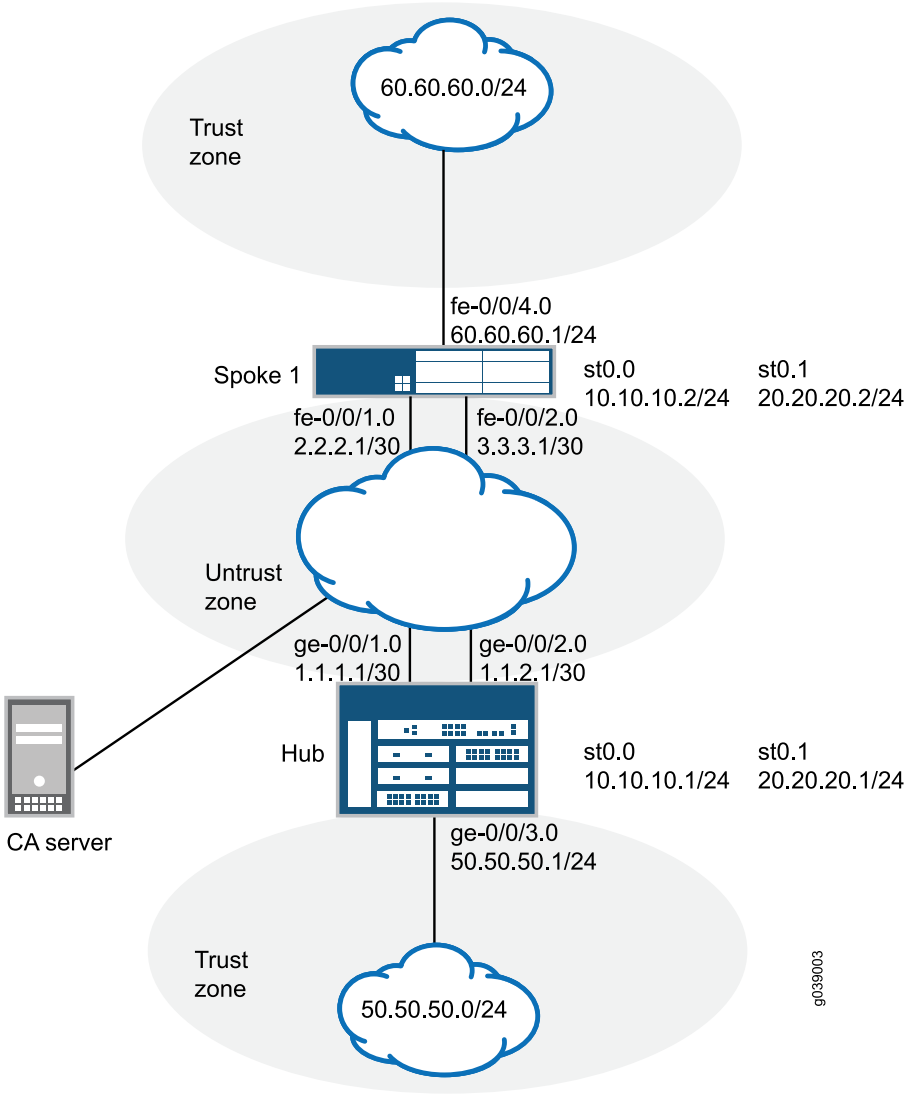

Configuration

To configure AutoVPN, perform these tasks:

NOTE: The first section describes how to obtain CA and local certificates online using the Simple Certificate Enrollment Protocol (SCEP) on the hub and spoke devices.

- Enroll Device Certificates with SCEP on page 6510
- Configuring the Hub on page 6513
- Configuring Spoke 1 on page 6521

Copyright © 2016, Juniper Networks, Inc.

Enroll Device Certificates with SCEP

Step-by-Step Procedure

To enroll digital certificates with SCEP on the hub:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair for each certificate.

```
user@host> request security pki generate-key-pair certificate-id Local1
user@host> request security pki generate-key-pair certificate-id Local2
```

4. Enroll the local certificates.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email hub@example.net ip-address
  1.1.1.1 subject
  DC=example.net,CN=hub,OU=SLT,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local2 domain-name example.net email hub_backup@example.net
  ip-address 1.1.2.1 subject
  DC=example.net,CN=hub_backup,OU=SBU,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificates.

```
user@host> show security pki local-certificate certificate-id Local1 detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a6d5f300000000258d
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Bangalore, Common name: hub, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SLT, CN=hub
  Alternate subject: "hub@example.net", example.net, 1.1.1.1
  Validity:
    Not before: 11- 6-2012 09:39
    Not after: 11- 6-2013 09:49
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:c9:c9:cc:30:b6:7a:86:12:89:b5:18:b3:76
    01:2d:cc:65:a8:a8:42:78:cd:d0:9a:a2:c0:aa:c4:bd:da:af:88:f3
    2a:78:1f:0a:58:e6:11:2c:81:8f:0e:7c:de:86:fc:48:4c:28:5b:8b
```


Copyright © 2016, Juniper Networks, Inc.

```
    34:91:ff:2e:91:e7:b5:bd:79:12:de:39:46:d9:fb:5c:91:41:d1:da
    90:f5:09:00:9b:90:07:9d:50:92:7d:ff:fb:3f:3c:bc:34:e7:e3:c8
    ea:cb:99:18:b4:b6:1d:a8:99:d3:36:b9:1b:36:ef:3e:a1:fd:48:82
    6a:da:22:07:da:e0:d2:55:ef:57:be:09:7a:0e:17:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    e1:f7:a1:a6:1e:c3:97:69:a5:07:9b:09:14:1a:c7:ae:09:f1:f6:35 (sha1)
    a0:02:fa:8d:5c:63:e5:6d:f7:f4:78:56:ac:4e:b2:c4 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started

user@host> show security pki local-certificate certificate-id Local2 detail

Certificate identifier: Local2
  Certificate version: 3
  Serial number: 505efdf900000000259a
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SBU, Country: IN, State: KA,

    Locality: Bangalore, Common name: hub_backup, Domain component:
example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SBU, CN=hub_backup

  Alternate subject: "hub_backup@example.net", example.net, 1.1.2.1
  Validity:
    Not before: 11- 9-2012 10:55
    Not after: 11- 9-2013 11:05
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:d5:44:08:96:f6:77:05:e6:91:50:8a:8a:2a
    4e:95:43:1e:88:ea:43:7c:c5:ac:88:d7:a0:8d:b5:d9:3f:41:db:db
    44:34:1f:56:a5:38:4b:b2:c5:85:f9:f1:bf:b2:7b:d4:b2:af:98:a0
    95:50:02:ad:f5:dd:4d:dc:67:85:dd:84:09:df:9c:68:a5:58:65:e7
    2c:72:cc:47:4b:d0:cc:4a:28:ca:09:db:ad:6e:5a:13:6c:e6:cc:f0
    29:ed:2b:2d:d1:38:38:bc:68:84:de:ae:86:39:c9:dd:06:d5:36:f0
    e6:2a:7b:46:4c:cd:a5:24:1c:e0:92:8d:ad:35:29:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    98:96:2f:ff:ca:af:33:ee:d7:4c:c8:4f:f7:71:53:c0:5d:5f:c5:59 (sha1)
    c9:87:e3:a4:5c:47:b5:aa:90:22:e3:06:b2:0b:e1:ea (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 1:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
```

```
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair for each certificate.

```
user@host> rrequest security pki generate-key-pair certificate-id Local1
user@host> request security pki generate-key-pair certificate-id Local2
```

4. Enroll the local certificates.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email spoke1@example.net
  ip-address 2.2.2.1 subject
  DC=example.net,CN=spoke1,OU=SLT,O=example,L=Mysore,ST=KA,C=IN
  challenge-password <password>
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local2 domain-name example.net email
  spoke1_backup@example.net ip-address 3.3.3.1 subject
  DC=example.net,CN=spoke1_backup,OU=SBU,O=example,L=Mysore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificates.

```
user@host> show security pki local-certificate certificate-id Local1 detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a7975f00000000258e
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Mysore, Common name: spoke1, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
  Alternate subject: "spoke1@example.net", example.net, 2.2.2.1
  Validity:
    Not before: 11- 6-2012 09:40
    Not after: 11- 6-2013 09:50
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:d8:45:09:77:cd:36:9a:6f:58:44:18:91:db
    b0:c7:8a:ee:c8:d7:a6:d2:e2:e7:20:46:2b:26:1a:92:e2:4e:8a:ce
    c9:25:d9:74:a2:81:ad:ea:e0:38:a0:2f:2d:ab:a6:58:ac:88:35:f4
    90:01:08:33:33:75:2c:44:26:f8:25:18:97:96:e4:28:de:3b:35:f2
    4a:f5:92:b7:57:ae:73:4f:8e:56:71:ab:81:54:1d:75:88:77:13:64
    1b:6b:01:96:15:0a:1c:54:e3:db:f8:ec:ec:27:5b:86:39:c1:09:a1
    e4:24:1a:19:0d:14:2c:4b:94:a4:04:91:3f:cb:ef:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
```


Copyright © 2016, Juniper Networks, Inc.

```
    b6:24:2a:0e:96:5d:8c:4a:11:f3:5a:24:89:7c:df:ea:d5:c0:80:56 (sha1)
    31:58:7f:15:bb:d4:66:b8:76:1a:42:4a:8a:16:b3:a9 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

```
user@host> show security pki local-certificate certificate-id Local2 detail

Certificate identifier: Local2
  Certificate version: 3
  Serial number: 506c3d0600000000259b
  Issuer:
   Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
   Organization: example, Organizational unit: SBU, Country: IN, State: KA,

    Locality: Mysore, Common name: spoke1_backup, Domain component:
example.net
  Subject string:
   C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SBU, CN=spoke1_backup

  Alternate subject: "spoke1_backup@example.net", example.net, 3.3.3.1
  Validity:
    Not before: 11- 9-2012 11:09
    Not after: 11- 9-2013 11:19
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:a7:02:b5:e2:cd:79:24:f8:97:a3:8d:4d:27
    8c:2b:dd:f1:57:72:4d:2b:6d:d5:95:0d:9c:1b:5c:e2:a4:b0:84:2e
    31:82:3c:91:08:a2:58:b9:30:4c:5f:a3:6b:e6:2b:9c:b1:42:dd:1c
    cd:a2:7a:84:ea:7b:a6:b7:9a:13:33:c6:27:2b:79:2a:b1:0c:fe:08
    4c:a7:35:fc:da:4f:df:1f:cf:f4:ba:bc:5a:05:06:63:92:41:b4:f2
    54:00:3f:ef:ff:41:e6:ca:74:10:56:f7:2b:5f:d3:1a:33:7e:49:74
    1c:42:cf:c2:23:ea:4b:8f:50:2c:eb:1c:a6:37:89:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    d6:7f:52:a3:b6:f8:ae:cb:70:3f:a9:79:ea:8a:da:9e:ba:83:e4:5f (sha1)
    76:0b:72:73:cf:51:ee:58:81:2d:f7:b4:e2:5c:f4:5c (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT for Local1 and SBU for Local2. The IKE configurations on the hub include OU=SLT and OU=SBU to identify the spoke.

Configuring the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/2 unit 0 family inet address 1.1.2.1/30
set interfaces ge-0/0/3 unit 0 family inet address 50.50.50.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.1/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 20.20.20.1/24
set policy-options policy-statement lan_nw from interface ge-0/0/3.0
set policy-options policy-statement lan_nw then accept
set policy-options policy-statement load_balance then load-balance per-packet
set protocols bgp group ibgp-1 type internal
set protocols bgp group ibgp-1 local-address 10.10.10.1
set protocols bgp group ibgp-1 export lan_nw
set protocols bgp group ibgp-1 cluster 1.2.3.4
set protocols bgp group ibgp-1 multipath
set protocols bgp group ibgp-1 allow 10.10.10.0/24
set protocols bgp group ibgp-2 type internal
set protocols bgp group ibgp-2 local-address 20.20.20.1
set protocols bgp group ibgp-2 export lan_nw
set protocols bgp group ibgp-2 cluster 1.2.3.5
set protocols bgp group ibgp-2 multipath
set protocols bgp group ibgp-2 allow 20.20.20.0/24
set routing-options static route 2.2.2.0/30 next-hop 1.1.1.2
set routing-options static route 3.3.3.0/30 next-hop 1.1.2.2
set routing-options autonomous-system 10
set routing-options forwarding-table export load_balance
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy-1 mode main
set security ike policy ike-policy-1 proposals ike-proposal
set security ike policy ike-policy-1 certificate local-certificate Local1
set security ike policy ike-policy-2 mode main
set security ike policy ike-policy-2 proposals ike-proposal
set security ike policy ike-policy-2 certificate local-certificate Local2
set security ike gateway hub-to-spoke-gw-1 ike-policy ike-policy-1
set security ike gateway hub-to-spoke-gw-1 dynamic distinguished-name wildcard OU=SLT
set security ike gateway hub-to-spoke-gw-1 dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw-1 local-identity distinguished-name
set security ike gateway hub-to-spoke-gw-1 external-interface ge-0/0/1.0
set security ike gateway hub-to-spoke-gw-2 ike-policy ike-policy-2
set security ike gateway hub-to-spoke-gw-2 dynamic distinguished-name wildcard
  OU=SBU
set security ike gateway hub-to-spoke-gw-2 dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw-2 local-identity distinguished-name
set security ike gateway hub-to-spoke-gw-2 external-interface ge-0/0/2.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy proposals ipsec-proposal
set security ipsec vpn hub-to-spoke-vpn-1 bind-interface st0.0
set security ipsec vpn hub-to-spoke-vpn-1 ike gateway hub-to-spoke-gw-1
set security ipsec vpn hub-to-spoke-vpn-1 ike ipsec-policy vpn-policy
set security ipsec vpn hub-to-spoke-vpn-2 bind-interface st0.1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec vpn hub-to-spoke-vpn-2 ike gateway hub-to-spoke-gw-2
set security ipsec vpn hub-to-spoke-vpn-2 ike ipsec-policy vpn-policy
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces ge-0/0/2.0
set security zones security-zone untrust interfaces st0.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the hub:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 1.1.1.1/30
user@host# set ge-0/0/2 unit 0 family inet address 1.1.2.1/30
user@host# set ge-0/0/3 unit 0 family inet address 50.50.50.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.1/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 20.20.20.1/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface ge-0/0/3.0
user@host# set policy-statement lan_nw then accept
user@host# set policy-statement load_balance then load-balance per-packet

[edit protocols bgp]
user@host# set group ibgp-1 type internal
user@host# set group ibgp-1 local-address 10.10.10.1
user@host# set group ibgp-1 export lan_nw
user@host# set group ibgp-1 cluster 1.2.3.4
user@host# set group ibgp-1 multipath
user@host# set group ibgp-1 allow 10.10.10.0/24

user@host# set group ibgp-2 type internal
user@host# set group ibgp-2 local-address 20.20.20.1
user@host# set group ibgp-2 export lan_nw
user@host# set group ibgp-2 cluster 1.2.3.5
user@host# set group ibgp-2 multipath
user@host# set group ibgp-2 allow 20.20.20.0/24
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit routing-options]
user@host# set static route 2.2.2.0/30 next-hop 1.1.1.2
user@host# set static route 3.3.3.0/30 next-hop 1.1.2.2
user@host# set autonomous-system 10
user@host# set forwarding-table export load_balance
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy-1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike policy ike-policy-2]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local2

[edit security ike gateway hub-to-spoke-gw-1]
user@host# set ike-policy ike-policy-1
user@host# set dynamic distinguished-name wildcard OU=SLT
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/1.0

[edit security ike gateway hub-to-spoke-gw-2]
user@host# set ike-policy ike-policy-2
user@host# set dynamic distinguished-name wildcard OU=SBU
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/2.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn hub-to-spoke-vpn-1]
user@host# set bind-interface st0.0
user@host# set ike gateway hub-to-spoke-gw-1
user@host# set ike ipsec-policy vpn-policy

[edit security ipsec vpn hub-to-spoke-vpn-2]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set bind-interface st0.1
user@host# set ike gateway hub-to-spoke-gw-2
user@host# set ike ipsec-policy vpn-policy
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.0
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces ge-0/0/2.0
user@host# set interfaces st0.1

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/3.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 1.1.1.1/30;
    }
  }
}
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 1.1.2.1/30;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
```

Copyright © 2016, Juniper Networks, Inc.

```
            address 50.50.50.1/24;
          }
        }
      }
      st0 {
        unit 0 {
          multipoint;
          family inet {
            address 10.10.10.1/24;
          }
        }
        unit 1 {
          multipoint;
          family inet {
            address 20.20.20.1/24;
          }
        }
      }
[edit]
user@host# show policy-options
policy-statement lan_nw {
  from interface ge-0/0/3.0;
  then accept;
}
  policy-statement load_balance {
    then {
      load-balance per-packet;
    }
  }
[edit]
user@host# show protocols
bgp {
  group ibgp-1 {
    type internal;
    local-address 10.10.10.1;
    export lan_nw;
    cluster 1.2.3.4;
    multipath;
    allow 10.10.10.0/24;
  }
  group ibgp-2 {
    type internal;
    local-address 20.20.20.1;
    export lan_nw;
    cluster 1.2.3.5;
    multipath;
    allow 20.20.20.0/24;
  }
}
[edit]
user@host# show routing-options
static {
  route 2.2.2.0/30 next-hop 1.1.1.2;
  route 3.3.3.0/30 next-hop 1.1.2.2;
  }
autonomous-system 10;
```


Copyright © 2016, Juniper Networks, Inc.

```
    forwarding-table {
      export load_balance;
    }
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy-1 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local1;
      }
    }
    policy ike-policy-2 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local2;
      }
    }
    gateway hub-to-spoke-gw-1 {
      ike-policy ike-policy-1;
      dynamic {
        distinguished-name {
          wildcard OU=SLT;
        }
        ike-user-type group-ike-id;
      }
      local-identity distinguished-name;
      external-interface ge-0/0/1.0;
    }
    gateway hub-to-spoke-gw-2 {
      ike-policy ike-policy-2;
      dynamic {
        distinguished-name {
          wildcard OU=SBU;
        }
        ike-user-type group-ike-id;
      }
      local-identity distinguished-name;
      external-interface ge-0/0/2.0;
    }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
    protocol esp;
    authentication-algorithm hmac-md5-96;
    encryption-algorithm des-cbc;
}
    policy vpn-policy {
      perfect-forward-secrecy {
```

Copyright © 2016, Juniper Networks, Inc.

```
        keys group14;
      }
      proposals ipsec-proposal;
    }
    vpn hub-to-spoke-vpn-1 {
      bind-interface st0.0;
      ike {
        gateway hub-to-spoke-gw-1;
        ipsec-policy vpn-policy;
      }
    }
    vpn hub-to-spoke-vpn-2 {
      bind-interface st0.1;
      ike {
        gateway hub-to-spoke-gw-2;
        ipsec-policy vpn-policy;
      }
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      st0.0;
      ge-0/0/1.0;
      ge-0/0/2.0;
      st0.1;
    }
}
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        ge-0/0/3.0;
      }
    }
[edit]
user@host# show security policies
default-policy {
    permit-all;
}
[edit]
user@host# show security pki
```


Copyright © 2016, Juniper Networks, Inc.

```
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/1 unit 0 family inet address 2.2.2.1/30
set interfaces fe-0/0/2 unit 0 family inet address 3.3.3.1/30
set interfaces fe-0/0/4 unit 0 family inet address 60.60.60.1/24
set interfaces st0 unit 0 family inet address 10.10.10.2/24
set interfaces st0 unit 1 family inet address 20.20.20.2/24
set policy-options policy-statement lan_nw from interface fe-0/0/4.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp-1 type internal
set protocols bgp group ibgp-1 local-address 10.10.10.2
set protocols bgp group ibgp-1 export lan_nw
set protocols bgp group ibgp-1 neighbor 10.10.10.1
set protocols bgp group ibgp-2 type internal
set protocols bgp group ibgp-2 local-address 20.20.20.2
set protocols bgp group ibgp-2 export lan_nw
set protocols bgp group ibgp-2 neighbor 20.20.20.1
set routing-options static route 1.1.1.0/30 next-hop 2.2.2.2
set routing-options static route 1.1.2.0/30 next-hop 3.3.3.2
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy-1 mode main
set security ike policy ike-policy-1 proposals ike-proposal
set security ike policy ike-policy-1 certificate local-certificate Local1
set security ike policy ike-policy-2 mode main
set security ike policy ike-policy-2 proposals ike-proposal
set security ike policy ike-policy-2 certificate local-certificate Local2
set security ike gateway spoke-to-hub-gw-1 ike-policy ike-policy-1
set security ike gateway spoke-to-hub-gw-1 address 1.1.1.1
set security ike gateway spoke-to-hub-gw-1 local-identity distinguished-name
set security ike gateway spoke-to-hub-gw-1 remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw-1 external-interface fe-0/0/1.0
set security ike gateway spoke-to-hub-gw-2 ike-policy ike-policy-2
set security ike gateway spoke-to-hub-gw-2 address 1.1.2.1
set security ike gateway spoke-to-hub-gw-2 local-identity distinguished-name
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ike gateway spoke-to-hub-gw-2 remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw-2 external-interface fe-0/0/2.0
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy proposals ipsec-proposal
set security ipsec vpn spoke-to-hub-1 bind-interface st0.0
set security ipsec vpn spoke-to-hub-1 ike gateway spoke-to-hub-gw-1
set security ipsec vpn spoke-to-hub-1 ike ipsec-policy vpn-policy
set security ipsec vpn spoke-to-hub-1 establish-tunnels immediately
set security ipsec vpn spoke-to-hub-2 bind-interface st0.1
set security ipsec vpn spoke-to-hub-2 ike gateway spoke-to-hub-gw-2
set security ipsec vpn spoke-to-hub-2 ike ipsec-policy vpn-policy
set security ipsec vpn spoke-to-hub-2 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces fe-0/0/2.0
set security zones security-zone untrust interfaces st0.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 1:

1. Configure interfaces.

```
[edit interfaces]
user@host# set fe-0/0/1 unit 0 family inet address 2.2.2.1/30
user@host# set fe-0/0/2 unit 0 family inet address 3.3.3.1/30
user@host# set fe-0/0/4 unit 0 family inet address 60.60.60.1/24
user@host# set st0 unit 0 family inet address 10.10.10.2/24
user@host# set st0 unit 1 family inet address 20.20.20.2/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface fe-0/0/4.0
user@host# set policy-statement lan_nw then accept

[edit protocols bgp]
user@host# set group ibgp-1 type internal
user@host# set group ibgp-1 local-address 10.10.10.2
user@host# set group ibgp-1 export lan_nw
user@host# set group ibgp-1 neighbor 10.10.10.1
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set group ibgp-2 type internal
user@host# set group ibgp-2 local-address 20.20.20.2
user@host# set group ibgp-2 export lan_nw
user@host# set group ibgp-2 neighbor 20.20.20.1

[edit routing-options]
user@host# set static route 1.1.1.0/30 next-hop 2.2.2.2
user@host# set static route 1.1.2.0/30 next-hop 3.3.3.2
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy-1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike policy ike-policy-2]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local2

[edit security ike gateway spoke-to-hub-gw-1]
user@host# set ike-policy ike-policy-1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/1.0

[edit security ike gateway spoke-to-hub-gw-2]
user@host# set ike-policy ike-policy-2
user@host# set address 1.1.2.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/2.0
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub-1]
user@host# set bind-interface st0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ike gateway spoke-to-hub-gw-1
user@host# set ike ipsec-policy vpn-policy
user@host# set establish-tunnels immediately

[edit security ipsec vpn spoke-to-hub-2]
user@host# set bind-interface st0.1
user@host# set ike gateway spoke-to-hub-gw-2
user@host# set ike ipsec-policy vpn-policy
user@host# set establish-tunnels immediately
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/1.0
user@host# set interfaces st0.0
user@host# set interfaces fe-0/0/2.0
user@host# set interfaces st0.1

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
fe-0/0/1 {
   unit 0 {
      family inet {
         address 2.2.2.1/30;
      }
   }
}
   fe-0/0/2 {
      unit 0 {
         family inet {
```


Copyright © 2016, Juniper Networks, Inc.

```
            address 3.3.3.1/30;
          }
        }
      }
      fe-0/0/4 {
        unit 0 {
          family inet {
            address 60.60.60.1/24;
          }
        }
      }
      st0 {
        unit 0 {
          family inet {
            address 10.10.10.2/24;
          }
        }
        unit 1 {
          family inet {
            address 20.20.20.2/24;
          }
        }
      }
[edit]
user@host# show policy-options
policy-statement lan_nw {
  from interface fe-0/0/4.0;
  then accept;
}
[edit]
user@host# show protocols
bgp {
  group ibgp-1 {
    type internal;
    local-address 10.10.10.2;
    export lan_nw;
    neighbor 10.10.10.1;
  }
  group ibgp-2 {
    type internal;
    local-address 20.20.20.2;
    export lan_nw;
    neighbor 20.20.20.1;
  }
}
[edit]
user@host# show routing-options
static {
  route 1.1.1.0/30 next-hop 2.2.2.2;
  route 1.1.2.0/30 next-hop 3.3.3.2;
  }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
  authentication-method rsa-signatures;
```

Copyright © 2016, Juniper Networks, Inc.

```
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy-1 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local1;
        }
    }
    policy ike-policy-2 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local2;
        }
    }
    gateway spoke-to-hub-gw-1 {
        ike-policy ike-policy-1;
        address 1.1.1.1;
        local-identity distinguished-name;
        remote-identity distinguished-name;
        external-interface fe-0/0/1.0;
    }
    gateway spoke-to-hub-gw-2 {
        ike-policy ike-policy-2;
        address 1.1.2.1;
        local-identity distinguished-name;
        remote-identity distinguished-name;
        external-interface fe-0/0/2.0;
    }
[edit]
user@host# show security ipsec
proposal ipsec-proposal {
    protocol esp;
    authentication-algorithm hmac-md5-96;
    encryption-algorithm des-cbc;
}
    policy vpn-policy {
        perfect-forward-secrecy {
            keys group14;
        }
        proposals ipsec-proposal;
    }
    vpn spoke-to-hub-1 {
        bind-interface st0.0;
        ike {
            gateway spoke-to-hub-gw-1;
            ipsec-policy vpn-policy;
        }
        establish-tunnels immediately;
    }
    vpn spoke-to-hub-2 {
        bind-interface st0.1;
        ike {
```


Copyright © 2016, Juniper Networks, Inc.

```
        gateway spoke-to-hub-gw-2;
        ipsec-policy vpn-policy;
      }
      establish-tunnels immediately;
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      fe-0/0/1.0;
      st0.0;
      fe-0/0/2.0;
      st0.1;
    }
}
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        fe-0/0/4.0;
      }
    }
[edit]
user@host# show security policies
default-policy {
    permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
    ca-identity ca-profile1;
    enrollment {
      url http://pc4/certsrv/mscep/mscep.dll;
    }
    revocation-check {
      disable;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Verifying IKE Phase 1 Status on page 6528
- Verifying IPsec Phase 2 Status on page 6528
- Verifying IPsec Next-Hop Tunnels on page 6529
- Verifying BGP on page 6529
- Verifying Learned Routes on page 6529
- Verifying Route Installation in Forwarding Table on page 6530

Verifying IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

3733049 UP     bc9686796c2e52e9  1fbe46eee168f24e  Main           2.2.2.1

3733048 UP     a88db7ed23ec5f6b  c88b81dff52617a5  Main           3.3.3.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spoke.

Verifying IPsec Phase 2 Status

Purpose Verify the IPsec Phase 2 status.

Action From operational mode, enter the **security ipsec security-associations** command.

```
user@host> security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <268173315 ESP:des/ md5 93cfb417 1152/ unlim -   root 500   2.2.2.1
  >268173315 ESP:des/ md5 101de6f7 1152/ unlim -   root 500   2.2.2.1
  <268173313 ESP:des/ md5 272e29c0 1320/ unlim -   root 500   3.3.3.1
  >268173313 ESP:des/ md5 a3bf8fad 1320/ unlim -   root 500   3.3.3.1
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spoke.

Copyright © 2016, Juniper Networks, Inc.

Verifying IPsec Next-Hop Tunnels

Purpose Verify the IPsec next-hop tunnels.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPSec VPN name                    Flag     IKE-ID
                              XAUTH username
10.10.10.2        st0.0       hub-to-spoke-vpn-1               Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
20.20.20.2        st0.1       hub-to-spoke-vpn-2               Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SBU, CN=spoke1_backup
```

Meaning The next-hop gateways are the IP addresses for the **st0** interfaces of the spokes. The next hop should be associated with the correct IPsec VPN name.

Verifying BGP

Purpose Verify that BGP references the IP addresses for the **st0** interfaces of the spoke.

Action From operational mode, enter the **show bgp summary** command.

```
user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Unconfigured peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2               10       4819       4820       0       2 1d 12:15:14
1/1/1/0              0/0/0/0
20.20.20.2               10       4926       4928       0       0 1d 13:03:03
1/1/1/0              0/0/0/0
```

Verifying Learned Routes

Purpose Verify that routes to the spoke have been learned.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route 60.60.60.0 detail** command.

```
user@host> show route 60.60.60.0 detail
inet.0: 47 destinations, 48 routes (46 active, 0 holddown, 1 hidden)
60.60.60.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Address: 0x167407c
                Next-hop reference count: 3
                Source: 10.10.10.2
                Next hop type: Router
                Next hop: 10.10.10.2 via st0.0
                Next hop type: Router
                Next hop: 20.20.20.2 via st0.1, selected
                Protocol next hop: 10.10.10.2
                Indirect next hop: 15c8000 262142
                Protocol next hop: 20.20.20.2
                Indirect next hop: 15c80e8 262143
                State: <Act Int Ext>
                Local AS:    10 Peer AS:    10
                Age: 1d 12:16:25  Metric2: 0
                Task: BGP_10.10.10.10.2+53120
                Announcement bits (2): 0-KRT 3-Resolve tree 1
                AS path: I
                Accepted Multipath
                Localpref: 100
                Router ID: 10.207.36.182
         BGP    Preference: 170/-101
                Next hop type: Indirect
                Address: 0x15b8ac0
                Next-hop reference count: 1
                Source: 20.20.20.2
                Next hop type: Router
                Next hop: 20.20.20.2 via st0.1, selected
                Protocol next hop: 20.20.20.2
                Indirect next hop: 15c80e8 262143
                State: <NotBest Int Ext>
                Inactive reason: Not Best in its group - Update source
                Local AS:    10 Peer AS:    10
                Age: 1d 13:04:14  Metric2: 0
                Task: BGP_10.20.20.20.2+50733
                AS path: I
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 10.207.36.182
```

Verifying Route Installation in Forwarding Table

Purpose Verify that routes to the spoke have been installed in the forwarding table.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route forwarding-table matching 60.60.60.0** command.

```
user@host> show route forwarding-table matching 60.60.60.0
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
60.60.60.0/24      user     0                    ulst 262144     1
                                                 indr 262142     2
                              10.10.10.2         ucst   572     3 st0.0
                                                 indr 262143     2
                              20.20.20.2         ucst   573     3 st0.1
```

Related Documentation

- Example: Configuring a Route-Based VPN on page 6074
- *Routing Protocols Overview for Security Devices*

Example: Configuring AutoVPN with iBGP and Active-Backup Tunnels

This example shows how to configure active and backup IPsec VPN tunnels between an AutoVPN hub and spoke. This example configures iBGP to forward traffic through the VPN tunnels.

- Requirements on page 6531
- Overview on page 6531
- Configuration on page 6534
- Verification on page 6553

Requirements

This example uses the following hardware and software components:

- Two supported SRX Series devices as AutoVPN hub and spoke
- Junos OS Release 12.1X44-D10 and later that support AutoVPN

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates.

NOTE: You should be familiar with the dynamic routing protocol that is used to forward packets through the VPN tunnels. For more information about specific requirements for a dynamic routing protocol, see the *Routing Protocols Overview for Security Devices*.

Overview

This example shows the configuration of an AutoVPN hub and a spoke with two IPsec VPN tunnels.

Copyright © 2016, Juniper Networks, Inc.

In this example, the first step is to enroll digital certificates in each device using the Simple Certificate Enrollment Protocol (SCEP). Certificates are enrolled in the hub and in the spoke for each IPsec VPN tunnel. One of the certificates for the spoke contains the organizational unit (OU) value "SLT" in the distinguished name (DN); the hub is configured with a group IKE ID to match the value "SLT" in the OU field. The other certificate for the spoke contains the OU value "SBU" in the DN; the hub is configured with a group IKE ID to match the value "SBU" in the OU field.

The spoke establishes IPsec VPN connections to the hub, which allows it to access resources on the hub. Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hub and the spoke must have the same values. Table 562 shows the options used in this example.

Table 562: Phase 1 and Phase 2 Options for AutoVPN Hub and Spoke iBGP Active-Backup Tunnel Configurations

Option	Value
IKE proposal:	
Authentication method	RSA digital certificates
Diffie-Hellman (DH) group	2
Authentication algorithm	SHA-1
Encryption algorithm	AES 128 CBC
IKE policy:	
Mode	Main
IPsec proposal:	
Protocol	ESP
Authentication algorithm	HMAC MD5 96
Encryption algorithm	DES CBC
IPsec policy:	
Perfect Forward Secrecy (PFS) group	14

The same certificate authority (CA) is configured on all devices.

NOTE: **Junos OS only supports a single level of certificate hierarchy.**

Table 563 shows the options configured on the hub and on the spoke.

Copyright © 2016, Juniper Networks, Inc.

Table 563: AutoVPN IBGP Active-Backup Tunnel Configuration for Hub and Spoke 1

Option	Hub	Spoke 1
IKE gateway:		
Remote IP address	hub-to-spoke-gw-1: Dynamic hub-to-spoke-gw-2: Dynamic	spoke-to-hub-gw-1: 1.1.1.1 spoke-to-hub-gw-2: 1.1.2.1
Remote IKE ID	hub-to-spoke-gw-1: DN on the spoke's certificate with the string **SLT** in the OU field hub-to-spoke-gw-2: DN on the spoke's certificate with the string **SBU** in the OU field	spoke-to-hub-gw-1: DN on the hub's certificate spoke-to-hub-gw-2: DN on the hub's certificate
Local IKE ID	DN on the hub's certificate	DN on the spoke's certificate
External interface	hub-to-spoke-gw-1: ge-0/0/1.0 hub-to-spoke-gw-2: ge-0/0/2.0	spoke-to-hub-gw-1: fe-0/0/1.0 spoke-to-hub-gw-2: fe-0/0/2.0
VPN:		
Bind interface	hub-to-spoke-vpn-1: st0.0 hub-to-spoke-vpn-2: st0.1	spoke-to-hub-1: st0.0 spoke-to-hub-2: st0.1
VPN monitor	hub-to-spoke-vpn-1: ge-0/0/1.0 (source interface) hub-to-spoke-vpn-2: ge-0/0/2.0 (source interface)	spoke-to-hub-1: 1.1.1.1 (destination IP) spoke-to-hub-2: 1.1.2.1 (destination IP)
Establish tunnels	(not configured)	Immediately on configuration commit

Routing information for all devices is exchanged through the VPN tunnels.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Topology

Figure 286 shows the SRX Series devices to be configured for AutoVPN in this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 286: AutoVPN Deployment with iBGP and Active-Backup Tunnels

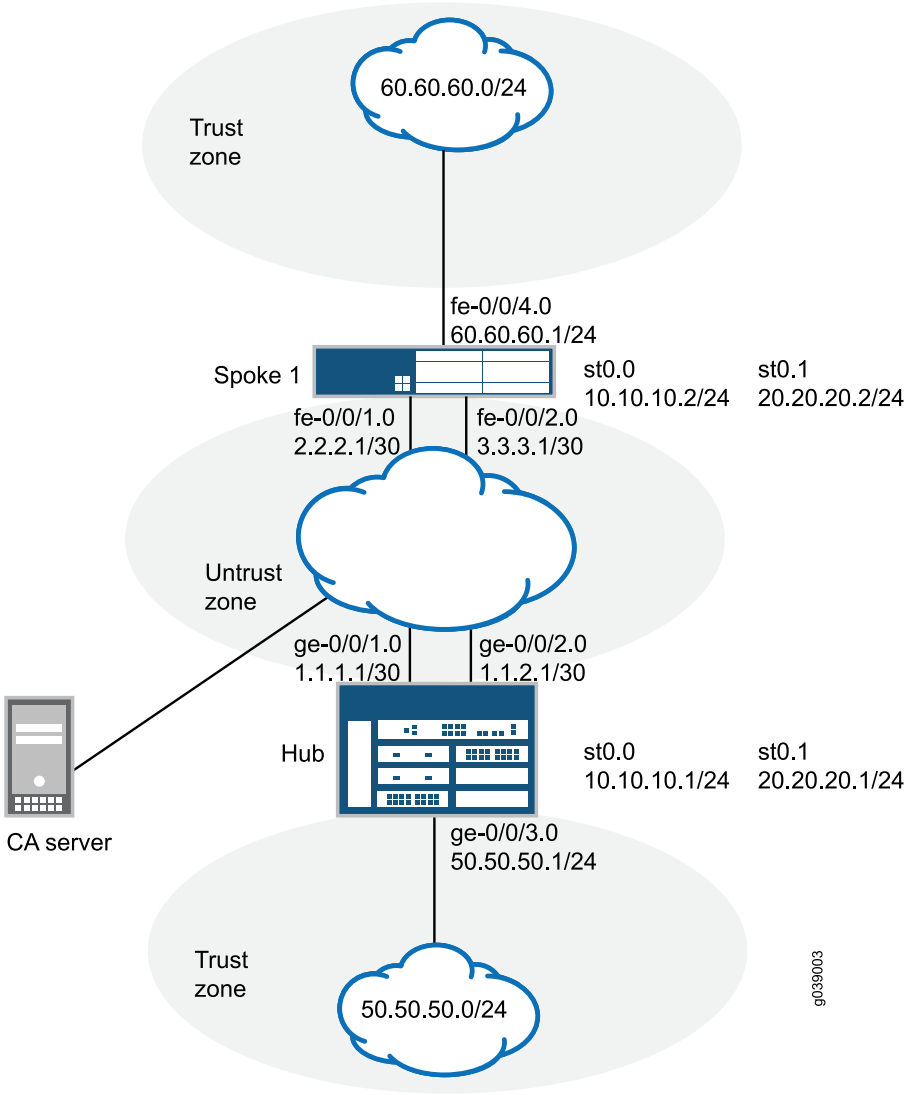

In this example, two IPsec VPN tunnels are established between the hub and spoke 1. Routing information is exchanged through iBGP sessions in each tunnel. The longest prefix match for the route to 60.60.60.0/24 is through the st0.0 interface on the hub. Thus, the primary tunnel for the route is through the st0.0 interfaces on the hub and spoke 1. The default route is through the backup tunnel on the st0.1 interfaces on the hub and spoke 1.

VPN monitoring checks the status of the tunnels. If there is a problem with the primary tunnel (for example, the remote tunnel gateway is not reachable), the tunnel status changes to down and data destined for 60.60.60.0/24 is rerouted through the backup tunnel.

Configuration

To configure AutoVPN, perform these tasks:

Copyright © 2016, Juniper Networks, Inc.

NOTE: The first section describes how to obtain CA and local certificates online using the Simple Certificate Enrollment Protocol (SCEP) on the hub and spoke devices.

- Enroll Device Certificates with SCEP on page 6535
- Configuring the Hub on page 6539
- Configuring Spoke 1 on page 6546

Enroll Device Certificates with SCEP

Step-by-Step Procedure

To enroll digital certificates with SCEP on the hub:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

   Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair for each certificate.

```
user@host> request security pki generate-key-pair certificate-id Local1
user@host> request security pki generate-key-pair certificate-id Local2
```

4. Enroll the local certificates.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email hub@example.net ip-address
  1.1.1.1 subject
  DC=example.net,CN=hub,OU=SLT,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local2 domain-name example.net email hub_backup@example.net
  ip-address 1.1.2.1 subject
  DC=example.net,CN=hub_backup,OU=SBU,O=example,L=Bangalore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificates.

```
user@host> show security pki local-certificate certificate-id Local1 detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a6d5f300000000258d
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,
```

Copyright © 2016, Juniper Networks, Inc.

```
    Locality: Bangalore, Common name: hub, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SLT, CN=hub
  Alternate subject: "hub@example.net", example.net, 1.1.1.1
  Validity:
    Not before: 11- 6-2012 09:39
    Not after: 11- 6-2013 09:49
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:c9:c9:cc:30:b6:7a:86:12:89:b5:18:b3:76
    01:2d:cc:65:a8:a8:42:78:cd:d0:9a:a2:c0:aa:c4:bd:da:af:88:f3
    2a:78:1f:0a:58:e6:11:2c:81:8f:0e:7c:de:86:fc:48:4c:28:5b:8b
    34:91:ff:2e:91:e7:b5:bd:79:12:de:39:46:d9:fb:5c:91:41:d1:da
    90:f5:09:00:9b:90:07:9d:50:92:7d:ff:fb:3f:3c:bc:34:e7:e3:c8
    ea:cb:99:18:b4:b6:1d:a8:99:d3:36:b9:1b:36:ef:3e:a1:fd:48:82
    6a:da:22:07:da:e0:d2:55:ef:57:be:09:7a:0e:17:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    e1:f7:a1:a6:1e:c3:97:69:a5:07:9b:09:14:1a:c7:ae:09:f1:f6:35 (sha1)
    a0:02:fa:8d:5c:63:e5:6d:f7:f4:78:56:ac:4e:b2:c4 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started

user@host> show security pki local-certificate certificate-id Local2 detail

Certificate identifier: Local2
  Certificate version: 3
  Serial number: 505efdf900000000259a
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SBU, Country: IN, State: KA,

    Locality: Bangalore, Common name: hub_backup, Domain component:
example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Bangalore, O=example, OU=SBU, CN=hub_backup

  Alternate subject: "hub_backup@example.net", example.net, 1.1.2.1
  Validity:
    Not before: 11- 9-2012 10:55
    Not after: 11- 9-2013 11:05
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:d5:44:08:96:f6:77:05:e6:91:50:8a:8a:2a
    4e:95:43:1e:88:ea:43:7c:c5:ac:88:d7:a0:8d:b5:d9:3f:41:db:db
    44:34:1f:56:a5:38:4b:b2:c5:85:f9:f1:bf:b2:7b:d4:b2:af:98:a0
    95:50:02:ad:f5:dd:4d:dc:67:85:dd:84:09:df:9c:68:a5:58:65:e7
    2c:72:cc:47:4b:d0:cc:4a:28:ca:09:db:ad:6e:5a:13:6c:e6:cc:f0
    29:ed:2b:2d:d1:38:38:bc:68:84:de:ae:86:39:c9:dd:06:d5:36:f0
    e6:2a:7b:46:4c:cd:a5:24:1c:e0:92:8d:ad:35:29:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    98:96:2f:ff:ca:af:33:ee:d7:4c:c8:4f:f7:71:53:c0:5d:5f:c5:59 (sha1)
    c9:87:e3:a4:5c:47:b5:aa:90:22:e3:06:b2:0b:e1:ea (md5)
```


Copyright © 2016, Juniper Networks, Inc.

```
Auto-re-enrollment:
  Status: Disabled
  Next trigger time: Timer not started
```

Step-by-Step Procedure

To enroll digital certificates with SCEP on spoke 1:

1. Configure the CA.

```
[edit]
user@host# set security pki ca-profile ca-profile1 ca-identity ca-profile1
user@host# set security pki ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set security pki ca-profile ca-profile1 revocation-check disable
user@host# commit
```

2. Enroll the CA certificate.

```
user@host> request security pki ca-certificate enroll ca-profile ca-profile1
```

Type **yes** at the prompt to load the CA certificate.

3. Generate a key pair for each certificate.

```
user@host> rrequest security pki generate-key-pair certificate-id Local1
user@host> request security pki generate-key-pair certificate-id Local2
```

4. Enroll the local certificates.

```
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local1 domain-name example.net email spoke1@example.net
  ip-address 2.2.2.1 subject
  DC=example.net,CN=spoke1,OU=SLT,O=example,L=Mysore,ST=KA,C=IN
  challenge-password <password>
user@host> request security pki local-certificate enroll ca-profile ca-profile1
  certificate-id Local2 domain-name example.net email
  spoke1_backup@example.net ip-address 3.3.3.1 subject
  DC=example.net,CN=spoke1_backup,OU=SBU,O=example,L=Mysore,ST=KA,C=IN
  challenge-password <password>
```

5. Verify the local certificates.

```
user@host> show security pki local-certificate certificate-id Local1 detail

Certificate identifier: Local1
  Certificate version: 3
  Serial number: 40a7975f00000000258e
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SLT, Country: IN, State: KA,

    Locality: Mysore, Common name: spoke1, Domain component: example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
  Alternate subject: "spoke1@example.net", example.net, 2.2.2.1
  Validity:
    Not before: 11- 6-2012 09:40
    Not after: 11- 6-2013 09:50
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:d8:45:09:77:cd:36:9a:6f:58:44:18:91:db
```

Copyright © 2016, Juniper Networks, Inc.

```
    b0:c7:8a:ee:c8:d7:a6:d2:e2:e7:20:46:2b:26:1a:92:e2:4e:8a:ce
    c9:25:d9:74:a2:81:ad:ea:e0:38:a0:2f:2d:ab:a6:58:ac:88:35:f4
    90:01:08:33:33:75:2c:44:26:f8:25:18:97:96:e4:28:de:3b:35:f2
    4a:f5:92:b7:57:ae:73:4f:8e:56:71:ab:81:54:1d:75:88:77:13:64
    1b:6b:01:96:15:0a:1c:54:e3:db:f8:ec:ec:27:5b:86:39:c1:09:a1
    e4:24:1a:19:0d:14:2c:4b:94:a4:04:91:3f:cb:ef:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    b6:24:2a:0e:96:5d:8c:4a:11:f3:5a:24:89:7c:df:ea:d5:c0:80:56 (sha1)
    31:58:7f:15:bb:d4:66:b8:76:1a:42:4a:8a:16:b3:a9 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

```
user@host> show security pki local-certificate certificate-id Local2 detail

Certificate identifier: Local2
  Certificate version: 3
  Serial number: 506c3d0600000000259b
  Issuer:
    Common name: CASERVER1, Domain component: net, Domain component: internal

  Subject:
    Organization: example, Organizational unit: SBU, Country: IN, State: KA,

    Locality: Mysore, Common name: spoke1_backup, Domain component:
example.net
  Subject string:
    C=IN, DC=example.net, ST=KA, L=Mysore, O=example, OU=SBU, CN=spoke1_backup

  Alternate subject: "spoke1_backup@example.net", example.net, 3.3.3.1
  Validity:
    Not before: 11- 9-2012 11:09
    Not after: 11- 9-2013 11:19
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:a7:02:b5:e2:cd:79:24:f8:97:a3:8d:4d:27
    8c:2b:dd:f1:57:72:4d:2b:6d:d5:95:0d:9c:1b:5c:e2:a4:b0:84:2e
    31:82:3c:91:08:a2:58:b9:30:4c:5f:a3:6b:e6:2b:9c:b1:42:dd:1c
    cd:a2:7a:84:ea:7b:a6:b7:9a:13:33:c6:27:2b:79:2a:b1:0c:fe:08
    4c:a7:35:fc:da:4f:df:1f:cf:f4:ba:bc:5a:05:06:63:92:41:b4:f2
    54:00:3f:ef:ff:41:e6:ca:74:10:56:f7:2b:5f:d3:1a:33:7e:49:74
    1c:42:cf:c2:23:ea:4b:8f:50:2c:eb:1c:a6:37:89:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    http://ca-server1/CertEnroll/CASERVER1.crl
    file://\\ca-server1\CertEnroll\CASERVER1.crl
  Fingerprint:
    d6:7f:52:a3:b6:f8:ae:cb:70:3f:a9:79:ea:8a:da:9e:ba:83:e4:5f (sha1)
    76:0b:72:73:cf:51:ee:58:81:2d:f7:b4:e2:5c:f4:5c (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

NOTE: The organizational unit (OU) shown in the subject field is SLT for Local1 and SBU for Local2. The IKE configurations on the hub include OU=SLT and OU=SBU to identify the spoke.

Copyright © 2016, Juniper Networks, Inc.

Configuring the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/2 unit 0 family inet address 1.1.2.1/30
set interfaces ge-0/0/3 unit 0 family inet address 50.50.50.1/24
set interfaces st0 unit 0 multipoint
set interfaces st0 unit 0 family inet address 10.10.10.1/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 20.20.20.1/24
set policy-options policy-statement lan_nw from interface ge-0/0/3.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp-1 type internal
set protocols bgp group ibgp-1 local-address 10.10.10.1
set protocols bgp group ibgp-1 export lan_nw
set protocols bgp group ibgp-1 cluster 1.2.3.4
set protocols bgp group ibgp-1 allow 10.10.10.0/24
set protocols bgp group ibgp-2 type internal
set protocols bgp group ibgp-2 local-address 20.20.20.1
set protocols bgp group ibgp-2 export lan_nw
set protocols bgp group ibgp-2 cluster 1.2.3.5
set protocols bgp group ibgp-2 allow 20.20.20.0/24
set routing-options static route 2.2.2.0/30 next-hop 1.1.1.2
set routing-options static route 3.3.3.0/30 next-hop 1.1.2.2
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy-1 mode main
set security ike policy ike-policy-1 proposals ike-proposal
set security ike policy ike-policy-1 certificate local-certificate Local1
set security ike policy ike-policy-2 mode main
set security ike policy ike-policy-2 proposals ike-proposal
set security ike policy ike-policy-2 certificate local-certificate Local2
set security ike gateway hub-to-spoke-gw-1 ike-policy ike-policy-1
set security ike gateway hub-to-spoke-gw-1 dynamic distinguished-name wildcard OU=SLT
set security ike gateway hub-to-spoke-gw-1 dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw-1 local-identity distinguished-name
set security ike gateway hub-to-spoke-gw-1 external-interface ge-0/0/1.0
set security ike gateway hub-to-spoke-gw-2 ike-policy ike-policy-2
set security ike gateway hub-to-spoke-gw-2 dynamic distinguished-name wildcard
  OU=SBU
set security ike gateway hub-to-spoke-gw-2 dynamic ike-user-type group-ike-id
set security ike gateway hub-to-spoke-gw-2 local-identity distinguished-name
set security ike gateway hub-to-spoke-gw-2 external-interface ge-0/0/2.0
set security ipsec vpn-monitor-options interval 5
set security ipsec vpn-monitor-options threshold 2
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ipsec policy vpn-policy perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy proposals ipsec-proposal
set security ipsec vpn hub-to-spoke-vpn-1 bind-interface st0.0
set security ipsec vpn hub-to-spoke-vpn-1 vpn-monitor source-interface ge-0/0/1.0
set security ipsec vpn hub-to-spoke-vpn-1 ike gateway hub-to-spoke-gw-1
set security ipsec vpn hub-to-spoke-vpn-1 ike ipsec-policy vpn-policy
set security ipsec vpn hub-to-spoke-vpn-2 bind-interface st0.1
set security ipsec vpn hub-to-spoke-vpn-2 vpn-monitor source-interface ge-0/0/2.0
set security ipsec vpn hub-to-spoke-vpn-2 ike gateway hub-to-spoke-gw-2
set security ipsec vpn hub-to-spoke-vpn-2 ike ipsec-policy vpn-policy
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces ge-0/0/2.0
set security zones security-zone untrust interfaces st0.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/3.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the hub:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 1.1.1.1/30
user@host# set ge-0/0/2 unit 0 family inet address 1.1.2.1/30
user@host# set ge-0/0/3 unit 0 family inet address 50.50.50.1/24
user@host# set st0 unit 0 multipoint
user@host# set st0 unit 0 family inet address 10.10.10.1/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 20.20.20.1/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement lan_nw from interface ge-0/0/3.0
user@host# set policy-statement lan_nw then accept


[edit protocols bgp]
user@host# set group ibgp-1 type internal
user@host# set group ibgp-1 local-address 10.10.10.1
user@host# set group ibgp-1 export lan_nw
user@host# set group ibgp-1 cluster 1.2.3.4
user@host# set group ibgp-1 allow 10.10.10.0/24


user@host# set group ibgp-2 type internal
user@host# set group ibgp-2 local-address 20.20.20.1
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set group ibgp-2 export lan_nw
user@host# set group ibgp-2 cluster 1.2.3.5
user@host# set group ibgp-2 allow 20.20.20.0/24

[edit routing-options]
user@host# set static route 2.2.2.0/30 next-hop 1.1.1.2
user@host# set static route 3.3.3.0/30 next-hop 1.1.2.2
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy-1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike policy ike-policy-2]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local2

[edit security ike gateway hub-to-spoke-gw-1]
user@host# set ike-policy ike-policy-1
user@host# set dynamic distinguished-name wildcard OU=SLT
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/1.0

[edit security ike gateway hub-to-spoke-gw-2]
user@host# set ike-policy ike-policy-2
user@host# set dynamic distinguished-name wildcard OU=SBU
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/2.0
```

4. Configure Phase 2 options.

```
[edit security ipsec vpn-monitor]
user@host# set options interval 5
user@host# set options threshold 2

[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit security ipsec vpn hub-to-spoke-vpn-1]
user@host# set bind-interface st0.0
user@host# set vpn-monitor source-interface ge-0/0/1.0
user@host# set ike gateway hub-to-spoke-gw-1
user@host# set ike ipsec-policy vpn-policy

[edit security ipsec vpn hub-to-spoke-vpn-2]
user@host# set bind-interface st0.1
user@host# set vpn-monitor source-interface ge-0/0/2.0
user@host# set ike gateway hub-to-spoke-gw-2
user@host# set ike ipsec-policy vpn-policy
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.0
user@host# set interfaces ge-0/0/1.0
user@host# set interfaces ge-0/0/2.0
user@host# set interfaces st0.1

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/3.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
   http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
   unit 0 {
      family inet {
         address 1.1.1.1/30;
      }
   }
}
   ge-0/0/2 {
```


Copyright © 2016, Juniper Networks, Inc.

```
      unit 0 {
        family inet {
          address 1.1.2.1/30;
        }
      }
    }
    ge-0/0/3 {
      unit 0 {
        family inet {
          address 50.50.50.1/24;
        }
      }
    }
    st0 {
      unit 0 {
        multipoint;
        family inet {
          address 10.10.10.1/24;
        }
      }
      unit 1 {
        multipoint;
        family inet {
          address 20.20.20.1/24;
        }
      }
    }
[edit]
user@host# show policy-options
policy-statement lan_nw {
  from interface ge-0/0/3.0;
  then accept;
}
[edit]
user@host# show protocols
bgp {
  group ibgp-1 {
    type internal;
    local-address 10.10.10.1;
    export lan_nw;
    cluster 1.2.3.4;
    allow 10.10.10.0/24;
  }
  group ibgp-2 {
    type internal;
    local-address 20.20.20.1;
    export lan_nw;
    cluster 1.2.3.5;
    allow 20.20.20.0/24;
  }
}
[edit]
user@host# show routing-options
static {
  route 2.2.2.0/30 next-hop 1.1.1.2;
  route 3.3.3.0/30 next-hop 1.1.2.2;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy-1 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local1;
        }
    }
    policy ike-policy-2 {
        mode main;
        proposals ike-proposal;
        certificate {
            local-certificate Local2;
        }
    }
    gateway hub-to-spoke-gw-1 {
        ike-policy ike-policy-1;
        dynamic {
            distinguished-name {
                wildcard OU=SLT;
            }
            ike-user-type group-ike-id;
        }
        local-identity distinguished-name;
        external-interface ge-0/0/1.0;
    }
    gateway hub-to-spoke-gw-2 {
        ike-policy ike-policy-2;
        dynamic {
            distinguished-name {
                wildcard OU=SBU;
            }
            ike-user-type group-ike-id;
        }
        local-identity distinguished-name;
        external-interface ge-0/0/2.0;
    }
[edit]
user@host# show security ipsec
vpn-monitor-options {
    interval 5;
    threshold 2;
}
    proposal ipsec-proposal {
        protocol esp;
        authentication-algorithm hmac-md5-96;
        encryption-algorithm des-cbc;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    policy vpn-policy {
      perfect-forward-secrecy {
        keys group14;
      }
      proposals ipsec-proposal;
    }
    vpn hub-to-spoke-vpn-1 {
      bind-interface st0.0;
      vpn-monitor {
        source-interface ge-0/0/1.0;
      }
      ike {
        gateway hub-to-spoke-gw-1;
        ipsec-policy vpn-policy;
      }
    }
    vpn hub-to-spoke-vpn-2 {
      bind-interface st0.1;
      vpn-monitor {
        source-interface ge-0/0/2.0;
      }
      ike {
        gateway hub-to-spoke-gw-2;
        ipsec-policy vpn-policy;
      }
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      st0.0;
      ge-0/0/1.0;
      ge-0/0/2.0;
      st0.1;
    }
}
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        ge-0/0/3.0;
```

```
    }
  }
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Spoke 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/1 unit 0 family inet address 2.2.2.1/30
set interfaces fe-0/0/2 unit 0 family inet address 3.3.3.1/30
set interfaces fe-0/0/4 unit 0 family inet address 60.60.60.1/24
set interfaces st0 unit 0 family inet address 10.10.10.2/24
set interfaces st0 unit 1 family inet address 20.20.20.2/24
set policy-options policy-statement default_route from protocol static
set policy-options policy-statement default_route from route-filter 0.0.0.0/0 exact
set policy-options policy-statement default_route then accept
set policy-options policy-statement lan_nw from interface fe-0/0/4.0
set policy-options policy-statement lan_nw then accept
set protocols bgp group ibgp-1 type internal
set protocols bgp group ibgp-1 local-address 10.10.10.2
set protocols bgp group ibgp-1 export lan_nw
set protocols bgp group ibgp-1 neighbor 10.10.10.1
set protocols bgp group ibgp-2 type internal
set protocols bgp group ibgp-2 local-address 20.20.20.2
set protocols bgp group ibgp-2 export default_route
set protocols bgp group ibgp-2 neighbor 20.20.20.1
set routing-options static route 1.1.1.0/30 next-hop 2.2.2.2
set routing-options static route 1.1.2.0/30 next-hop 3.3.3.2
set routing-options static route 0.0.0.0/0 next-hop st0.1
set routing-options autonomous-system 10
set security ike proposal ike-proposal authentication-method rsa-signatures
set security ike proposal ike-proposal dh-group group2
set security ike proposal ike-proposal authentication-algorithm sha1
set security ike proposal ike-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-policy-1 mode main
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ike policy ike-policy-1 proposals ike-proposal
set security ike policy ike-policy-1 certificate local-certificate Local1
set security ike policy ike-policy-2 mode main
set security ike policy ike-policy-2 proposals ike-proposal
set security ike policy ike-policy-2 certificate local-certificate Local2
set security ike gateway spoke-to-hub-gw-1 ike-policy ike-policy-1
set security ike gateway spoke-to-hub-gw-1 address 1.1.1.1
set security ike gateway spoke-to-hub-gw-1 local-identity distinguished-name
set security ike gateway spoke-to-hub-gw-1 remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw-1 external-interface fe-0/0/1.0
set security ike gateway spoke-to-hub-gw-2 ike-policy ike-policy-2
set security ike gateway spoke-to-hub-gw-2 address 1.1.2.1
set security ike gateway spoke-to-hub-gw-2 local-identity distinguished-name
set security ike gateway spoke-to-hub-gw-2 remote-identity distinguished-name
set security ike gateway spoke-to-hub-gw-2 external-interface fe-0/0/2.0
set security ipsec vpn-monitor-options interval 5
set security ipsec vpn-monitor-options threshold 2
set security ipsec proposal ipsec-proposal protocol esp
set security ipsec proposal ipsec-proposal authentication-algorithm hmac-md5-96
set security ipsec proposal ipsec-proposal encryption-algorithm des-cbc
set security ipsec policy vpn-policy perfect-forward-secrecy keys group14
set security ipsec policy vpn-policy proposals ipsec-proposal
set security ipsec vpn spoke-to-hub-1 bind-interface st0.0
set security ipsec vpn spoke-to-hub-1 vpn-monitor destination-ip 1.1.1.1
set security ipsec vpn spoke-to-hub-1 ike gateway spoke-to-hub-gw-1
set security ipsec vpn spoke-to-hub-1 ike ipsec-policy vpn-policy
set security ipsec vpn spoke-to-hub-1 establish-tunnels immediately
set security ipsec vpn spoke-to-hub-2 bind-interface st0.1
set security ipsec vpn spoke-to-hub-2 vpn-monitor destination-ip 1.1.2.1
set security ipsec vpn spoke-to-hub-2 ike gateway spoke-to-hub-gw-2
set security ipsec vpn spoke-to-hub-2 ike ipsec-policy vpn-policy
set security ipsec vpn spoke-to-hub-2 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces fe-0/0/1.0
set security zones security-zone untrust interfaces st0.0
set security zones security-zone untrust interfaces fe-0/0/2.0
set security zones security-zone untrust interfaces st0.1
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces fe-0/0/4.0
set security policies default-policy permit-all
set security pki ca-profile ca-profile1 ca-identity ca-profile1
set security pki ca-profile ca-profile1 enrollment url http://pc4/certsrv/mscep/mscep.dll
set security pki ca-profile ca-profile1 revocation-check disable
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure spoke 1:

1. Configure interfaces.

```
[edit interfaces]
user@host# set fe-0/0/1 unit 0 family inet address 2.2.2.1/30
```

```
user@host# set fe-0/0/2 unit 0 family inet address 3.3.3.1/30
user@host# set fe-0/0/4 unit 0 family inet address 60.60.60.1/24
user@host# set st0 unit 0 family inet address 10.10.10.2/24
user@host# set st0 unit 1 family inet address 20.20.20.2/24
```

2. Configure routing protocol.

```
[edit policy-options]
user@host# set policy-statement default_route from protocol static
user@host# set policy-statement default_route from route-filter 0.0.0.0/0 exact
user@host# set policy-statement default_route then accept
user@host# set policy-statement lan_nw from interface fe-0/0/4.0
user@host# set policy-statement lan_nw then accept

[edit protocols bgp]
user@host# set group ibgp-1 type internal
user@host# set group ibgp-1 local-address 10.10.10.2
user@host# set group ibgp-1 export lan_nw
user@host# set group ibgp-1 neighbor 10.10.10.1

user@host# set group ibgp-2 type internal
user@host# set group ibgp-2 local-address 20.20.20.2
user@host# set group ibgp-2 export default_route
user@host# set group ibgp-2 neighbor 20.20.20.1

[edit routing-options]
user@host# set static route 1.1.1.0/30 next-hop 2.2.2.2
user@host# set static route 1.1.2.0/30 next-hop 3.3.3.2
user@host# set static route 0.0.0.0/0 next-hop st0.1
user@host# set autonomous-system 10
```

3. Configure Phase 1 options.

```
[edit security ike proposal ike-proposal]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group2
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-128-cbc

[edit security ike policy ike-policy-1]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local1

[edit security ike policy ike-policy-2]
user@host# set mode main
user@host# set proposals ike-proposal
user@host# set certificate local-certificate Local2

[edit security ike gateway spoke-to-hub-gw-1]
user@host# set ike-policy ike-policy-1
user@host# set address 1.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security ike gateway spoke-to-hub-gw-2]
user@host# set ike-policy ike-policy-2
user@host# set address 1.1.2.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name
user@host# set external-interface fe-0/0/2.0
```

4. Configure Phase 2 options.

```
[edit security ipsec vpn-monitor]
user@host# set options interval 5
user@host# set options threshold 2

[edit security ipsec proposal ipsec-proposal]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-md5-96
user@host# set encryption-algorithm des-cbc

[edit security ipsec policy vpn-policy]
user@host# set perfect-forward-secrecy keys group14
user@host# set proposals ipsec-proposal

[edit security ipsec vpn spoke-to-hub-1]
user@host# set bind-interface st0.0
user@host# set vpn-monitor destination-ip 1.1.1.1
user@host# set ike gateway spoke-to-hub-gw-1
user@host# set ike ipsec-policy vpn-policy
user@host# set establish-tunnels immediately

[edit security ipsec vpn spoke-to-hub-2]
user@host# set bind-interface st0.1
user@host# set vpn-monitor destination-ip 1.1.2.1
user@host# set ike gateway spoke-to-hub-gw-2
user@host# set ike ipsec-policy vpn-policy
user@host# set establish-tunnels immediately
```

5. Configure zones.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/1.0
user@host# set interfaces st0.0
user@host# set interfaces fe-0/0/2.0
user@host# set interfaces st0.1

[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces fe-0/0/4.0
```

6. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

7. Configure the CA profile.

```
[edit security pki]
user@host# set ca-profile ca-profile1 ca-identity ca-profile1
user@host# set ca-profile ca-profile1 enrollment url
  http://pc4/certsrv/mscep/mscep.dll
user@host# set ca-profile ca-profile1 revocation-check disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security zones**, **show security policies**, and **show security pki** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 2.2.2.1/30;
    }
  }
}
  fe-0/0/2 {
    unit 0 {
      family inet {
        address 3.3.3.1/30;
      }
    }
  }
  fe-0/0/4 {
    unit 0 {
      family inet {
        address 60.60.60.1/24;
      }
    }
  }
  st0 {
    unit 0 {
      family inet {
        address 10.10.10.2/24;
      }
    }
    unit 1 {
      family inet {
        address 20.20.20.2/24;
      }
    }
  }
[edit]
user@host# show policy-options
policy-statement default_route {
  from {
    protocol static;
    route-filter 0.0.0.0/0 exact;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    then accept;
}
    policy-statement lan_nw {
      from interface fe-0/0/4.0;
      then accept;
    }
[edit]
user@host# show protocols
bgp {
    group ibgp-1 {
      type internal;
      local-address 10.10.10.2;
      export lan_nw;
      neighbor 10.10.10.1;
    }
    group ibgp-2 {
      type internal;
      local-address 20.20.20.2;
      export default_route;
      neighbor 20.20.20.1;
    }
}
[edit]
user@host# show routing-options
static {
    route 1.1.1.0/30 next-hop 2.2.2.2;
    route 1.1.2.0/30 next-hop 3.3.3.2;
    route 0.0.0.0/0 next-hop st0.1;
    }
autonomous-system 10;
[edit]
user@host# show security ike
proposal ike-proposal {
    authentication-method rsa-signatures;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
}
    policy ike-policy-1 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local1;
      }
    }
    policy ike-policy-2 {
      mode main;
      proposals ike-proposal;
      certificate {
        local-certificate Local2;
      }
    }
    gateway spoke-to-hub-gw-1 {
      ike-policy ike-policy-1;
      address 1.1.1.1;
```

Copyright © 2016, Juniper Networks, Inc.

```
        local-identity distinguished-name;
        remote-identity distinguished-name;
        external-interface fe-0/0/1.0;
    }
    gateway spoke-to-hub-gw-2 {
        ike-policy ike-policy-2;
        address 1.1.2.1;
        local-identity distinguished-name;
        remote-identity distinguished-name;
        external-interface fe-0/0/2.0;
    }
[edit]
user@host# show security ipsec
vpn-monitor-options {
    interval 5;
    threshold 2;
}
    proposal ipsec-proposal {
        protocol esp;
        authentication-algorithm hmac-md5-96;
        encryption-algorithm des-cbc;
    }
    policy vpn-policy {
        perfect-forward-secrecy {
            keys group14;
        }
        proposals ipsec-proposal;
    }
    vpn spoke-to-hub-1 {
        bind-interface st0.0;
        vpn-monitor {
            destination-ip 1.1.1.1;
        }
        ike {
            gateway spoke-to-hub-gw-1;
            ipsec-policy vpn-policy;
        }
        establish-tunnels immediately;
    }
    vpn spoke-to-hub-2 {
        bind-interface st0.1;
        vpn-monitor {
            destination-ip 1.1.2.1;
        }
        ike {
            gateway spoke-to-hub-gw-2;
            ipsec-policy vpn-policy;
        }
        establish-tunnels immediately;
    }
[edit]
user@host# show security zones
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    protocols {
      all;
    }
  }
  interfaces {
    fe-0/0/1.0;
    st0.0;
    fe-0/0/2.0;
    st0.1;
  }
}
  security-zone trust {
    host-inbound-traffic {
      system-services {
        all;
      }
      protocols {
        all;
      }
    }
    interfaces {
      fe-0/0/4.0;
    }
  }
[edit]
user@host# show security policies
default-policy {
  permit-all;
}
[edit]
user@host# show security pki
ca-profile ca-profile1 {
  ca-identity ca-profile1;
  enrollment {
    url http://pc4/certsrv/mscep/mscep.dll;
  }
  revocation-check {
    disable;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying IKE Phase 1 Status (Both Tunnels Are Up) on page 6554
- Verifying IPsec Phase 2 Status (Both Tunnels Are Up) on page 6554
- Verifying IPsec Next-Hop Tunnels (Both Tunnels Are Up) on page 6554
- Verifying BGP (Both Tunnels Are Up) on page 6555
- Verifying Learned Routes (Both Tunnels Are Up) on page 6555

Copyright © 2016, Juniper Networks, Inc.

- Verifying IKE Phase 1 Status (Primary Tunnel Is Down) on page 6556
- Verifying IPsec Phase 2 Status (Primary Tunnel Is Down) on page 6556
- Verifying IPsec Next-Hop Tunnels (Primary Tunnel Is Down) on page 6557
- Verifying BGP (Primary Tunnel Is Down) on page 6557
- Verifying Learned Routes (Primary Tunnel Is Down) on page 6557

Verifying IKE Phase 1 Status (Both Tunnels Are Up)

Purpose Verify the IKE Phase 1 status when both IPSec VPN tunnels are up.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

3733075 UP     d4f51c28c0a82101  05b125993a864d3c  Main           3.3.3.1

3733076 UP     d53c8a0b7d4c319b  c23c5f7a26388247  Main           2.2.2.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spoke.

Verifying IPsec Phase 2 Status (Both Tunnels Are Up)

Purpose Verify the IPsec Phase 2 status when both IPsec VPN tunnels are up.

Action From operational mode, enter the **security ipsec security-associations** command.

```
user@host> security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <268173316 ESP:des/ md5 3cd96946 3555/ unlim U   root 500   2.2.2.1
  >268173316 ESP:des/ md5 1c09b9b 3555/ unlim  U   root 500   2.2.2.1
  <268173313 ESP:des/ md5 7c6ffca3 3340/ unlim U   root 500   3.3.3.1
  >268173313 ESP:des/ md5 33bf6f2f 3340/ unlim U   root 500   3.3.3.1
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spoke.

Verifying IPsec Next-Hop Tunnels (Both Tunnels Are Up)

Purpose Verify the IPsec next-hop tunnels.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPSec VPN name                    Flag     IKE-ID
                            XAUTH username
10.10.10.2        st0.0       hub-to-spoke-vpn-1               Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SLT, CN=spoke1
20.20.20.2        st0.1       hub-to-spoke-vpn-2               Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SBU, CN=spoke1_backup
```

Meaning The next-hop gateways are the IP addresses for the **st0** interfaces of the spoke. The next hop should be associated with the correct IPsec VPN name.

Verifying BGP (Both Tunnels Are Up)

Purpose Verify that BGP references the IP addresses for the **st0** interfaces of the spoke when both IPsec VPN tunnels are up.

Action From operational mode, enter the **show bgp summary** command.

```
user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Unconfigured peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2               10          5          6       0       0          54
1/1/1/0              0/0/0/0
20.20.20.2               10         13         16       0       0        4:29
1/1/1/0              0/0/0/0
```

Verifying Learned Routes (Both Tunnels Are Up)

Purpose Verify that routes to the spoke have been learned when both tunnels are up. The route to 60.60.60.0/24 is through the st0.0 interface and the default route is through the st0.1 interface.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route 60.60.60.0** command.

```
user@host> show route 60.60.60.0
inet.0: 48 destinations, 48 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

60.60.60.0/24      *[BGP/170] 00:01:11, localpref 100
                      AS path: I
                    > to 10.10.10.2 via st0.0
```

From operational mode, enter the **show route 0.0.0.0** command.

```
user@host> show route 0.0.0.0
inet.0: 48 destinations, 48 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:04:55, localpref 100
                      AS path: I
                    > to 20.20.20.2 via st0.1
```

Verifying IKE Phase 1 Status (Primary Tunnel Is Down)

Purpose Verify the IKE Phase 1 status when the primary tunnel is down.

Action From operational mode, enter the **show security ike security-associations** command.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

3733075 UP     d4f51c28c0a82101  05b125993a864d3c  Main           3.3.3.1

3733076 UP     d53c8a0b7d4c319b  c23c5f7a26388247  Main           2.2.2.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. If no SAs are listed, there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spoke.

Verifying IPsec Phase 2 Status (Primary Tunnel Is Down)

Purpose Verify the IPsec Phase 2 status when the primary tunnel is down.

Action From operational mode, enter the **security ipsec security-associations** command.

```
user@host> security ipsec security-associations
   Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  <268173313 ESP:des/ md5 7c6ffca3 3156/ unlim U   root 500   3.3.3.1
  >268173313 ESP:des/ md5 33bf6f2f 3156/ unlim U   root 500   3.3.3.1
```

Meaning The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. If no SAs are listed, there was a problem with Phase 2 establishment. Check the IKE policy

Copyright © 2016, Juniper Networks, Inc.

parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spoke.

Verifying IPsec Next-Hop Tunnels (Primary Tunnel Is Down)

Purpose Verify the IPsec next-hop tunnel.

Action From operational mode, enter the **show security ipsec next-hop-tunnels** command.

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPSec VPN name                    Flag     IKE-ID
                              XAUTH username
20.20.20.2        st0.1       hub-to-spoke-vpn-2                Auto     C=IN,
DC=example.net, ST=KA, L=Mysore, O=example, OU=SBU, CN=spoke1_backup
```

Meaning The next-hop gateways are the IP addresses for the **st0** interfaces of the spoke. The next hop should be associated with the correct IPsec VPN name, in this case the backup VPN tunnel.

Verifying BGP (Primary Tunnel Is Down)

Purpose Verify that BGP references the IP addresses for the **st0** interfaces of the spoke when the primary tunnel is down.

Action From operational mode, enter the **show bgp summary** command.

```
user@host> show bgp summary
Groups: 2 Peers: 1 Down peers: 0
Unconfigured peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
20.20.20.2               10         20         24       0       0        7:24
1/1/1/0             0/0/0/0
```

Verifying Learned Routes (Primary Tunnel Is Down)

Purpose Verify that routes to the spoke have been learned when the primary tunnel is down. Both the route to 60.60.60.0/24 and the default route are through the st0.1 interface.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route 60.60.60.0** command.

```
user@host> show route 60.60.60.0
inet.0: 46 destinations, 46 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:07:41, localpref 100
                      AS path: I
                    > to 20.20.20.2 via st0.1
```

From operational mode, enter the **show route 0.0.0.0** command.

```
user@host> show route 0.0.0.0
inet.0: 46 destinations, 46 routes (45 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/170] 00:07:47, localpref 100
                      AS path: I
                    > to 20.20.20.2 via st0.1
```

Related Documentation

- Example: Configuring a Route-Based VPN on page 6074
- *Routing Protocols Overview for Security Devices*

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 273

Configuring Auto Discovery VPNs

- Understanding Auto Discovery VPN on page 6559
- Understanding Traffic Routing with Shortcut Tunnels on page 6564
- Example: Improving Network Resource Utilization with Auto Discovery VPN Dynamic Tunnels on page 6566
- Enabling OSPF to Update Routes Quickly After ADVPN Shortcut Tunnels Are Established on page 6604

Understanding Auto Discovery VPN

AutoVPN deployments can use the Auto Discovery VPN (ADVPN) protocol to dynamically establish spoke-to-spoke VPN tunnels. When passing traffic from one spoke to another spoke, the hub can suggest that the spokes establish a direct security association (SA), called a shortcut, between each other. Shortcuts can be established and torn down dynamically between spokes, resulting in better network resource utilization and less reliance on a centrally located hub.

- ADVPN Protocol on page 6559
- Establishing a Shortcut on page 6560
- Shortcut Initiator and Responder Roles on page 6561
- Shortcut Attributes on page 6561
- Shortcut Termination on page 6562
- ADVPN Configuration Limitations on page 6563

ADVPN Protocol

The ADVPN protocol is an extension of IKEv2 that allows a shortcut to be created between two VPN peers. Devices that support the ADVPN protocol send an ADVPN_SUPPORTED notification in the IKEv2 Notify payload during the initial IKE exchange. A device that supports ADVPN can act as either a shortcut suggester or a shortcut partner, but not both. This shortcut capability information, along with the ADVPN version number, is also exchanged.

Copyright © 2016, Juniper Networks, Inc.

Establishing a Shortcut

An IPsec VPN gateway can act as a shortcut suggester when it notices that traffic is exiting a tunnel with one of its peers and entering a tunnel with another peer. Figure 287 shows traffic from Spoke 1 to Spoke 3 passing through the hub.

Figure 287: Spoke-to-Spoke Traffic Passing Through Hub

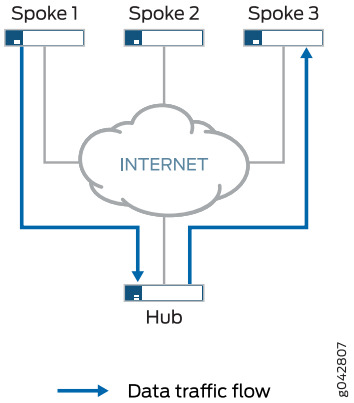

When ADVPN is configured on the devices, ADVPN shortcut capability information is exchanged between the hub and spokes. As long as Spokes 1 and 3 have previously advertised ADVPN shortcut partner capability to the hub, the hub can suggest that Spokes 1 and 3 establish a shortcut between each other.

The shortcut suggester uses its already established IKEv2 SAs with the peers to begin a shortcut exchange with one of the two peers. If the peer accepts the shortcut exchange, then the shortcut suggester begins a shortcut exchange with the other peer. The shortcut exchange includes information to allow the peers (referred to as shortcut partners) to establish IKE and IPsec SAs with each other. The creation of the shortcut between the shortcut partners starts only after both peers accept the shortcut exchange.

Figure 288 shows traffic passing through a shortcut between Spokes 1 and 3. Traffic from Spoke 1 to Spoke 3 does not need to traverse the hub.

Copyright © 2016, Juniper Networks, Inc.

Figure 288: Spoke-to-Spoke Traffic Passing Through Shortcut

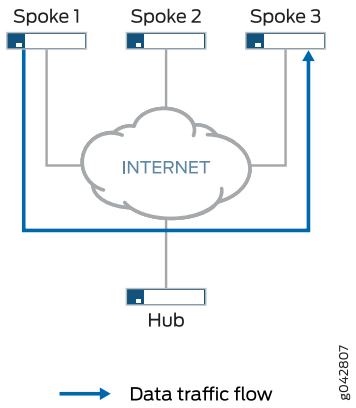

Shortcut Initiator and Responder Roles

The shortcut suggester chooses one of the shortcut partners to act as the initiator for the shortcut; the other partner acts as the responder. If one of the partners is behind a NAT device, then the partner behind the NAT device is chosen as the initiator. If none of the partners is behind a NAT device, then the suggester randomly chooses one of the partners as the initiator; the other partner acts as the responder. If both partners are behind NAT devices, then a shortcut cannot be created between them; the suggester does not send a shortcut exchange to any of the peers.

The shortcut suggester begins the shortcut exchange with the responder first. If the responder accepts the shortcut suggestion, then the suggester notifies the initiator.

Using information contained in the shortcut suggester's notification, the shortcut initiator establishes an IKEv2 exchange with the responder, and a new IPsec SA is established between the two partners. On each partner, the route to the network behind its partner now points to the shortcut instead of to the tunnel between the partner and the suggester. Traffic originating behind one of the partners that is destined to a network behind the other shortcut partner flows over the shortcut.

If the partners decline the shortcut suggestion, then the partners notify the suggester with the reason for the rejection. In this case, traffic between the partners continues to flow through the shortcut suggester.

Shortcut Attributes

The shortcut receives some of its attributes from the shortcut suggester while other attributes are inherited from the suggester-partner VPN tunnel configuration. Table 564 shows the parameters of the shortcut.

Table 564: Shortcut Parameters

Attributes	Received/Inherited From
ADVPN	Configuration

Copyright © 2016, Juniper Networks, Inc.

Table 564: Shortcut Parameters *(continued)*

Attributes	Received/Inherited From
Antireplay	Configuration
Authentication algorithm	Configuration
Dead peer detection	Configuration
DF bit	Configuration
Encryption algorithm	Configuration
Establish tunnels	Suggester
External interface	Configuration
Gateway policy	Configuration
General IKE ID	Configuration
IKE version	Configuration
Install interval	Configuration
Local address	Configuration
Local identity	Suggester
NAT traversal	Configuration
Perfect forward secrecy	Configuration
Protocol	Configuration
Proxy ID	Not applicable
Remote address	Suggester
Remote identity	Suggester
Respond bad SPI	Configuration
Traffic selector	Not applicable

Shortcut Termination

By default, the shortcut lasts indefinitely. Shortcut partners terminate the shortcut if traffic falls below a specified rate for a specified time. By default, the shortcut is terminated if traffic falls below 5 packets per second for 900 seconds; the idle time and

Copyright © 2016, Juniper Networks, Inc.

idle threshold values are configurable for partners. The shortcut can be manually deleted on either shortcut partner with the **clear security ike security-association** or **clear security ipsec security-association** commands to clear the corresponding IKE or IPsec SA. Either of the shortcut partners can terminate the shortcut at any time by sending an IKEv2 delete payload to the other shortcut partner.

When the shortcut is terminated, the corresponding IKE SA and all child IPsec SAs are deleted. After the shortcut is terminated, the corresponding route is deleted on both shortcut partners and traffic between the two peers again flows through the suggester. Shortcut termination information is sent from a partner to the suggester.

The lifetime of a shortcut is independent of the tunnel between the shortcut suggester and shortcut partner. The shortcut is not terminated simply because the tunnel between the suggester and partner is terminated.

ADVPN Configuration Limitations

Note the following limitations when configuring ADVPN:

- Configuring an ADVPN partner is only allowed on site-to-site VPNs. Configuring an ADVPN suggester is only allowed on AutoVPN hubs.
- You cannot configure both suggester and partner roles on the same gateway. When ADVPN is enabled on a gateway, you cannot disable both suggester and partner roles on the gateway.
- As mentioned previously, you cannot create a shortcut between partners that are both behind NAT devices. The suggester can initiate a shortcut exchange if only one of the partners is behind a NAT device or if no partners are behind NAT devices.
- Only the OSPF dynamic routing protocol is supported with ADVPN; RIP and BGP are not supported

The following configurations are not supported with ADVPN:

- IKEv1
- Policy-based VPN
- IKEv2 configuration payload
- Traffic selectors
- Preshared key
- Point-to-point secure tunnel interfaces

Related Documentation

- Understanding Traffic Routing with Shortcut Tunnels on page 6564
- Example: Improving Network Resource Utilization with Auto Discovery VPN Dynamic Tunnels on page 6566
- Enabling OSPF to Update Routes Quickly After ADVPN Shortcut Tunnels Are Established on page 6604

Copyright © 2016, Juniper Networks, Inc.

Understanding Traffic Routing with Shortcut Tunnels

Tunnel flaps or catastrophic changes can cause both static tunnels and shortcut tunnels to go down. When this happens, traffic to a specific destination might be routed through an unexpected shortcut tunnel instead of through an expected static tunnel.

In Figure 289, static tunnels exist between the hub and each of the spokes. OSPF adjacencies are established between the hub and spokes. Spoke A also has a shortcut tunnel with Spoke B and OSPF adjacencies are established between the spokes. The hub (the shortcut suggester) recognizes that if connectivity between the hub and Spoke A goes down, Spoke A's network can be reached through the shortcut tunnel between Spoke B and Spoke A.

Figure 289: Static Tunnels and Shortcut Tunnel Established in Hub-and-Spoke Network

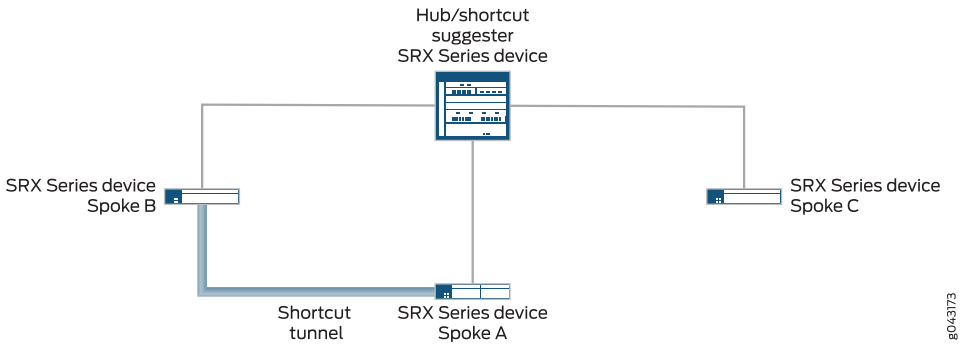

In Figure 290, the static tunnel between the hub and Spoke A is down. If there is new traffic from Spoke C to Spoke A, Spoke C forwards the traffic to the hub because it does not have a shortcut tunnel with Spoke A. The hub does not have an active static tunnel with Spoke A but it recognizes that there is a shortcut tunnel between Spoke A and Spoke B, so it forwards the traffic from Spoke C to Spoke B.

Figure 290: Traffic Path from Spoke C to Spoke A

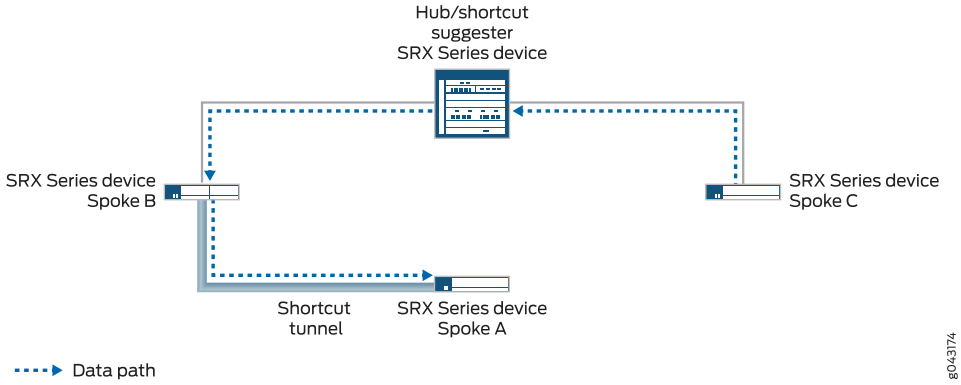

Copyright © 2016, Juniper Networks, Inc.

As long as both Spoke B and Spoke C support Auto Discovery VPN (ADVPN) partner capability, the hub can suggest that the spokes establish a direct shortcut between each other. This occurs even though there is no direct traffic between the two spokes. Traffic from Spoke C to Spoke A travels through the shortcut tunnel between Spoke C and Spoke B, and then through the shortcut tunnel between Spoke B and Spoke A (see Figure 291).

Figure 291: Traffic Path from Spoke C to Spoke A Through Shortcut Tunnels

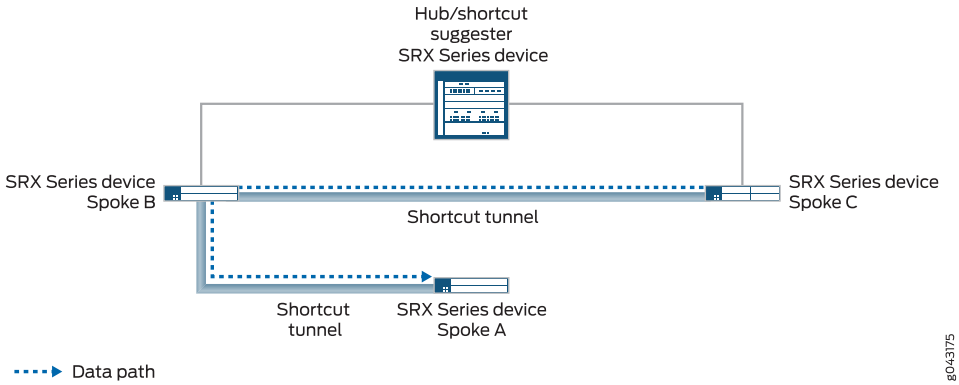

When the static tunnel between the hub and Spoke A is reestablished, the tunnel is advertised to all spokes. Spoke C learns that there is a better route to reach Spoke A; instead of passing traffic through Spoke B, it forwards traffic for Spoke A to the hub. The hub suggests that a shortcut tunnel be established between Spoke C and Spoke A. When the shortcut tunnel is established between Spoke C and Spoke A, traffic flows through the shortcut tunnel (see Figure 292). Traffic between Spoke C and Spoke A no longer travels through Spoke B, and the shortcut tunnel between Spoke B and Spoke C eventually disappears.

Figure 292: Traffic Path from Spoke C to Spoke A Through Shortcut Tunnel

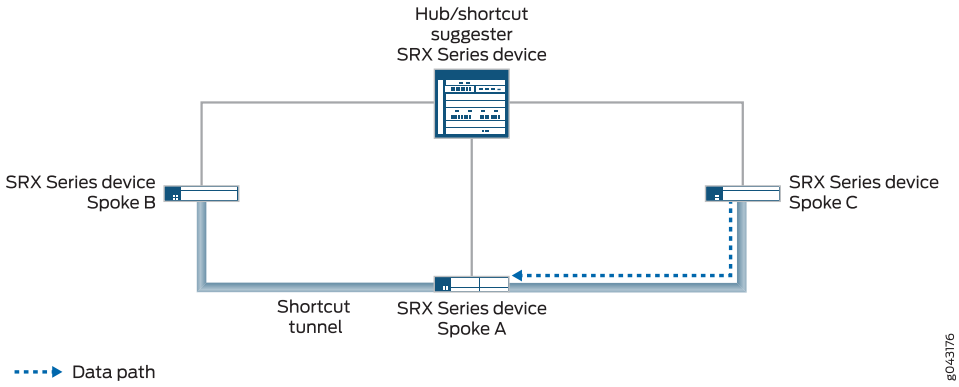

NOTE: You can use the connection-limit option at the [edit security ike gateway *gateway-name* advpn partner] hierarchy level to set the maximum number of shortcut tunnels that can be created with different shortcut partners using a particular gateway. The maximum number, which is also the default, is platform-dependent.

Related Documentation

- Understanding Auto Discovery VPN on page 6559

Example: Improving Network Resource Utilization with Auto Discovery VPN Dynamic Tunnels

If you are deploying an AutoVPN network, you might be able to increase your network resource utilization by configuring Auto Discovery VPN (ADVPN). In AutoVPN networks, VPN traffic flows through the hub even when the traffic is travelling from one spoke to another. ADVPN allows VPN tunnels to be established dynamically between spokes, which can result in better network resource utilization. Use this example to configure ADVPN to enable dynamic spoke-to-spoke VPN tunnels in your AutoVPN network.

- Requirements on page 6566
- Overview on page 6567
- Configuration on page 6569
- Verification on page 6588

Requirements

This example uses the following hardware and software components:

- Three supported SRX Series devices as AutoVPN hub and spokes.
- Junos OS Release 12.3X48-D10 or later releases that support ADVPN.
- Digital certificates enrolled in the hub and spokes that allow the devices to authenticate each other.

Before you begin:

1. Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates. See "Understanding Local Certificate Requests" on page 6371.
2. Enroll the digital certificates in each device. See "Understanding Certificate Loading" on page 6377.

Copyright © 2016, Juniper Networks, Inc.

NOTE: This example uses the OSPF dynamic routing protocol as well as static route configurations to forward packets through VPN tunnels. You should be familiar with the OSPF dynamic routing protocol that is used to forward packets through the VPN tunnels. For more information about specific requirements for using dynamic routing protocols, see the *Routing Protocols Overview for Security Devices*.

Overview

This example shows the configurations of an AutoVPN hub and two spokes for ADVPN. The spokes establish IPsec VPN connections to the hub, which allows them to communicate with each other as well as to access resources on the hub. While traffic is initially passed from one spoke to the other through the hub, ADVPN allows the spokes to establish a direct security association between each other. The hub acts as the shortcut suggester. On the hub, the ADVPN configuration disables the **partner** role. On the spokes, ADVPN configuration disables the **suggester** role.

Certain Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hub and spokes must have the same values. Table 565 shows the values used in this example.

Table 565: Phase 1 and Phase 2 Options for AutoVPN Hub and Spokes for ADVPN Example

Option	Value
IKE proposal:	
Authentication method	rsa-signatures
Diffie-Hellman (DH) group	group5
Authentication algorithm	sha1
Encryption algorithm	aes-256-cbc
IKE policy:	
Certificate	local-certificate
IKE gateway:	
Version	v2-only
IPsec proposal:	
Protocol	esp
Authentication algorithm	hmac-sha1-96
Encryption algorithm	aes-256-cbc

Copyright © 2016, Juniper Networks, Inc.

Table 565: Phase 1 and Phase 2 Options for AutoVPN Hub and Spokes for ADVPN Example *(continued)*

Option	Value
IPsec policy:	
Perfect Forward Secrecy (PFS) group	group5

The IKE gateway configuration on the hub and spokes include remote and local values that identify VPN peers. Table 566 shows the IKE gateway configuration for the hub and spokes in this example.

Table 566: IKE Gateway Configuration for ADVPN Example

Option	Hub	Spokes
Remote IP address	Dynamic	Spoke 1: 11.1.1.1 Spoke 2: 11.1.1.1
Local IP address	11.1.1.1	Spoke 1: 21.1.1.2 Spoke 2: 31.1.1.2
Remote IKE ID	Distinguished name (DN) with the string "XYZ" in the organization (O) field and "Sales" in the organization unit (OU) field in the spokes' certificates	DN with the string "Sales" in the OU field in the hub's certificate
Local IKE ID	DN on the hub's certificate	DN on the spokes' certificate

The hub authenticates the spokes' IKE ID if the subject fields of the spokes' certificates contain the string "XYZ" in the O field and "Sales" in the OU field.

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085.

Topology

Figure 293 shows the SRX Series devices to be configured for this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 293: AutoVPN Deployment with ADVPN

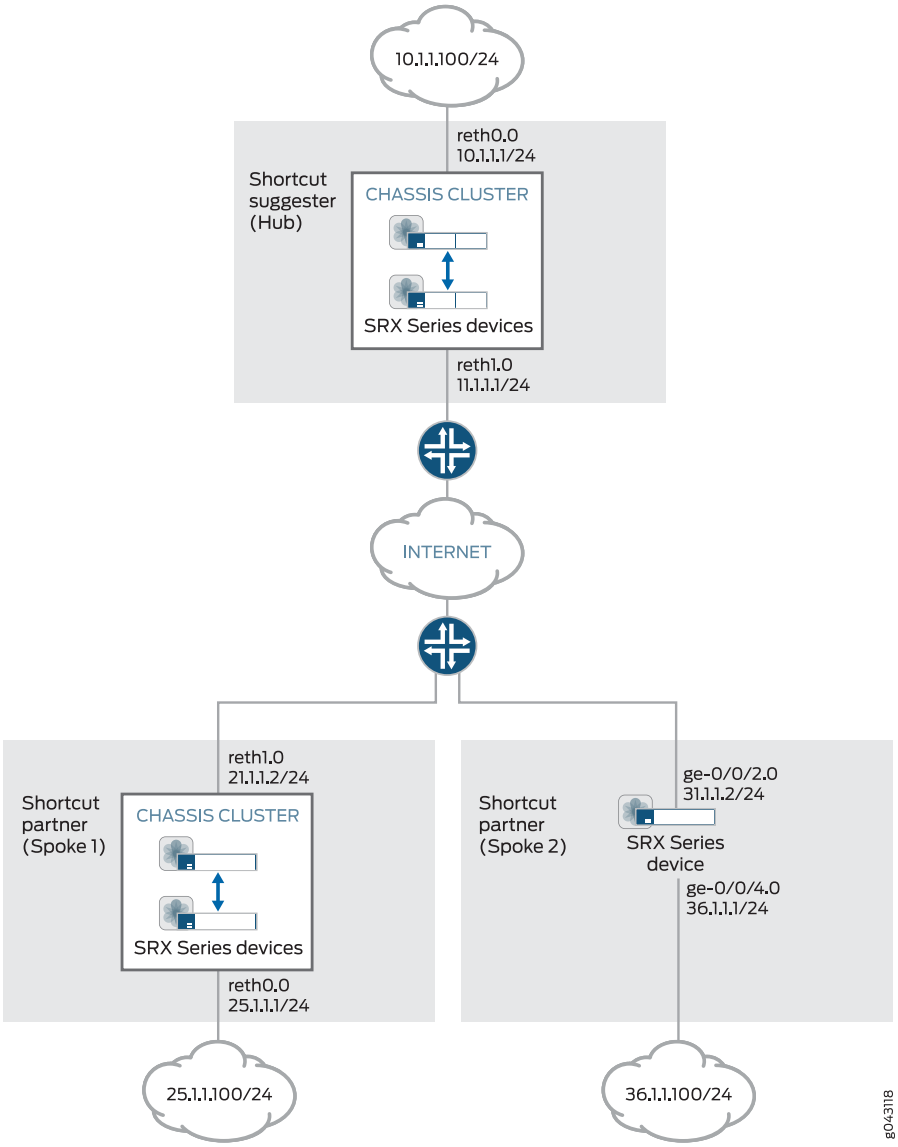

Configuration

- Configuring the Suggester (Hub) on page 6569
- Configuring the Partner (Spoke 1) on page 6576
- Configuring the Partner (Spoke 2) on page 6582

Configuring the Suggester (Hub)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/3 gigether-options redundant-parent reth0
set interfaces ge-0/0/4 gigether-options redundant-parent reth1
set interfaces ge-7/0/3 gigether-options redundant-parent reth0
set interfaces ge-7/0/4 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 10.1.1.1/24
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 11.1.1.1/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 172.16.1.1/24
set protocols ospf graceful-restart restart-duration 300
set protocols ospf graceful-restart notify-duration 300
set protocols ospf graceful-restart no-strict-lsa-checking
set protocols ospf area 0.0.0.0 interface st0.1 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.1 metric 10
set protocols ospf area 0.0.0.0 interface st0.1 retransmit-interval 1
set protocols ospf area 0.0.0.0 interface st0.1 dead-interval 40
set protocols ospf area 0.0.0.0 interface st0.1 demand-circuit
set protocols ospf area 0.0.0.0 interface st0.1 dynamic-neighbors
set protocols ospf area 0.0.0.0 interface reth0.0
set routing-options graceful-restart
set routing-options static route 21.1.1.0/24 next-hop 11.1.1.2
set routing-options static route 31.1.1.0/24 next-hop 11.1.1.2
set routing-options router-id 172.16.1.1
set security ike proposal IKE_PROP authentication-method rsa-signatures
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL certificate local-certificate Suggester_Certificate_ID
set security ike gateway SUGGESTER_GW ike-policy IKE_POL
set security ike gateway SUGGESTER_GW dynamic distinguished-name wildcard
  O=XYZ,OU=Sales
set security ike gateway SUGGESTER_GW dynamic ike-user-type group-ike-id
set security ike gateway SUGGESTER_GW dead-peer-detection
set security ike gateway SUGGESTER_GW local-identity distinguished-name
set security ike gateway SUGGESTER_GW external-interface reth1.0
set security ike gateway SUGGESTER_GW local-address 11.1.1.1
set security ike gateway SUGGESTER_GW advpn partner disable
set security ike gateway SUGGESTER_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn SUGGESTER_VPN bind-interface st0.1
set security ipsec vpn SUGGESTER_VPN ike gateway SUGGESTER_GW
set security ipsec vpn SUGGESTER_VPN ike ipsec-policy IPSEC_POL
set security pki ca-profile advpn ca-identity advpn
set security pki ca-profile advpn enrollment url http://10.157.92.176:8080/scep/advpn/
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces reth0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces reth1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the suggester:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/3 gigether-options redundant-parent reth0
user@host# set ge-0/0/4 gigether-options redundant-parent reth1
user@host# set ge-7/0/3 gigether-options redundant-parent reth0
user@host# set ge-7/0/4 gigether-options redundant-parent reth1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 10.1.1.1/24
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 11.1.1.1/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 172.16.1.1/24
```

2. Configure the routing protocol and static routes.

```
[edit protocols ospf]
user@host# set graceful-restart restart-duration 300
user@host# set graceful-restart notify-duration 300
user@host# set graceful-restart no-strict-lsa-checking
user@host# set area 0.0.0.0 interface st0.1 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.1 metric 10
user@host# set area 0.0.0.0 interface st0.1 retransmit-interval 1
user@host# set area 0.0.0.0 interface st0.1 dead-interval 40
user@host# set area 0.0.0.0 interface st0.1 demand-circuit
user@host# set area 0.0.0.0 interface st0.1 dynamic-neighbors
user@host# set area 0.0.0.0 interface reth0.0

[edit routing-options]
user@host# set graceful-restart
user@host# set static route 21.1.1.0/24 next-hop 11.1.1.2
user@host# set static route 31.1.1.0/24 next-hop 11.1.1.2
user@host# set router-id 172.16.1.1
```

3. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set proposals IKE_PROP
user@host# set certificate local-certificate Suggester_Certificate_ID

[edit security ike gateway SUGGESTER_GW]
user@host# set ike-policy IKE_POL
user@host# set dynamic distinguished-name wildcard O=XYZ,OU=Sales
user@host# set dynamic ike-user-type group-ike-id
user@host# set dead-peer-detection
user@host# set local-identity distinguished-name
user@host# set external-interface reth1.0
user@host# set local-address 11.1.1.1
user@host# set advpn partner disable
user@host# set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy IPSEC_POL]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals IPSEC_PROP

[edit security isec vpn SUGGESTER_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway SUGGESTER_GW
user@host# set ike ipsec-policy IPSEC_POL
```

5. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile advpn ca-identity advpn
user@host# set ca-profile advpn enrollment url
  http://10.157.92.176:8080/scep/advpn/
```

6. Configure zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces reth0.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces reth1.0
```

7. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```


Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
  user@host# show interfaces
  ge-0/0/3 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-0/0/4 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-7/0/3 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/4 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  reth0 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 10.1.1.1/24;
      }
    }
  }
  reth1 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 11.1.1.1/24;
      }
    }
  }
  st0 {
    unit 1 {
      multipoint;
      family inet {
        address 172.16.1.1/24;
      }
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show protocols
ospf {
  graceful-restart {
    restart-duration 300;
    notify-duration 300;
    no-strict-lsa-checking;
  }
  area 0.0.0.0 {
    interface st0.1 {
      interface-type p2mp;
      metric 10;
      retransmit-interval 1;
      dead-interval 40;
      demand-circuit;
      dynamic-neighbors;
    }
    interface reth0.0;
  }
}
[edit]
user@host# show routing-options
graceful-restart;
static {
  route 21.1.1.0/24 next-hop 11.1.1.2;
  route 31.1.1.0/24 next-hop 11.1.1.2;
}
router-id 172.16.1.1;
[edit]
user@host# show security ike
proposal IKE_PROP {
  authentication-method rsa-signatures;
  dh-group group5;
  authentication-algorithm sha1;
  encryption-algorithm aes-256-cbc;
}
policy IKE_POL {
  proposals IKE_PROP;
  certificate {
    local-certificate Suggester_Certificate_ID;
  }
}
gateway SUGGESTER_GW {
  ike-policy IKE_POL;
  dynamic {
    distinguished-name {
      wildcard O=XYZ,OU=Sales;
    }
    ike-user-type group-ike-id;
  }
  dead-peer-detection {
  }
  local-identity distinguished-name;
  external-interface reth1.0
  local-address 11.1.1.1;
```


Copyright © 2016, Juniper Networks, Inc.

```
    advpn {
        partner {
            disable;
        }
    }
    version v2-only;
}
[edit]
user@host# show security ipsec
proposal IPSEC_PROP {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
}
policy IPSEC_POL {
    perfect-forward-secrecy {
        keys group5;
    }
    proposals IPSEC_PROP;
}
vpn SUGGESTER_VPN {
    bind-interface st0.1;
    ike {
        gateway SUGGESTER_GW;
        ipsec-policy IPSEC_POL;
    }
}
[edit]
user@host# show security pki
ca-profile advpn {
    ca-identity advpn;
    enrollment {
        url http://10.157.92.176:8080/scep/advpn/;
    }
}
[edit]
user@host# show security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        st0.1;
        reth0.0;
    }
}
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
        protocols {
            all;
        }
    }
    interfaces {
        reth1.0;
    }
}
[edit]
user@host# show security policies
default-policy {
    permit-all;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Partner (Spoke 1)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/3 gigether-options redundant-parent reth0
set interfaces ge-0/0/4 gigether-options redundant-parent reth1
set interfaces ge-7/0/3 gigether-options redundant-parent reth0
set interfaces ge-7/0/4 gigether-options redundant-parent reth1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 25.1.1.1/24
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 21.1.1.2/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 172.16.1.2/24
set protocols ospf graceful-restart restart-duration 300
set protocols ospf graceful-restart notify-duration 300
set protocols ospf graceful-restart no-strict-lsa-checking
set protocols ospf area 0.0.0.0 interface st0.1 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.1 metric 15
set protocols ospf area 0.0.0.0 interface st0.1 retransmit-interval 1
set protocols ospf area 0.0.0.0 interface st0.1 dead-interval 40
set protocols ospf area 0.0.0.0 interface st0.1 demand-circuit
set protocols ospf area 0.0.0.0 interface st0.1 dynamic-neighbors
set protocols ospf area 0.0.0.0 interface reth0.0
set routing-options graceful-restart
set routing-options static route 11.1.1.0/24 next-hop 21.1.1.1
set routing-options static route 31.1.1.0/24 next-hop 21.1.1.1
set routing-options router-id 172.16.1.2
set security ike proposal IKE_PROP authentication-method rsa-signatures
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL certificate local-certificate Partner1_Certificate_ID
set security ike gateway PARTNER_GW ike-policy IKE_POL
set security ike gateway PARTNER_GW address 11.1.1.1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ike gateway PARTNER_GW local-identity distinguished-name
set security ike gateway PARTNER_GW remote-identity distinguished-name container
  OU=Sales
set security ike gateway PARTNER_GW external-interface reth1
set security ike gateway PARTNER_GW local-address 21.1.1.2
set security ike gateway PARTNER_GW advpn suggester disable
set security ike gateway PARTNER_GW advpn partner
set security ike gateway PARTNER_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn PARTNER_VPN bind-interface st0.1
set security ipsec vpn PARTNER_VPN ike gateway PARTNER_GW
set security ipsec vpn PARTNER_VPN ike ipsec-policy IPSEC_POL
set security ipsec vpn PARTNER_VPN establish-tunnels immediately
set security pki ca-profile advpn ca-identity advpn
set security pki ca-profile advpn enrollment url http://10.157.92.176:8080/scep/advpn/
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces reth0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces reth1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure spoke 1:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/3 gigether-options redundant-parent reth0
user@host# set ge-0/0/4 gigether-options redundant-parent reth1
user@host# set ge-7/0/3 gigether-options redundant-parent reth0
user@host# set ge-7/0/4 gigether-options redundant-parent reth1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 25.1.1.1/24
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 21.1.1.2/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 172.16.1.2/24
```

2. Configure the routing protocol and static routes.

```
[edit protocols ospf]
user@host# set graceful-restart restart-duration 300
user@host# set graceful-restart notify-duration 300
user@host# set graceful-restart no-strict-lsa-checking
user@host# set area 0.0.0.0 interface st0.1 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.1 metric 15
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set area 0.0.0.0 interface st0.1 retransmit-interval 1
user@host# set area 0.0.0.0 interface st0.1 dead-interval 40
user@host# set area 0.0.0.0 interface st0.1 demand-circuit
user@host# set area 0.0.0.0 interface st0.1 dynamic-neighbors
user@host# set protocols ospf area 0.0.0.0 interface reth0.0

[edit routing-options]
user@host# set graceful-restart
user@host# set static route 11.1.1.0/24 next-hop 21.1.1.1
user@host# set static route 31.1.1.0/24 next-hop 21.1.1.1
user@host# set router-id 172.16.1.2
```

3. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
user@host# set proposals IKE_PROP
user@host# set certificate local-certificate Partner1_Certificate_ID

[edit security ike gateway PARTNER_GW]
user@host# set ike-policy IKE_POL
user@host# set address 11.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name container OU=Sales
user@host# set external-interface reth1
user@host# set local-address 21.1.1.2
user@host# set advpn suggester disable
user@host# set advpn partner
user@host# set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy IPSEC_POL]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals IPSEC_PROP

[edit security isec vpn PARTNER_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway PARTNER_GW
user@host# set ike ipsec-policy IPSEC_POL
user@host# set establish-tunnels immediately
```

5. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile advpn ca-identity advpn
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ca-profile advpn enrollment url
  http://10.157.92.176:8080/scep/advpn/
```

6. Configure zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces reth0.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces reth1.0
```

7. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
  user@host# show interfaces
  ge-0/0/3 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-0/0/4 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-7/0/3 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/4 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  reth0 {
    redundant-ether-options {
      redundancy-group 1;
    }
    unit 0 {
      family inet {
        address 25.1.1.1/24;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
reth1 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    family inet {
      address 21.1.1.2/24;
    }
  }
}
st0 {
  unit 1 {
    multipoint;
    family inet {
      address 172.16.1.2/24;
    }
  }
}
[edit]
user@host# show protocols
ospf {
  graceful-restart {
    restart-duration 300;
    notify-duration 300;
    no-strict-lsa-checking;
  }
  area 0.0.0.0 {
    interface st0.1 {
      interface-type p2mp;
      metric 15;
      retransmit-interval 1;
      dead-interval 40;
      demand-circuit;
      dynamic-neighbors;
    }
    interface reth0.0;
  }
}
[edit]
user@host# show routing-options
graceful-restart;
static {
  route 11.1.1.0/24 next-hop 21.1.1.1;
  route 31.1.1.0/24 next-hop 21.1.1.1;
}
router-id 172.16.1.2;
[edit]
user@host# show security ike
proposal IKE_PROP {
  authentication-method rsa-signatures;
  dh-group group5;
  authentication-algorithm sha1;
  encryption-algorithm aes-256-cbc;
```


Copyright © 2016, Juniper Networks, Inc.

```
}
policy IKE_POL {
  proposals IKE_PROP;
  certificate {
    local-certificate Partner1_Certificate_ID;
  }
}
gateway PARTNER_GW {
  ike-policy IKE_POL;
  address 11.1.1.1;
  local-identity distinguished-name;
  remote-identity distinguished-name container OU=Sales;
  external-interface reth1;
  local-address 21.1.1.2;
  advpn {
    suggester {
      disable;
    }
    partner {
    }
  }
  version v2-only;
}
[edit]
user@host# show security ipsec
proposal IPSEC_PROP {
  protocol esp;
  authentication-algorithm hmac-sha1-96;
  encryption-algorithm aes-256-cbc;
}
policy IPSEC_POL {
  perfect-forward-secrecy {
    keys group5;
  }
  proposals IPSEC_PROP;
}
vpn PARTNER_VPN {
  bind-interface st0.1;
  ike {
    gateway PARTNER_GW;
    ipsec-policy IPSEC_POL;
  }
  establish-tunnels immediately;
}
[edit]
user@host# show security pki
ca-profile advpn {
  ca-identity advpn;
  enrollment {
    url http://10.157.92.176:8080/scep/advpn/;
  }
}
[edit]
user@host# show security zones
security-zone trust {
  host-inbound-traffic {
```

Copyright © 2016, Juniper Networks, Inc.

```
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        st0.1;
        reth0.0;
      }
    }
    security-zone untrust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        reth1.0;
      }
    }
    [edit]
  user@host# show security policies
    default-policy {
      permit-all;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Partner (Spoke 2)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/2 unit 0 family inet address 31.1.1.2/24
set interfaces ge-0/0/4 unit 0 family inet address 36.1.1.1/24
set interfaces st0 unit 1 multipoint
set interfaces st0 unit 1 family inet address 172.16.1.3/24
set protocols ospf graceful-restart restart-duration 300
set protocols ospf graceful-restart notify-duration 300
set protocols ospf graceful-restart no-strict-lsa-checking
set protocols ospf area 0.0.0.0 interface st0.1 interface-type p2mp
set protocols ospf area 0.0.0.0 interface st0.1 metric 15
set protocols ospf area 0.0.0.0 interface st0.1 retransmit-interval 1
set protocols ospf area 0.0.0.0 interface st0.1 dead-interval 40
set protocols ospf area 0.0.0.0 interface st0.1 demand-circuit
set protocols ospf area 0.0.0.0 interface st0.1 dynamic-neighbors
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set routing-options graceful-restart
```


Copyright © 2016, Juniper Networks, Inc.

```
set routing-options static route 11.1.1.0/24 next-hop 31.1.1.1
set routing-options static route 21.1.1.0/24 next-hop 31.1.1.1
set routing-options router-id 172.16.1.3
set security ike proposal IKE_PROP authentication-method rsa-signatures
set security ike proposal IKE_PROP dh-group group5
set security ike proposal IKE_PROP authentication-algorithm sha1
set security ike proposal IKE_PROP encryption-algorithm aes-256-cbc
set security ike policy IKE_POL proposals IKE_PROP
set security ike policy IKE_POL certificate local-certificate Partner2_Certificate_ID
set security ike gateway PARTNER_GW ike-policy IKE_POL
set security ike gateway PARTNER_GW address 11.1.1.1
set security ike gateway PARTNER_GW dead-peer-detection
set security ike gateway PARTNER_GW local-identity distinguished-name
set security ike gateway PARTNER_GW remote-identity distinguished-name container
  OU=Sales
set security ike gateway PARTNER_GW external-interface ge-0/0/2.0
set security ike gateway PARTNER_GW local-address 31.1.1.2
set security ike gateway PARTNER_GW advpn suggester disable
set security ike gateway PARTNER_GW advpn partner
set security ike gateway PARTNER_GW version v2-only
set security ipsec proposal IPSEC_PROP protocol esp
set security ipsec proposal IPSEC_PROP authentication-algorithm hmac-sha1-96
set security ipsec proposal IPSEC_PROP encryption-algorithm aes-256-cbc
set security ipsec policy IPSEC_POL perfect-forward-secrecy keys group5
set security ipsec policy IPSEC_POL proposals IPSEC_PROP
set security ipsec vpn PARTNER_VPN bind-interface st0.1
set security ipsec vpn PARTNER_VPN ike gateway PARTNER_GW
set security ipsec vpn PARTNER_VPN ike ipsec-policy IPSEC_POL
set security ipsec vpn PARTNER_VPN establish-tunnels immediately
set security pki ca-profile advpn ca-identity advpn
set security pki ca-profile advpn enrollment url http://10.157.92.176:8080/scep/advpn/
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces ge-0/0/4.0
set security zones security-zone trust interfaces st0.1
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/2.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure spoke 2:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/2 unit 0 family inet address 31.1.1.2/24
user@host# set ge-0/0/4 unit 0 family inet address 36.1.1.1/24
user@host# set st0 unit 1 multipoint
user@host# set st0 unit 1 family inet address 172.16.1.3/24
```

2. Configure the routing protocol and static routes.

Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols ospf
user@host# set graceful-restart restart-duration 300
user@host# set graceful-restart notify-duration 300
user@host# set graceful-restart no-strict-lsa-checking
user@host# set area 0.0.0.0 interface st0.1 interface-type p2mp
user@host# set area 0.0.0.0 interface st0.1 metric 15
user@host# set area 0.0.0.0 interface st0.1 retransmit-interval 1
user@host# set area 0.0.0.0 interface st0.1 dead-interval 40
user@host# set area 0.0.0.0 interface st0.1 demand-circuit
user@host# set area 0.0.0.0 interface st0.1 dynamic-neighbors
user@host# set area 0.0.0.0 interface ge-0/0/4.0

[edit routing-options]
user@host# set graceful-restart
user@host# set static route 11.1.1.0/24 next-hop 31.1.1.1
user@host# set static route 21.1.1.0/24 next-hop 31.1.1.1
user@host# set router-id 172.16.1.3
```

3. Configure Phase 1 options.

```
[edit security ike proposal IKE_PROP]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy IKE_POL]
user@host# set proposals IKE_PROP
user@host# set certificate local-certificate Partner2_Certificate_ID

[edit security ike gateway PARTNER_GW]
user@host# set ike-policy IKE_POL
user@host# set address 11.1.1.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name container OU=Sales
user@host# set external-interface ge-0/0/2.0
user@host# set local-address 31.1.1.2
user@host# set advpn suggester disable
user@host# set advpn partner
user@host# set version v2-only
```

4. Configure Phase 2 options.

```
[edit security ipsec proposal IPSEC_PROP]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy IPSEC_POL]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals IPSEC_PROP

[edit security isec vpn PARTNER_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway PARTNER_GW
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ike ipsec-policy IPSEC_POL
user@host# set establish-tunnels immediately
```

5. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile advpn ca-identity advpn
user@host# set ca-profile advpn enrollment url
  http://10.157.92.176:8080/scep/advpn/
```

6. Configure zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/4.0
user@host# set interfaces st0.1

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces ge-0/0/2.0
```

7. Configure the default security policy.

```
[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, **show security ike**, **show security ipsec**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
  user@host# show interfaces
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 31.1.1.2/24;
      }
    }
  }
  ge-0/0/4{
    unit 0 {
      family inet {
        address 36.1.1.1/24;
      }
    }
  }
  st0 {
    unit 1 {
      multipoint;
      family inet {
        address 172.16.1.3/24;
      }
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
}
[edit]
user@host# show protocols
ospf {
  graceful-restart {
    restart-duration 300;
    notify-duration 300;
    no-strict-lsa-checking;
  }
  area 0.0.0.0 {
    interface st0.1 {
      interface-type p2mp;
      metric 15;
      retransmit-interval 1;
      dead-interval 40;
      demand-circuit;
      dynamic-neighbors;
    }
    interface ge-0/0/4.0;
  }
}
[edit]
user@host# show routing-options
graceful-restart;
static {
  route 11.1.1.0/24 next-hop 31.1.1.1;
  route 21.1.1.0/24 next-hop 31.1.1.1;
}
router-id 172.16.1.3;
[edit]
user@host# show security ike
proposal IKE_PROP {
  authentication-method rsa-signatures;
  dh-group group5;
  authentication-algorithm sha1;
  encryption-algorithm aes-256-cbc;
}
policy IKE_POL {
  proposals IKE_PROP;
  certificate {
    local-certificate Partner2_Certificate_ID
  }
}
gateway PARTNER_GW {
  ike-policy IKE_POL;
  address 11.1.1.1;
  local-identity distinguished-name;
  remote-identity distinguished-name container OU=Sales;
  external-interface ge-0/0/2.0;
  local-address 31.1.1.2;
  advpn {
    suggester{
      disable;
    }
    partner {
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    version v2-only;
}
[edit]
user@host# show security ipsec
proposal IPSEC_PROP {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
}
policy IPSEC_POL {
    perfect-forward-secrecy {
        keys group5;
    }
    proposals IPSEC_PROP;
}
vpn PARTNER_VPN {
    bind-interface st0.1;
    ike {
        gateway PARTNER_GW;
        ipsec-policy IPSEC_POL;
    }
    establish-tunnels immediately;
}
[edit]
user@host# show security pki
ca-profile advpn {
    ca-identity advpn;
    enrollment {
        url http://10.157.92.176:8080/scep/advpn/;
    }
}
[edit]
user@host# show security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        ge-0/0/4.0;
        st0.1;
    }
}
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
```

```
    }
    interfaces {
      ge-0/0/2.0;
    }
  }
  [edit]
user@host# show security policies
  default-policy {
    permit-all;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly. First, verify that tunnels are established between the AutoVPN hub and spokes. When traffic is passed from one spoke to another through the hub, a shortcut can be established between the spokes. Verify that the shortcut partners have established a tunnel between them and that a route to the peer is installed on the partners.

- Verifying Tunnels Between the Hub and Spokes on page 6588
- Verifying the Shortcut Tunnel Between Partners on page 6595

Verifying Tunnels Between the Hub and Spokes

Purpose Verify that tunnels are established between the AutoVPN hub and spokes. Initial traffic from one spoke to another must travel through the hub.

Action From operational mode, enter the **show security ike security-associations** and **show security ipsec security-associations** commands on the hub and spokes.

The following commands are entered on the hub:

```
user@host> show security ike security-associations
node1:
--------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

10957048 UP    2d58d8fbc396762d  46145be580c68be0  IKEv2          31.1.1.2

10957049 UP    fa05ee6d0f2cfb22  16f5ca836b118c0e  IKEv2          21.1.1.2


user@host> show security ike security-associations detail
node1:
--------------------------------------------------------------------------
IKE peer 31.1.1.2, Index 10957048, Gateway Name: SUGGESTER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    0
     Suggestions accepted:    0
     Suggestions declined:    0
  Role: Responder, State: UP
  Initiator cookie: 2d58d8fbc396762d, Responder cookie: 46145be580c68be0
```

 Copyright © 2016, Juniper Networks, Inc.

```
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 28196 seconds
  Peer ike-id: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                2030
   Output bytes  :                2023
   Input  packets:                   4
   Output packets:                   4
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 31.1.1.2:500
    Local identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Remote identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Flags: IKE SA is created

IKE peer 21.1.1.2, Index 10957049, Gateway Name: SUGGESTER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    0
     Suggestions accepted:    0
     Suggestions declined:    0
  Role: Responder, State: UP
  Initiator cookie: fa05ee6d0f2cfb22, Responder cookie: 16f5ca836b118c0e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 21.1.1.2:500
  Lifetime: Expires in 28219 seconds
Peer ike-id: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                2030
   Output bytes  :                2023
   Input  packets:                   4
   Output packets:                   4
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 21.1.1.2:500
    Local identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Remote identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
```

```
    Flags: IKE SA is created

user@host> show security ipsec security-associations
node1:
--------------------------------------------------------------------------
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <201326593 ESP:aes-cbc-256/sha1 44ccf265 2999/ unlim - root 500 31.1.1.2

  >201326593 ESP:aes-cbc-256/sha1 a9d301b0 2999/ unlim - root 500 31.1.1.2

  <201326594 ESP:aes-cbc-256/sha1 98a2b155 3022/ unlim - root 500 21.1.1.2

  >201326594 ESP:aes-cbc-256/sha1 de912bcd 3022/ unlim - root 500 21.1.1.2

user@host> show security ipsec security-associations detail
node1:
--------------------------------------------------------------------------

ID: 201326593 Virtual-system: root, VPN Name: SUGGESTER_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
    Tue Jan 13 2015 12:57:48 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:57:48 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 12:57:48 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: 44ccf265, AUX-SPI: 0
    Hard lifetime: Expires in 2991 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2414 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: a9d301b0, AUX-SPI: 0
    Hard lifetime: Expires in 2991 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2414 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

ID: 201326594 Virtual-system: root, VPN Name: SUGGESTER_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 21.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 3, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
    Tue Jan 13 2015 12:58:11 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:58:11 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
```


Copyright © 2016, Juniper Networks, Inc.

```
    Tue Jan 13 2015 12:58:11 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: 98a2b155, AUX-SPI: 0
    Hard lifetime: Expires in 3014 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2436 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: de912bcd, AUX-SPI: 0
    Hard lifetime: Expires in 3014 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2436 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

```
user@host> show route protocol ospf
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

25.1.1.0/24        *[OSPF/10] 00:00:27, metric 11
                    > to 172.16.1.2 via st0.1
36.1.1.0/24        *[OSPF/10] 00:00:27, metric 11
                    > to 172.16.1.3 via st0.1
172.16.1.2/32      *[OSPF/10] 00:00:27, metric 10
                    > to 172.16.1.2 via st0.1
172.16.1.3/32      *[OSPF/10] 00:00:27, metric 10
                    > to 172.16.1.3 via st0.1
224.0.0.5/32       *[OSPF/10] 00:00:48, metric 1
                       MultiRecv
```

```
user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.3       st0.1                  Full      172.16.1.3       128     -
172.16.1.2       st0.1                  Full      172.16.1.2       128     -
```

The following commands are entered on spoke 1:

```
user@host> show security ike security-associations
node0:
--------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

578872  UP     fa05ee6d0f2cfb22  16f5ca836b118c0e  IKEv2          11.1.1.1
```

```
user@host> show security ike security-associations detail
node0:
--------------------------------------------------------------------------
IKE peer 11.1.1.1, Index 578872, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    0
     Suggestions accepted:    0
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: fa05ee6d0f2cfb22, Responder cookie: 16f5ca836b118c0e
  Exchange type: IKEv2, Authentication method: RSA-signatures
```

Copyright © 2016, Juniper Networks, Inc.

```
  Local: 21.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 28183 seconds
  Peer ike-id: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 2023
   Output bytes  :                 2030
   Input  packets:                    4
   Output packets:                    4
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 21.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Flags: IKE SA is created

user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <67108866 ESP:aes-cbc-256/sha1 de912bcd 2985/ unlim - root 500 11.1.1.1

  >67108866 ESP:aes-cbc-256/sha1 98a2b155 2985/ unlim - root 500 11.1.1.1


user@host> show security ipsec security-associations detail
node0:
--------------------------------------------------------------------------

ID: 67108866 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
    Tue Jan 13 2015 12:58:11 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:58:11 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 12:58:11 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: de912bcd, AUX-SPI: 0
    Hard lifetime: Expires in 2980 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2358 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```


Copyright © 2016, Juniper Networks, Inc.

```
  Direction: outbound, SPI: 98a2b155, AUX-SPI: 0
    Hard lifetime: Expires in 2980 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2358 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

user@host> show route protocol ospf
inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.1.1.0/24        *[OSPF/10] 00:11:46, metric 16
                    > to 172.16.1.1 via st0.1
36.1.1.0/24        *[OSPF/10] 00:11:46, metric 26
                    > to 172.16.1.1 via st0.1
172.16.1.1/32      *[OSPF/10] 00:11:46, metric 15
                    > to 172.16.1.1 via st0.1
172.16.1.3/32      *[OSPF/10] 00:11:46, metric 25
                    > to 172.16.1.1 via st0.1
224.0.0.5/32       *[OSPF/10] 00:16:52, metric 1
                       MultiRecv

user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.1       st0.1                  Full      172.16.1.1       128     -
```

The following commands are entered on spoke 2:

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

2299162 UP     2d58d8fbc396762d  46145be580c68be0  IKEv2          11.1.1.1


user@host> show security ike security-associations detail
IKE peer 11.1.1.1, Index 2299162, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    0
     Suggestions accepted:    0
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 2d58d8fbc396762d, Responder cookie: 46145be580c68be0
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 28135 seconds
  Peer ike-id: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 2023
   Output bytes  :                 2030
   Input  packets:                    4
```

Copyright © 2016, Juniper Networks, Inc.

```
   Output packets:                    4
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 31.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Flags: IKE SA is created

user@host> show security ipsec security-associations
  Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <67108866 ESP:aes-cbc-256/sha1 a9d301b0 2936/ unlim - root 500 11.1.1.1

  >67108866 ESP:aes-cbc-256/sha1 44ccf265 2936/ unlim - root 500 11.1.1.1


user@host> show security ipsec security-associations detail
ID: 67108866 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
    Tue Jan 13 2015 12:57:48 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:57:48 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 12:57:48 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: a9d301b0, AUX-SPI: 0
    Hard lifetime: Expires in 2933 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2311 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 44ccf265, AUX-SPI: 0
    Hard lifetime: Expires in 2933 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2311 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

user@host> show route protocol ospf
inet.0: 36 destinations, 36 routes (35 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.1.1.0/24        *[OSPF/10] 00:00:09, metric 16
                    > to 172.16.1.1 via st0.1
25.1.1.0/24        *[OSPF/10] 00:00:09, metric 26
                    > to 172.16.1.1 via st0.1
172.16.1.1/32      *[OSPF/10] 00:00:09, metric 15
                    > to 172.16.1.1 via st0.1
```


Copyright © 2016, Juniper Networks, Inc.

```
172.16.1.2/32      *[OSPF/10] 00:00:09, metric 25
                    > to 172.16.1.1 via st0.1
224.0.0.5/32       *[OSPF/10] 00:17:52, metric 1
                      MultiRecv
```

```
user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.1       st0.1                  Full      172.16.1.1       128     -
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. The hub shows two active tunnels, one to each spoke. Each spoke shows an active tunnel to the hub.

If no SAs are listed for IKE Phase 1, then there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spokes.

If no SAs are listed for IKE Phase 2, then there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spokes.

The **show route protocol ospf** command displays entries in the routing table that were learned from the OSPF protocol. The **show ospf neighbor** command displays information about OSPF neighbors.

Verifying the Shortcut Tunnel Between Partners

Purpose The AutoVPN hub can act as a shortcut suggester when it notices that traffic is exiting a tunnel with one of its spokes and entering a tunnel with another spoke. A new IPsec SA, or shortcut, is established between the two shortcut partners. On each partner, the route to the network behind its partner now points to the shortcut tunnel instead of to the tunnel between the partner and the suggester (hub).

Action From operational mode, enter the **show security ike security-associations**, **show security ipsec security-associations**, **show route protocol ospf**, and **show ospf neighbor** commands on the spokes.

The following commands are entered on the hub:

```
user@host> show security ike security-associations
node0:
--------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

10957048 UP    2d58d8fbc396762d  46145be580c68be0  IKEv2          31.1.1.2

10957049 UP    fa05ee6d0f2cfb22  16f5ca836b118c0e  IKEv2          21.1.1.2
```

```
user@host> show security ike security-associations detail
node0:
--------------------------------------------------------------------------
IKE peer 31.1.1.2, Index 10957048, Gateway Name: SUGGESTER_GW
```

```
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    1
     Suggestions accepted:    1
     Suggestions declined:    0
  Role: Responder, State: UP
  Initiator cookie: 2d58d8fbc396762d, Responder cookie: 46145be580c68be0
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 27781 seconds
  Peer ike-id: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 260
   Output bytes  :                 548
   Input  packets:                   3
   Output packets:                   3
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 31.1.1.2:500
    Local identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Remote identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Flags: IKE SA is created

IKE peer 21.1.1.2, Index 10957049, Gateway Name: SUGGESTER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    1
     Suggestions accepted:    1
     Suggestions declined:    0
  Role: Responder, State: UP
  Initiator cookie: fa05ee6d0f2cfb22, Responder cookie: 16f5ca836b118c0e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 21.1.1.2:500
  Lifetime: Expires in 27804 seconds
  Peer ike-id: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 244
   Output bytes  :                 548
   Input  packets:                   3
   Output packets:                   3
  IPSec security associations: 0 created, 0 deleted
```


Copyright © 2016, Juniper Networks, Inc.

```
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 21.1.1.2:500
    Local identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Remote identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Flags: IKE SA is created

user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
s  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <201326593 ESP:aes-cbc-256/sha1 44ccf265 2584/ unlim - root 500 31.1.1.2

  >201326593 ESP:aes-cbc-256/sha1 a9d301b0 2584/ unlim - root 500 31.1.1.2

  <201326594 ESP:aes-cbc-256/sha1 98a2b155 2607/ unlim - root 500 21.1.1.2

  >201326594 ESP:aes-cbc-256/sha1 de912bcd 2607/ unlim - root 500 21.1.1.2


user@host> show security ipsec security-associations detail
node0:
--------------------------------------------------------------------------

ID: 201326593 Virtual-system: root, VPN Name: SUGGESTER_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
    Tue Jan 13 2015 13:09:48 -0800: Bind-interface's address received. Information
 updated (1 times)
    Tue Jan 13 2015 13:09:48 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
  Direction: inbound, SPI: 44ccf265, AUX-SPI: 0
    Hard lifetime: Expires in 2578 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2001 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: a9d301b0, AUX-SPI: 0
    Hard lifetime: Expires in 2578 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2001 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

ID: 201326594 Virtual-system: root, VPN Name: SUGGESTER_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 21.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
```

```
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
  Tunnel events:
    Tue Jan 13 2015 13:09:48 -0800: Bind-interface's address received. Information
 updated (1 times)
    Tue Jan 13 2015 13:09:48 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
  Direction: inbound, SPI: 98a2b155, AUX-SPI: 0
    Hard lifetime: Expires in 2601 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2023 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: de912bcd, AUX-SPI: 0
    Hard lifetime: Expires in 2601 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2023 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

user@host> show route protocol ospf
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

25.1.1.0/24        *[OSPF/10] 00:04:49, metric 11
                    > to 172.16.1.2 via st0.1
36.1.1.0/24        *[OSPF/10] 00:04:49, metric 11
                    > to 172.16.1.3 via st0.1
172.16.1.2/32      *[OSPF/10] 00:04:49, metric 10
                    > to 172.16.1.2 via st0.1
172.16.1.3/32      *[OSPF/10] 00:04:49, metric 10
                    > to 172.16.1.3 via st0.1
224.0.0.5/32       *[OSPF/10] 00:05:10, metric 1
                       MultiRecv

user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.3       st0.1                  Full      172.16.1.3       128     -
172.16.1.2       st0.1                  Full      172.16.1.2       128     -
```

The following commands are entered on spoke 1:

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

578872  UP     fa05ee6d0f2cfb22  16f5ca836b118c0e  IKEv2          11.1.1.1

578873  UP     895e4d9c7c5da7a4  17de7f18b45139b4  IKEv2          31.1.1.2


user@host> show security ike security-associations detail
node0:
--------------------------------------------------------------------------
IKE peer 11.1.1.1, Index 578872, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    1
     Suggestions accepted:    1
```


Copyright © 2016, Juniper Networks, Inc.

```
    Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: fa05ee6d0f2cfb22, Responder cookie: 16f5ca836b118c0e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 21.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 27906 seconds
  Peer ike-id: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                2495
   Output bytes  :                2274
   Input  packets:                   6
   Output packets:                   7
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 21.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Flags: IKE SA is created

IKE peer 31.1.1.2, Index 578873, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Initiator, State: UP
  Initiator cookie: 895e4d9c7c5da7a4, Responder cookie: 17de7f18b45139b4
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 21.1.1.2:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 28787 seconds
  Peer ike-id: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                1855
   Output bytes  :                1990
   Input  packets:                   2
   Output packets:                   2
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 21.1.1.2:500, Remote: 31.1.1.2:500
    Local identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Flags: IKE SA is created
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <67108866 ESP:aes-cbc-256/sha1 de912bcd 2709/ unlim - root 500 11.1.1.1

  >67108866 ESP:aes-cbc-256/sha1 98a2b155 2709/ unlim - root 500 11.1.1.1

  <67108868 ESP:aes-cbc-256/sha1 75d0177b 3590/ unlim - root 500 31.1.1.2

  >67108868 ESP:aes-cbc-256/sha1 e4919d73 3590/ unlim - root 500 31.1.1.2


user@host> show security ipsec security-associations detail
node0:
--------------------------------------------------------------------------

ID: 67108866 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
    Tue Jan 13 2015 12:58:11 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:58:11 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 12:58:11 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: de912bcd, AUX-SPI: 0
    Hard lifetime: Expires in 2701 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2079 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 98a2b155, AUX-SPI: 0
    Hard lifetime: Expires in 2701 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2079 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

ID: 67108868 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Jan 13 2015 13:12:52 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 13:12:52 -0800: Tunnel is ready. Waiting for trigger event
```


Copyright © 2016, Juniper Networks, Inc.

```
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 13:12:52 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: 75d0177b, AUX-SPI: 0
    Hard lifetime: Expires in 3582 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2959 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: e4919d73, AUX-SPI: 0
    Hard lifetime: Expires in 3582 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2959 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

user@host> show route protocol ospf
inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.1.1.0/24        *[OSPF/10] 00:03:29, metric 16
                    > to 172.16.1.1 via st0.1
36.1.1.0/24        *[OSPF/10] 00:00:35, metric 16
                    > to 172.16.1.3 via st0.1
172.16.1.1/32      *[OSPF/10] 00:03:29, metric 15
                    > to 172.16.1.1 via st0.1
172.16.1.3/32      *[OSPF/10] 00:00:35, metric 15
                    > to 172.16.1.3 via st0.1
224.0.0.5/32       *[OSPF/10] 00:20:22, metric 1
                       MultiRecv

user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.3       st0.1                  Full      172.16.1.3       128     -
172.16.1.1       st0.1                  Full      172.16.1.1       128
```

The following commands are entered on spoke 2:

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

2299162 UP     2d58d8fbc396762d  46145be580c68be0  IKEv2          11.1.1.1

2299163 UP     895e4d9c7c5da7a4  17de7f18b45139b4  IKEv2          21.1.1.2


user@host> show security ike security-associations detail
IKE peer 11.1.1.1, Index 2299162, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    1
     Suggestions accepted:    1
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 2d58d8fbc396762d, Responder cookie: 46145be580c68be0
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 11.1.1.1:500
```

Copyright © 2016, Juniper Networks, Inc.

```
  Lifetime: Expires in 27835 seconds
  Peer ike-id: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                2571
   Output bytes  :                2290
   Input  packets:                   7
   Output packets:                   7
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 31.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=suggester, OU=Sales, O=XYZ, L=Sunnyvale, ST=CA,
C=US
    Flags: IKE SA is created

IKE peer 21.1.1.2, Index 2299163, Gateway Name: PARTNER_GW
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: 895e4d9c7c5da7a4, Responder cookie: 17de7f18b45139b4
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 21.1.1.2:500
  Lifetime: Expires in 28739 seconds
  Peer ike-id: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                2066
   Output bytes  :                1931
   Input  packets:                   3
   Output packets:                   3
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 31.1.1.2:500, Remote: 21.1.1.2:500
    Local identity: DC=XYZ, CN=partner2, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Remote identity: DC=XYZ, CN=partner1, OU=Sales, O=XYZ, L=NewYork, ST=NY, C=US

    Flags: IKE SA is created

user@host> show security ipsec security-associations
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <67108866 ESP:aes-cbc-256/sha1 a9d301b0 2638/ unlim - root 500 11.1.1.1
```


Copyright © 2016, Juniper Networks, Inc.

```
  >67108866 ESP:aes-cbc-256/sha1 44ccf265 2638/ unlim - root 500 11.1.1.1

  <67108868 ESP:aes-cbc-256/sha1 e4919d73 3542/ unlim - root 500 21.1.1.2

  >67108868 ESP:aes-cbc-256/sha1 75d0177b 3542/ unlim - root 500 21.1.1.2
```

```
user@host> show security ipsec security-associations detail
ID: 67108866 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
    Tue Jan 13 2015 12:57:48 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 12:57:48 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 12:57:48 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: a9d301b0, AUX-SPI: 0
    Hard lifetime: Expires in 2632 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2010 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 44ccf265, AUX-SPI: 0
    Hard lifetime: Expires in 2632 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2010 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

ID: 67108868 Virtual-system: root, VPN Name: PARTNER_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 21.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Responder
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608aa9
  Tunnel events:
    Tue Jan 13 2015 13:12:52 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Jan 13 2015 13:12:52 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Jan 13 2015 13:12:52 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: e4919d73, AUX-SPI: 0
    Hard lifetime: Expires in 3536 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2958 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Copyright © 2016, Juniper Networks, Inc.

```
  Direction: outbound, SPI: 75d0177b, AUX-SPI: 0
    Hard lifetime: Expires in 3536 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2958 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (256 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64

user@host> show route protocol ospf
inet.0: 36 destinations, 36 routes (35 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.1.1.0/24        *[OSPF/10] 00:03:55, metric 16
                    > to 172.16.1.1 via st0.1
25.1.1.0/24        *[OSPF/10] 00:01:02, metric 16
                    > to 172.16.1.2 via st0.1
172.16.1.1/32      *[OSPF/10] 00:03:55, metric 15
                    > to 172.16.1.1 via st0.1
172.16.1.2/32      *[OSPF/10] 00:01:02, metric 15
                    > to 172.16.1.2 via st0.1
224.0.0.5/32       *[OSPF/10] 00:21:38, metric 1
                       MultiRecv

user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.16.1.2       st0.1                  Full      172.16.1.2       128    -
172.16.1.1       st0.1                  Full      172.16.1.1       128    -
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. The hub still shows two active tunnels, one to each spoke. Each spoke shows two active tunnels, one to the hub and one to its shortcut partner.

The **show route protocol ospf** command shows the addition of routes to the partner and to the hub.

Related Documentation

- Understanding Auto Discovery VPN on page 6559
- Understanding Traffic Routing with Shortcut Tunnels on page 6564
- Enabling OSPF to Update Routes Quickly After ADVPN Shortcut Tunnels Are Established on page 6604

Enabling OSPF to Update Routes Quickly After ADVPN Shortcut Tunnels Are Established

Problem **Description:** OSPF can take up to 9 seconds to update a shortcut route in the routing table. It can take up to 10 seconds before traffic is forwarded to the shortcut tunnel.

Symptoms: When a shortcut tunnel is established between two shortcut partners, OSPF initiates an OSPF hello packet. Because of the timing of the shortcut tunnel establishment and the OSPF neighbor installation, the first packet in the tunnel might be dropped. This can cause OSPF to try again to establish an OSPF adjacency.

Copyright © 2016, Juniper Networks, Inc.

By default, the interval at which the OSPF retries to establish an adjacency is 10 seconds. After a shortcut tunnel is established, it can take more than 10 seconds for OSPF to establish an adjacency between the partners.

Solution Configuring a smaller retry interval, such as 1 or 2 seconds, can enable OSPF to establish adjacencies faster over the shortcut tunnel. For example, use the following configurations:

```
[edit]
set protocols ospf area 0.0.0.0 interface st0.1 retransmit-interval 1
set protocols ospf area 0.0.0.0 interface st0.1 dead-interval 40
```

Related Documentation

- Understanding Auto Discovery VPN on page 6559
- Understanding Traffic Routing with Shortcut Tunnels on page 6564
- Example: Improving Network Resource Utilization with Auto Discovery VPN Dynamic Tunnels on page 6566

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 274

Configuring AutoVPN and Traffic Selectors

- Understanding AutoVPN with Traffic Selectors on page 6607
- Example: Forwarding Traffic Through an AutoVPN Tunnel with Traffic Selectors on page 6608
- Example: Ensuring VPN Tunnel Availability with AutoVPN and Traffic Selectors on page 6624

Understanding AutoVPN with Traffic Selectors

AutoVPN hubs can be configured with multiple traffic selectors to protect traffic to spokes. This feature provides the following benefits:

- A single VPN configuration can support many different peers.
- VPN peers can be non-SRX Series devices.
- A single peer can establish multiple tunnels with the same VPN.
- A larger number of tunnels can be supported than with AutoVPN with dynamic routing protocols.

When the hub-to-spoke tunnel is established, the hub uses auto route insertion (ARI), known in previous releases as *reverse route insertion (RRI)*, to insert the route to the spoke prefix in its routing table. The ARI route can then be imported to routing protocols and distributed to the core network.

AutoVPN with traffic selectors can be configured with the secure tunnel (st0) interface in point-to-point mode for both IKEv1 and IKEv2.

NOTE: **Dynamic routing protocols are not supported on st0 interfaces when traffic selectors are configured.**

Copyright © 2016, Juniper Networks, Inc.

Note the following caveats when configuring AutoVPN with traffic selectors:

- Dynamic routing protocols are not supported with traffic selectors with st0 interfaces in point-to-point mode.
- IPv6 addresses cannot be configured for traffic selectors on AutoVPN hubs. Only IPv4-in-IPv4 tunnel encapsulation is supported for traffic selectors on AutoVPN hubs; IPv4-in-IPv6, IPv6-in-IPv4, and IPv6-in-IPv6 tunnels are not supported.
- Auto Discovery VPN and IKEv2 configuration payload cannot be configured with AutoVPN with traffic selectors.
- Spokes can be non-SRX Series devices; however, note the following differences:
  - In IKEv2, a non-SRX Series spoke can propose multiple traffic selectors in a single SA negotiation. This is not supported on SRX Series devices and the negotiation is rejected.
  - A non-SRX Series spoke can identify specific ports or protocols for traffic selector use. Ports and protocols are not supported with traffic selectors on SRX Series devices and the negotiation is rejected.

Related Documentation

- Understanding Traffic Selectors in Route-Based VPNs on page 6195
- Understanding Auto Route Insertion on page 6213
- Example: Ensuring VPN Tunnel Availability with AutoVPN and Traffic Selectors on page 6624
- Example: Configuring Traffic Selectors in a Route-Based VPN on page 6198

Example: Forwarding Traffic Through an AutoVPN Tunnel with Traffic Selectors

This example shows how to configure traffic selectors, instead of dynamic routing protocols, to forward packets through a VPN tunnel in an AutoVPN deployment. When traffic selectors are configured, the secure tunnel (st0) interface must be in point-to-point mode. Traffic selectors are configured on both the hub and spoke devices.

- Requirements on page 6608
- Overview on page 6609
- Configuration on page 6611
- Verification on page 6621

Requirements

This example uses the following hardware and software components:

- Two SRX Series devices connected and configured in a chassis cluster. The chassis cluster is the AutoVPN hub.
- An SRX Series device configured as an AutoVPN spoke.

Copyright © 2016, Juniper Networks, Inc.

- Junos OS Release 12.3X48-D10 or later.
- Digital certificates enrolled in the hub and the spoke devices that allow the devices to authenticate each other.

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates. See "Understanding Local Certificate Requests" on page 6371.
- Enroll the digital certificates in each device. See "Understanding Certificate Loading" on page 6377.

Overview

In this example, traffic selectors are configured on the AutoVPN hub and spoke. Only traffic that conforms to the configured traffic selector is forwarded through the tunnel. On the hub, the traffic selector is configured with the local IP address 192.0.0.0/8 and the remote IP address 172.0.0.0/8. On the spoke, the traffic selector is configured with the local IP address 172.0.0.0/8 and the remote IP address 192.0.0.0/8.

NOTE: **The traffic selector IP addresses configured on the spoke can be a subset of the traffic selector IP addresses configured on the hub. This is known as traffic selector flexible match.**

Certain Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hubs and spokes must have the same values. Table 567 shows the values used in this example:

Table 567: Phase 1 and Phase 2 Options for AutoVPN Hubs and Spokes with Traffic Selectors

Option	Value
IKE proposal:	
Authentication method	rsa-signatures
Diffie-Hellman (DH) group	group5
Authentication algorithm	sha-1
Encryption algorithm	aes-256-cbc
IKE policy:	
Mode	main
Certificate	local-certificate
IKE gateway:	

Copyright © 2016, Juniper Networks, Inc.

Table 567: Phase 1 and Phase 2 Options for AutoVPN Hubs and Spokes with Traffic Selectors *(continued)*

Option	Value
Dynamic	distinguished name wildcard DC=Common_component
IKE user type	group IKE id
Local identity	distinguished name
Version	v1-only
IPsec proposal:	
Protocol	esp
Authentication algorithm	hmac-sha1-96
Encryption algorithm	aes-192-cbc
Lifetime	3600 seconds 150,000 kilobytes
IPsec policy:	
Perfect Forward Secrecy (PFS) group	group5

Topology

Figure 294 shows the SRX Series devices to be configured for this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 294: AutoVPN with Traffic Selectors

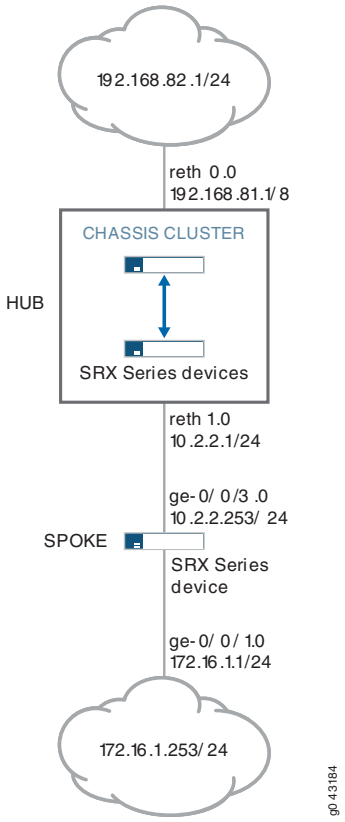

Configuration

- Configuring the Hub on page 6611
- Configuring the Spoke on page 6616

Configuring the Hub

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/2 gigether-options redundant-parent reth1
set interfaces ge-0/0/3 gigether-options redundant-parent reth0
set interfaces ge-8/0/2 gigether-options redundant-parent reth1
set interfaces ge-8/0/3 gigether-options redundant-parent reth0
set interfaces lo0 unit 0 family inet address 10.100.1.100/24
set interfaces lo0 redundant-pseudo-interface-options redundancy-group 1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 192.168.81.1/8
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 10.2.2.1/24
set interfaces st0 unit 1 family inet
```

Copyright © 2016, Juniper Networks, Inc.

```
set security ike proposal prop_ike authentication-method rsa-signatures
set security ike proposal prop_ike dh-group group5
set security ike proposal prop_ike authentication-algorithm sha1
set security ike proposal prop_ike encryption-algorithm aes-256-cbc
set security ike policy ikepol1 mode main
set security ike policy ikepol1 proposals prop_ike
set security ike policy ikepol1 certificate local-certificate Hub_ID
set security ike gateway HUB_GW ike-policy ikepol1
set security ike gateway HUB_GW dynamic distinguished-name wildcard
  DC=Domain_component
set security ike gateway HUB_GW dynamic ike-user-type group-ike-id
set security ike gateway HUB_GW local-identity distinguished-name
set security ike gateway HUB_GW external-interface reth1
set security ike gateway HUB_GW version v1-only
set security ipsec proposal prop_ipsec protocol esp
set security ipsec proposal prop_ipsec authentication-algorithm hmac-sha1-96
set security ipsec proposal prop_ipsec encryption-algorithm aes-192-cbc
set security ipsec proposal prop_ipsec lifetime-seconds 3600
set security ipsec proposal prop_ipsec lifetime-kilobytes 150000
set security ipsec policy ipsecpol1 perfect-forward-secrecy keys group5
set security ipsec policy ipsecpol1 proposals prop_ipsec
set security ipsec vpn HUB_VPN bind-interface st0.1
set security ipsec vpn HUB_VPN ike gateway HUB_GW
set security ipsec vpn HUB_VPN ike ipsec-policy ipsecpol1
set security ipsec vpn HUB_VPN traffic-selector ts1 local-ip 192.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts1 remote-ip 172.0.0.0/8
set security pki ca-profile rsa ca-identity rsa
set security pki ca-profile rsa revocation-check disable
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces reth0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces reth1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the hub:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/2 gigether-options redundant-parent reth1
user@host# set ge-0/0/3 gigether-options redundant-parent reth0
user@host# set ge-8/0/2 gigether-options redundant-parent reth1
user@host# set ge-8/0/3 gigether-options redundant-parent reth0
user@host# set lo0 unit 0 family inet address 10.100.1.100/24
user@host# set lo0 redundant-pseudo-interface-options redundancy-group 1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 192.168.81.1/8
user@host# set reth1 redundant-ether-options redundancy-group 1
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set reth1 unit 0 family inet address 10.2.2.1/24
user@host# set st0 unit 1 family inet
```

2. Configure Phase 1 options.

```
[edit security ike proposal prop_ike]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy ikepol1]
user@host# set mode main
user@host# set proposals prop_ike
user@host# set certificate local-certificate Hub_ID

[edit security ike gateway HUB_GW]
user@host# set ike-policy ikepol1
user@host# set dynamic distinguished-name wildcard DC=Domain_component
user@host# set dynamic ike-user-type group-ike-id
user@host# set local-identity distinguished-name
user@host# set external-interface reth1
user@host# set version v1-only
```

3. Configure Phase 2 options.

```
[edit security ipsec proposal prop_ipsec]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-192-cbc
user@host# set lifetime-seconds 3600
user@host# set lifetime-kilobytes 150000

[edit security ipsec policy ipsecpol1]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals prop_ipsec

[edit security ipsec HUB_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway HUB_GW
user@host# set ike ipsec-policy ipsecpol1
user@host# set traffic-selector ts1 local-ip 192.0.0.0/8
user@host# set traffic-selector ts1 remote-ip 172.0.0.0/8
```

4. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile rsa ca-identity rsa
user@host# set ca-profile rsa revocation-check disable
```

5. Configure security zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces reth0.0
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces lo0.0
user@host# set interfaces reth1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/2 {
    gigether-options {
        redundant-parent reth1;
    }
}
ge-0/0/3 {
    gigether-options {
        redundant-parent reth0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.100.1.100/24;
        }
    }
    redundant-pseudo-interface-options {
        redundancy-group 1;
    }
}
reth0 {
    redundant-ether-options {
        redundancy-group 1;
    }
    unit 0 {
        family inet {
            address 192.168.81.1/8;
        }
    }
}
reth1 {
    redundant-ether-options {
        redundancy-group 1;
    }
    unit 0 {
        family inet {
            address 10.2.2.1/24;
        }
```

 Copyright © 2016, Juniper Networks, Inc.

```
    }
}
st0 {
    unit 1 {
        family inet;
    }
}
[edit]
user@host# show security ike
proposal prop_ike {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
policy ikepol1 {
    mode main;
    proposals prop_ike;
    certificate {
        local-certificate Hub_ID;
    }
}
gateway HUB_GW {
    ike-policy ikepol1;
    dynamic distinguished-name wildcard DC=Domain_component;
    dynamic ike-user-type group-ike-id;
    local-identity distinguished-name;
    external-interface reth1;
    version v1-only;
}
[edit]
user@host# show security ipsec
proposal prop_ipsec {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-192-cbc;
    lifetime-seconds 3600;
    lifetime-kilobytes 150000;
}
policy ipsecpol1 {
    perfect-forward-secrecy {
        keys group5;
    }
    proposals prop_ipsec;
}
vpn HUB_VPN {
    bind-interface st0.1;
    ike {
        gateway HUB_GW;
        ipsec-policy ipsecpol1;
    }
    traffic-selector ts1 {
        local-ip 192.0.0.0/8;
        remote-ip 172.0.0.0/8;
    }
}
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security pki
ca-profile rsa {
    ca-identity rsa;
    revocation-check {
        disable;
    }
}
[edit]
user@host# show security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        st0.1;
        reth0.0;
    }
}
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        lo0.0;
        reth1.0;
    }
}
[edit]
user@host# show security policies
default-policy {
    permit-all;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Spoke

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 172.16.1.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.2.2.253/24
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces st0 unit 1 family inet
set security ike proposal prop_ike authentication-method rsa-signatures
set security ike proposal prop_ike dh-group group5
set security ike proposal prop_ike authentication-algorithm sha1
set security ike proposal prop_ike encryption-algorithm aes-256-cbc
set security ike policy ikepol1 mode main
set security ike policy ikepol1 proposals prop_ike
set security ike policy ikepol1 certificate local-certificate Spoke1_ID
set security ike gateway SPOKE_GW ike-policy ikepol1
set security ike gateway SPOKE_GW address 10.2.2.1
set security ike gateway SPOKE_GW local-identity distinguished-name
set security ike gateway SPOKE_GW remote-identity distinguished-name container
  DC=Domain_component
set security ike gateway SPOKE_GW external-interface ge-0/0/3.0
set security ike gateway SPOKE_GW version v1-only
set security ipsec proposal prop_ipsec protocol esp
set security ipsec proposal prop_ipsec authentication-algorithm hmac-sha1-96
set security ipsec proposal prop_ipsec encryption-algorithm aes-192-cbc
set security ipsec proposal prop_ipsec lifetime-seconds 3600
set security ipsec proposal prop_ipsec lifetime-kilobytes 150000
set security ipsec policy ipsecpol1 perfect-forward-secrecy keys group5
set security ipsec policy ipsecpol1 proposals prop_ipsec
set security ipsec vpn SPOKE_VPN bind-interface st0.1
set security ipsec vpn SPOKE_VPN ike gateway SPOKE_GW
set security ipsec vpn SPOKE_VPN ike ipsec-policy ipsecpol1
set security ipsec vpn SPOKE_VPN traffic-selector ts1 local-ip 172.0.0.0/8
set security ipsec vpn SPOKE_VPN traffic-selector ts1 remote-ip 192.0.0.0/8
set security ipsec vpn SPOKE_VPN establish-tunnels immediately
set security pki ca-profile rsa ca-identity rsa
set security pki ca-profile rsa revocation-check disable
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces ge-0/0/3.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces ge-0/0/1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the hub:

1. Configure interfaces.

   ```
   [edit interfaces]
   user@host# set ge-0/0/1 unit 0 family inet address 172.16.1.1/24
   user@host# set ge-0/0/3 unit 0 family inet address 10.2.2.253/24
   user@host# set st0 unit 1 family inet
   ```

2. Configure Phase 1 options.

   ```
   [edit security ike proposal prop_ike]
   user@host# set authentication-method rsa-signatures
   ```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy ikepol1]
user@host# set mode main
user@host# set proposals prop_ike
user@host# set certificate local-certificate Spoke1_ID

[edit security ike gateway SPOKE_GW]
user@host# set ike-policy ikepol1
user@host# set address 10.2.2.1
user@host# set local-identity distinguished-name
user@host# set remote-identity distinguished-name container
  DC=Domain_component
user@host# set external-interface ge-0/0/3.0
user@host# set version v1-only
```

3. Configure Phase 2 options.

```
[edit security ipsec proposal prop_ipsec]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-192-cbc
user@host# set lifetime-seconds 3600
user@host# set lifetime-kilobytes 150000

[edit security ipsec policy ipsecpol1]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals prop_ipsec

[edit security ipsec SPOKE_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway SPOKE_GW
user@host# set ike ipsec-policy ipsecpol1
user@host# set traffic-selector ts1 local-ip 172.0.0.0/8
user@host# set traffic-selector ts1 remote-ip 192.0.0.0/8
user@host# set establish-tunnels immediately
```

4. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile rsa ca-identity rsa
user@host# set ca-profile rsa revocation-check disable
```

5. Configure security zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces ge-0/0/3.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces ge-0/0/1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            address 172.16.1.1/24;
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 10.2.2.253/24;
        }
    }
}
st0 {
    unit 1 {
        family inet;
    }
}
[edit]
user@host# show security ike
proposal prop_ike {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
}
policy ikepol1 {
    mode main;
    proposals prop_ike;
    certificate {
        local-certificate Spoke1_ID;
    }
}
gateway SPOKE_GW {
    ike-policy ikepol1;
    address 10.2.2.1;
    local-identity distinguished-name;
    remote-identity distinguished-name container DC=Domain_component;
    external-interface ge-0/0/3.0;
    version v1-only;
}
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show security ipsec
proposal prop_ipsec {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-192-cbc;
    lifetime-seconds 3600;
    lifetime-kilobytes 150000;
}
policy ipsecpol1 {
    perfect-forward-secrecy {
        keys group5;
    }
    proposals prop_ipsec;
}
vpn SPOKE_VPN {
    bind-interface st0.1;
    ike {
        gateway SPOKE_GW;
        ipsec-policy ipsecpol1;
    }
    traffic-selector ts1 {
        local-ip 172.0.0.0/8;
        remote-ip 192.0.0.0/8;
    }
    establish-tunnels immediately;
}
[edit]
user@host# show security pki
ca-profile rsa {
    ca-identity rsa;
    revocation-check {
        disable;
    }
}
[edit]
user@host# show security zones
security-zone trust {
    host-inbound-traffic {
        system-services {
            all;
        }
        protocols {
            all;
        }
    }
    interfaces {
        st0.1;
        ge-0/0/3.0;
    }
}
security-zone untrust {
    host-inbound-traffic {
        system-services {
            all;
        }
```


Copyright © 2016, Juniper Networks, Inc.

```
      protocols {
        all;
      }
    }
    interfaces {
      ge-0/0/1.0;
    }
  }
  [edit]
  user@host# show security policies
  default-policy {
    permit-all;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Tunnels on page 6621
- Verifying Traffic Selectors on page 6623

Verifying Tunnels

Purpose Verify that tunnels are established between the AutoVPN hub and spoke.

Action From operational mode, enter the **show security ike security-associations** and **show security ipsec security-associations** commands on the hub.

```
user@host> show security ike security-associations
node0:
--------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

1350248074 UP  d195bce6ccfcf9af  8f1569c6592c8408  Main           10.2.2.253


user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <77594650 ESP:aes-cbc-192/sha1 ac97cb1 2799/  150000 - root 500 10.2.2.253

  >77594650 ESP:aes-cbc-192/sha1 828dc013 2798/  150000 - root 500 10.2.2.253


user@host> show security ipsec security-associations detail
node0:
--------------------------------------------------------------------------

ID: 77594650 Virtual-system: root, VPN Name: HUB_VPN
  Local Gateway: 10.2.2.1, Remote Gateway: 10.2.2.253
  Traffic Selector Name: ts1
  Local Identity: ipv4(192.0.0.0-192.255.255.255)
  Remote Identity: ipv4(172.0.0.0-172.255.255.255)
  Version: IKEv1
```

Copyright © 2016, Juniper Networks, Inc.

```
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 0x24608b29
  Tunnel events:
    Tue Dec 30 2014 11:30:21 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Dec 30 2014 11:30:20 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Dec 30 2014 11:30:20 -0800: IKE SA negotiation successfully completed (3
 times)
  Location: FPC 5, PIC 0, KMD-Instance 1
  Direction: inbound, SPI: ac97cb1, AUX-SPI: 0
    Hard lifetime: Expires in 2796 seconds
    Lifesize Remaining:  150000 kilobytes
    Soft lifetime: Expires in 2211 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Location: FPC 5, PIC 0, KMD-Instance 1
  Direction: outbound, SPI: 828dc013, AUX-SPI: 0
    Hard lifetime: Expires in 2796 seconds
    Lifesize Remaining:  150000 kilobytes
    Soft lifetime: Expires in 2211 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

From operational mode, enter the **show security ike security-associations** and **show security ipsec security-associations** commands on the spoke.

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

276505646 UP   d195bce6ccfcf9af  8f1569c6592c8408  Main           10.2.2.1


user@host> show security ipsec security-associations
  Total active tunnels: 1
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <69206018 ESP:aes-cbc-192/sha1 828dc013 2993/  150000 - root 500 10.2.2.1

  >69206018 ESP:aes-cbc-192/sha1 ac97cb1 2993/  150000 - root 500 10.2.2.1


user@host> show security ipsec security-associations detail
ID: 69206018 Virtual-system: root, VPN Name: SPOKE_VPN
  Local Gateway: 10.2.2.253, Remote Gateway: 10.2.2.1
  Traffic Selector Name: ts1
  Local Identity: ipv4(172.0.0.0-172.255.255.255)
  Remote Identity: ipv4(192.0.0.0-192.255.255.255)
  Version: IKEv1
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x2c608b29
  Tunnel events:
    Tue Dec 30 2014 11:30:20 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Dec 30 2014 11:30:20 -0800: IKE SA negotiation successfully completed (1
 times)
    Tue Dec 30 2014 11:26:11 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
  Location: FPC 1, PIC 0, KMD-Instance 1
  Direction: inbound, SPI: 828dc013, AUX-SPI: 0
    Hard lifetime: Expires in 2991 seconds
```


Copyright © 2016, Juniper Networks, Inc.

```
    Lifesize Remaining:  150000 kilobytes
    Soft lifetime: Expires in 2369 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Location: FPC 1, PIC 0, KMD-Instance 1
  Direction: outbound, SPI: ac97cb1, AUX-SPI: 0
    Hard lifetime: Expires in 2991 seconds
    Lifesize Remaining:  150000 kilobytes
    Soft lifetime: Expires in 2369 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. The hub shows one active tunnel to the spoke while the spoke shows one active tunnel to the hub.

If no SAs are listed for IKE Phase 1, then there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and spoke.

If no SAs are listed for IKE Phase 2, then there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and spoke.

Verifying Traffic Selectors

Purpose Verify the traffic selectors.

Action From operational mode, enter the **show security ipsec traffic-selector interface-name st0.1** command on the hub.

```
user@host> show security ipsec traffic-selector interface-name st0.1
node0:
--------------------------------------------------------------------------
 Source IP                          Destination IP                   Interface
Tunnel-id    IKE-ID
 192.0.0.0-192.255.255.255            172.0.0.0-172.255.255.255          st0.1
    77594650     DC=Domain_component, CN=Spoke1_ID, OU=Sales, O=XYZ, L=Sunnyvale,
 ST=CA, C=US
```

From operational mode, enter the **show security ipsec traffic-selector interface-name st0.1** command on the spoke.

```
user@host> show security ipsec traffic-selector interface-name st0.1
 Source IP                          Destination IP                   Interface
Tunnel-id    IKE-ID
 172.0.0.0-172.255.255.255            192.0.0.0-192.255.255.255          st0.1
    69206018     DC=Domain_component, CN=Hub_ID, OU=Sales, O=XYZ, L=Sunnyvale,
ST=CA, C=US
```

Meaning A traffic selector (also known as a proxy ID in IKEv1) is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses. Only traffic that conforms to a traffic selector is permitted through an SA.

Copyright © 2016, Juniper Networks, Inc.

Traffic selectors are negotiated between the initiator and the responder (the SRX Series hub).

Related Documentation

- Understanding AutoVPN with Traffic Selectors on page 6607
- Understanding Traffic Selectors in Route-Based VPNs on page 6195

Example: Ensuring VPN Tunnel Availability with AutoVPN and Traffic Selectors

Georedundancy is the deployment of multiple geographically distant sites so that traffic can continue to flow over a provider network even if there is a power outage, a natural disaster, or other catastrophic event that affects a site. In a mobile provider network, multiple Evolved Node B (eNodeB) devices can be connected to the core network through georedundant IPsec VPN gateways on SRX Series devices. The alternate routes to the eNodeB devices are distributed to the core network using a dynamic routing protocol.

This example configures AutoVPN hubs with multiple traffic selectors on SRX Series devices to ensure that there are georedundant IPsec VPN gateways to eNodeB devices. Auto route insertion (ARI) is used to automatically insert routes toward the eNodeB devices in the routing tables on the hubs. ARI routes are then distributed to the provider's core network through BGP.

- Requirements on page 6624
- Overview on page 6625
- Configuration on page 6627
- Verification on page 6642

Requirements

This example uses the following hardware and software components:

- Two SRX Series devices connected and configured in a chassis cluster. The chassis cluster is AutoVPN hub A.
- An SRX Series device configured as AutoVPN hub B.
- Junos OS Release 12.3X48-D10 or later.
- eNodeB devices that can establish IPsec VPN tunnels with AutoVPN hubs. eNodeB devices are third-party network equipment providers that initiate a VPN tunnel with AutoVPN hubs.
- Digital certificates enrolled in the hubs and the eNodeB devices that allow the devices to authenticate each other.

Before you begin:

- Obtain the address of the certificate authority (CA) and the information they require (such as the challenge password) when you submit requests for local certificates. See "Understanding Local Certificate Requests" on page 6371.

Copyright © 2016, Juniper Networks, Inc.

- Enroll the digital certificates in each device. See "Understanding Certificate Loading" on page 6377.

NOTE: This example uses the BGP dynamic routing protocol to advertise routes toward the eNodeB devices to the core network. For more information about specific requirements for using dynamic routing protocols, see the *Routing Protocols Overview for Security Devices*.

Overview

In this example, two AutoVPN hubs are configured with multiple traffic selectors on SRX Series devices to provide georedundant IPsec VPN gateways to eNodeB devices. ARI automatically inserts routes to the eNodeB devices in the routing tables on the hubs. ARI routes are then distributed to the provider's core network through BGP.

Certain Phase 1 and Phase 2 IKE tunnel options configured on the AutoVPN hubs and eNodeB devices must have the same values. Table 568 shows the values used in this example:

Table 568: Phase 1 and Phase 2 Options for Georedundant AutoVPN Hubs

Option	Value
IKE proposal:	
Authentication method	rsa-signatures
Diffie-Hellman (DH) group	group5
Authentication algorithm	sha-1
Encryption algorithm	aes-256-cbc
IKE policy:	
Certificate	local-certificate
IKE gateway:	
Dynamic	distinguished name wildcard DC=Common_component
IKE user type	group IKE id
Dead peer detection	probe-idle-tunnel
Local identity	distinguished name
Version	v2-only

Copyright © 2016, Juniper Networks, Inc.

Table 568: Phase 1 and Phase 2 Options for Georedundant AutoVPN Hubs *(continued)*

Option	Value
IPsec proposal:	
Protocol	esp
Authentication algorithm	hmac-sha1-96
Encryption algorithm	aes-256-cbc
IPsec policy:	
Perfect Forward Secrecy (PFS) group	group5

NOTE: In this example, the default security policy that permits all traffic is used for all devices. More restrictive security policies should be configured for production environments. See "Security Policies Overview" on page 1085. For simplicity, the configuration on the SRX Series devices allows all types of inbound traffic; this configuration is not recommended for production deployments.

Topology

Figure 295 shows the SRX Series devices to be configured for this example.

Figure 295: Georedundant IPsec VPN Gateways to eNodeB Devices

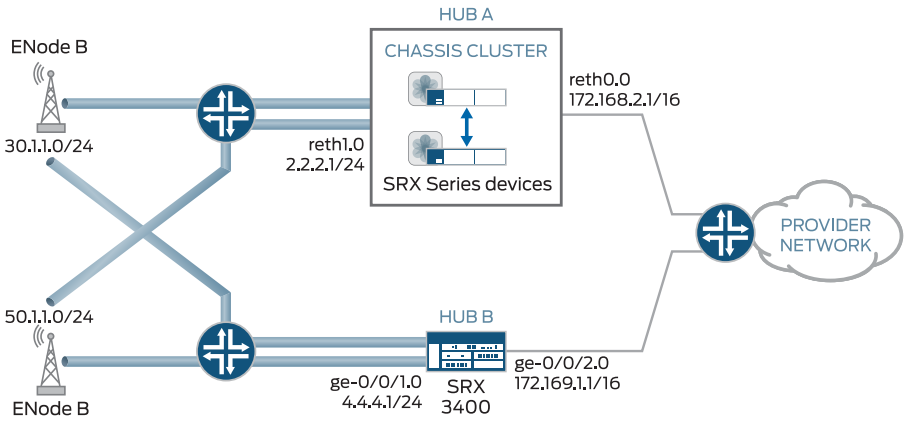

Copyright © 2016, Juniper Networks, Inc.

Configuration

- Configuring Hub A on page 6627
- Configuring Hub B on page 6634
- Configuring the eNodeB (Sample Configuration) on page 6641

Configuring Hub A

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/2 gigether-options redundant-parent reth1
set interfaces ge-0/0/3 gigether-options redundant-parent reth0
set interfaces ge-8/0/2 gigether-options redundant-parent reth1
set interfaces ge-8/0/3 gigether-options redundant-parent reth0
set interfaces lo0 unit 0 family inet address 100.100.1.100/24
set interfaces lo0 redundant-pseudo-interface-options redundancy-group 1
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 172.168.2.1/16
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 family inet address 2.2.2.1/24
set interfaces st0 unit 1 family inet
set security ike proposal prop_ike authentication-method rsa-signatures
set security ike proposal prop_ike dh-group group5
set security ike proposal prop_ike authentication-algorithm sha1
set security ike proposal prop_ike encryption-algorithm aes-256-cbc
set security ike policy ph1_ike_policy proposals prop_ike
set security ike policy ph1_ike_policy certificate local-certificate HubA_certificate
set security ike gateway HUB_GW ike-policy ph1_ike_policy
set security ike gateway HUB_GW dynamic distinguished-name wildcard
  DC=Common_component
set security ike gateway HUB_GW dynamic ike-user-type group-ike-id
set security ike gateway HUB_GW dead-peer-detection probe-idle-tunnel
set security ike gateway HUB_GW local-identity distinguished-name
set security ike gateway HUB_GW external-interface reth1
set security ike gateway HUB_GW version v2-only
set security ipsec proposal prop_ipsec protocol esp
set security ipsec proposal prop_ipsec authentication-algorithm hmac-sha1-96
set security ipsec proposal prop_ipsec encryption-algorithm aes-256-cbc
set security ipsec policy ph2_ipsec_policy perfect-forward-secrecy keys group5
set security ipsec policy ph2_ipsec_policy proposals prop_ipsec
set security ipsec vpn HUB_VPN bind-interface st0.1
set security ipsec vpn HUB_VPN ike gateway HUB_GW
set security ipsec vpn HUB_VPN ike ipsec-policy ph2_ipsec_policy
set security ipsec vpn HUB_VPN traffic-selector ts1 local-ip 172.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts1 remote-ip 50.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts2 local-ip 172.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts2 remote-ip 30.0.0.0/8
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 172.168.2.1
set protocols bgp group internal-peers export inject_ts1_routes
set protocols bgp group internal-peers export inject_ts2_routes
```

Copyright © 2016, Juniper Networks, Inc.

```
set protocols bgp group internal-peers export inject_up_routes
set protocols bgp group internal-peers neighbor 172.168.2.4
set policy-options policy-statement inject_ts1_routes term cp_allow from protocol static
set policy-options policy-statement inject_ts1_routes term cp_allow from route-filter
  30.1.2.0/24 orlonger
set policy-options policy-statement inject_ts1_routes term cp_allow from route-filter
  30.1.1.0/24 orlonger
set policy-options policy-statement inject_ts1_routes term cp_allow then next-hop self
set policy-options policy-statement inject_ts1_routes term cp_allow then accept
set policy-options policy-statement inject_ts2_routes term mp_allow from protocol static
set policy-options policy-statement inject_ts2_routes term mp_allow from route-filter
  50.1.1.0/24 orlonger
set policy-options policy-statement inject_ts2_routes term mp_net_allow from route-filter
  50.1.2.0/24 orlonger
set policy-options policy-statement inject_ts2_routes term mp_net_allow then next-hop
  self
set policy-options policy-statement inject_ts2_routes term mp_net_allow then accept
set policy-options policy-statement inject_up_routes term up_allow from protocol static
set policy-options policy-statement inject_up_routes term up_allow from route-filter
  172.168.1.0/24 orlonger
set policy-options policy-statement inject_up_routes term up_allow from route-filter
  172.168.2.0/24 orlonger
set policy-options policy-statement inject_up_routes term up_allow then next-hop self
set policy-options policy-statement inject_up_routes term up_allow then accept
set security pki ca-profile csa ca-identity csa
set security pki ca-profile csa revocation-check disable
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces reth0.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces reth1.0
set security policies default-policy permit-all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure hub A:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/2 gigether-options redundant-parent reth1
user@host# set ge-0/0/3 gigether-options redundant-parent reth0
user@host# set ge-8/0/2 gigether-options redundant-parent reth1
user@host# set ge-8/0/3 gigether-options redundant-parent reth0
user@host# set lo0 unit 0 family inet address 100.100.1.100/24
user@host# set lo0 redundant-pseudo-interface-options redundancy-group 1
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 172.168.2.1/16
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 family inet address 2.2.2.1/24
user@host# set st0 unit 1 family inet
```


Copyright © 2016, Juniper Networks, Inc.

2. Configure Phase 1 options.

```
[edit security ike proposal prop_ike]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy ph1_ike_policy]
user@host# set proposals prop_ike
user@host# set certificate local-certificate HubA_certificate

[edit security ike gateway HUB_GW]
user@host# set ike-policy ph1_ike_policy
user@host# set dynamic distinguished-name wildcard DC=Common_component
user@host# set dynamic ike-user-type group-ike-id
user@host# set dead-peer-detection probe-idle-tunnel
user@host# set local-identity distinguished-name
user@host# set external-interface reth1
user@host# set version v2-only
```

3. Configure Phase 2 options.

```
[edit security ipsec proposal prop_ipsec]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy ph2_ipsec_policy]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals prop_ipsec

[edit security ipsec vpn HUB_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway HUB_GW
user@host# set ike ipsec-policy ph2_ipsec_policy
user@host# set traffic-selector ts1 local-ip 172.0.0.0/8
user@host# set traffic-selector ts1 remote-ip 50.0.0.0/8
user@host# set traffic-selector ts2 local-ip 172.0.0.0/8
user@host# set traffic-selector ts2 remote-ip 30.0.0.0/8
```

4. Configure the BGP routing protocol.

```
[edit protocols bgp group internal-peers]
user@host# set type internal
user@host# set local-address 172.168.2.1
user@host# set export inject_ts1_routes
user@host# set export inject_ts2_routes
user@host# set export inject_up_routes
user@host# set neighbor 172.168.2.4
```

5. Configure routing options.

```
[edit policy-options policy-statement inject_ts1_routes]
user@host# set term cp_allow from protocol static
user@host# set term cp_allow from route-filter 30.1.2.0/24 orlonger
user@host# set term cp_allow from route-filter 30.1.1.0/24 orlonger
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set term cp_allow then next-hop self
user@host# set term cp_allow then accept

[edit policy-options policy-statement inject_ts2_routes]
user@host# set term mp_allow from protocol static
user@host# set term mp_allow from route-filter 50.1.1.0/24 orlonger
user@host# set term mp_allow from route-filter 50.1.2.0/24 orlonger
user@host# set term mp_allow then next-hop self
user@host# set term mp_allow then accept

[edit policy-options policy-statement inject_up_routes]
user@host# set term up_allow from protocol static
user@host# set term up_allow from route-filter 172.168.1.0/24 orlonger
user@host# set term up_allow from route-filter 172.168.2.0/24 orlonger
user@host# set term up_allow then next-hop self
user@host# set term up_allow then accept
```

6. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile csa ca-identity csa
user@host# set ca-profile csa revocation-check disable
```

7. Configure security zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces reth0.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces lo0.0
user@host# set interfaces reth1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security ike**, **show security ipsec**, **show protocols bgp**, **show policy-options**, **show security pki**, **show security zones**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
  ge-0/0/2 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-0/0/3 {
    gigether-options {
```

 Copyright © 2016, Juniper Networks, Inc.

```
        redundant-parent reth0;
      }
    }
    ge-8/0/2 {
      gigether-options {
        redundant-parent reth1;
      }
    }
    ge-8/0/3 {
      gigether-options {
        redundant-parent reth0;
      }
    }
    lo0 {
      unit 0 {
        family inet {
          address 100.100.1.100/24;
        }
      }
      redundant-pseudo-interface-options {
        redundancy-group 1;
      }
    }
    reth0 {
      redundant-ether-options {
        redundancy-group 1;
      }
      unit 0 {
        family inet {
          address 172.168.2.1/16;
        }
      }
    }
    reth1 {
      redundant-ether-options {
        redundancy-group 1;
      }
      unit 0 {
        family inet {
          address 2.2.2.1/24;
        }
      }
    }
    st0 {
      unit 1 {
        family inet;
      }
    }
  }
[edit]
user@host# show security ike
  proposal prop_ike {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
  }
```

Copyright © 2016, Juniper Networks, Inc.

```
    policy ph1_ike_policy {
      proposals prop_ike;
      certificate {
        local-certificate HubA_certificate;
      }
    }
    gateway HUB_GW {
      ike-policy ph1_ike_policy;
      dynamic {
        distinguished-name {
          wildcard DC=Common_component;
        }
        ike-user-type group-ike-id;
      }
      dead-peer-detection {
        probe-idle-tunnel;
      }
      local-identity distinguished-name;
      external-interface reth1;
      version v2-only;
    }
[edit]
user@host# show security ipsec
    proposal prop_ipsec {
      protocol esp;
      authentication-algorithm hmac-sha1-96;
      encryption-algorithm aes-256-cbc;
    }
    policy ph2_ipsec_policy {
      perfect-forward-secrecy {
        keys group5;
      }
      proposals prop_ipsec;
    }
    vpn HUB_VPN {
      bind-interface st0.1;
      ike {
        gateway HUB_GW;
        ipsec-policy ph2_ipsec_policy;
      }
      traffic-selector ts1 {
        local-ip 172.0.0.0/8;
        remote-ip 50.0.0.0/8;
      }
      traffic-selector ts2 {
        local-ip 172.0.0.0/8;
        remote-ip 30.0.0.0/8;
      }
    }
[edit]
user@host# show protocols bgp
    group internal-peers {
      type internal;
      local-address 172.168.2.1;
      export [ inject_ts1_routes inject_ts2_routes inject_up_routes ];
      neighbor 172.168.2.4;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
[edit]
user@host# show policy-options
policy-statement inject_ts1_routes {
    term cp_allow {
        from {
            protocol static;
            route-filter 30.1.2.0/24 orlonger;
            route-filter 30.1.1.0/24 orlonger;
        }
        then {
            next-hop self;
            accept;
        }
    }
}
policy-statement inject_ts2_routes {
    term mp_allow {
        from {
            protocol static;
            route-filter 50.1.1.0/24 orlonger;
            route-filter 50.1.2.0/24 orlonger;
        }
        then {
            next-hop self;
            accept;
        }
    }
}
policy-statement inject_up_routes {
    term up_allow {
        from {
            protocol static;
            route-filter 172.168.1.0/24 orlonger;
            route-filter 172.168.2.0/24 orlonger;
        }
        then {
            next-hop self;
            accept;
        }
    }
}
[edit]
user@host# show security pki
ca-profile csa {
    ca-identity csa;
    revocation-check {
        disable;
    }
}
[edit]
user@host# show security zones
    security-zone trust {
        host-inbound-traffic {
            system-services {
                all;
```

```
        }
        protocols {
          all;
        }
      }
      interfaces {
        st0.1;
        reth0.0;
      }
    }
    security-zone untrust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        lo0.0;
        reth1.0;
      }
    }
[edit]
user@host# show security policies
  default-policy {
    permit-all;
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Hub B

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 4.4.4.1/24
set interfaces ge-0/0/2 unit 0 family inet address 172.169.1.1/16
set interfaces lo0 unit 0 family inet address 100.100.1.101/24
set interfaces st0 unit 1 family inet
set security ike proposal prop_ike authentication-method rsa-signatures
set security ike proposal prop_ike dh-group group5
set security ike proposal prop_ike authentication-algorithm sha1
set security ike proposal prop_ike encryption-algorithm aes-256-cbc
set security ike policy ph1_ike_policy proposals prop_ike
set security ike policy ph1_ike_policy certificate local-certificate HubB_certificate
set security ike gateway HUB_GW ike-policy ph1_ike_policy
set security ike gateway HUB_GW dynamic distinguished-name wildcard
  DC=Common_component
set security ike gateway HUB_GW dynamic ike-user-type group-ike-id
set security ike gateway HUB_GW dead-peer-detection probe-idle-tunnel
set security ike gateway HUB_GW local-identity distinguished-name
```


Copyright © 2016, Juniper Networks, Inc.

```
set security ike gateway HUB_GW external-interface ge-0/0/1
set security ike gateway HUB_GW version v2-only
set security ipsec proposal prop_ipsec protocol esp
set security ipsec proposal prop_ipsec authentication-algorithm hmac-sha1-96
set security ipsec proposal prop_ipsec encryption-algorithm aes-256-cbc
set security ipsec policy ph2_ipsec_policy perfect-forward-secrecy keys group5
set security ipsec policy ph2_ipsec_policy proposals prop_ipsec
set security ipsec vpn HUB_VPN bind-interface st0.1
set security ipsec vpn HUB_VPN ike gateway HUB_GW
set security ipsec vpn HUB_VPN ike ipsec-policy ph2_ipsec_policy
set security ipsec vpn HUB_VPN traffic-selector ts1 local-ip 172.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts1 remote-ip 50.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts2 local-ip 172.0.0.0/8
set security ipsec vpn HUB_VPN traffic-selector ts2 remote-ip 30.0.0.0/8
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 172.169.1.1
set protocols bgp group internal-peers export inject_ts1_routes
set protocols bgp group internal-peers export inject_ts2_routes
set protocols bgp group internal-peers export inject_up_routes
set policy-options policy-statement inject_ts1_routes term cp_allow from protocol static
set policy-options policy-statement inject_ts1_routes term cp_allow from route-filter
  30.1.2.0/24 orlonger
set policy-options policy-statement inject_ts1_routes term cp_allow from route-filter
  30.1.1.0/24 orlonger
set policy-options policy-statement inject_ts1_routes term cp_allow then next-hop self
set policy-options policy-statement inject_ts1_routes term cp_allow then accept
set policy-options policy-statement inject_ts2_routes term mp_allow from protocol static
set policy-options policy-statement inject_ts2_routes term mp_allow from route-filter
  50.1.1.0/24 orlonger
set policy-options policy-statement inject_ts2_routes term mp_net_allow from route-filter
  50.1.2.0/24 orlonger
set policy-options policy-statement inject_ts2_routes term mp_net_allow then next-hop
  self
set policy-options policy-statement inject_ts2_routes term mp_net_allow then accept
set policy-options policy-statement inject_up_routes term up_allow from protocol static
set policy-options policy-statement inject_up_routes term up_allow from route-filter
  172.169.1.0/24 orlonger
set policy-options policy-statement inject_up_routes term up_allow from route-filter
  172.169.2.0/24 orlonger
set policy-options policy-statement inject_up_routes term up_allow then next-hop self
set policy-options policy-statement inject_up_routes term up_allow then accept
set security pki ca-profile csa ca-identity csa
set security pki ca-profile csa revocation-check disable
set security zones security-zone trust host-inbound-traffic system-services all
set security zones security-zone trust host-inbound-traffic protocols all
set security zones security-zone trust interfaces st0.1
set security zones security-zone trust interfaces ge-0/0/2.0
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security policies default-policy permit-all
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure hub B:

1. Configure interfaces.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 4.4.4.1/24
user@host# set ge-0/0/2 unit 0 family inet address 172.169.1.1/16
user@host# set lo0 unit 0 family inet address 100.100.1.101/24
user@host# set st0 unit 1 family inet
```

2. Configure Phase 1 options.

```
[edit security ike proposal prop_ike]
user@host# set authentication-method rsa-signatures
user@host# set dh-group group5
user@host# set authentication-algorithm sha1
user@host# set encryption-algorithm aes-256-cbc

[edit security ike policy ph1_ike_policy]
user@host# set proposals prop_ike
user@host# set certificate local-certificate HubB_certificate

[edit security ike gateway HUB_GW]
user@host# set ike-policy ph1_ike_policy
user@host# set dynamic distinguished-name wildcard DC=Common_component
user@host# set dynamic ike-user-type group-ike-id
user@host# set dead-peer-detection probe-idle-tunnel
user@host# set local-identity distinguished-name
user@host# set external-interface ge-0/0/1
user@host# set version v2-only
```

3. Configure Phase 2 options.

```
[edit security ipsec proposal prop_ipsec]
user@host# set protocol esp
user@host# set authentication-algorithm hmac-sha1-96
user@host# set encryption-algorithm aes-256-cbc

[edit security ipsec policy ph2_ipsec_policy]
user@host# set perfect-forward-secrecy keys group5
user@host# set proposals prop_ipsec

[edit security ipsec vpn HUB_VPN]
user@host# set bind-interface st0.1
user@host# set ike gateway HUB_GW
user@host# set ike ipsec-policy ph2_ipsec_policy
user@host# set traffic-selector ts1 local-ip 172.0.0.0/8
user@host# set traffic-selector ts1 remote-ip 50.0.0.0/8
user@host# set traffic-selector ts2 local-ip 172.0.0.0/8
user@host# set traffic-selector ts2 remote-ip 30.0.0.0/8
```

4. Configure the BGP routing protocol.

Copyright © 2016, Juniper Networks, Inc.

```
[edit protocols bgp group internal-peers]
user@host# set type internal
user@host# set local-address 172.169.1.1
user@host# set export inject_ts1_routes
user@host# set export inject_ts2_routes
user@host# set export inject_up_routes
user@host# set neighbor 172.169.1.2
```

5. Configure routing options.

```
[edit policy-options policy-statement inject_ts1_routes]
user@host# set term cp_allow from protocol static
user@host# set term cp_allow from route-filter 30.1.2.0/24 orlonger
user@host# set term cp_allow from route-filter 30.1.1.0/24 orlonger
user@host# set term cp_allow then next-hop self
user@host# set term cp_allow then accept

[edit policy-options policy-statement inject_ts2_routes]
user@host# set term mp_allow from protocol static
user@host# set term mp_allow from route-filter 50.1.1.0/24 orlonger
user@host# set term mp_allow from route-filter 50.1.2.0/24 orlonger
user@host# set term mp_allow then next-hop self
user@host# set term mp_allow then accept

[edit policy-options policy-statement inject_up_routes]
user@host# set term up_allow from protocol static
user@host# set term up_allow from route-filter 172.169.1.0/24 orlonger
user@host# set term up_allow from route-filter 172.169.2.0/24 orlonger
user@host# set term up_allow then next-hop self
user@host# set term up_allow then accept
```

6. Configure certificate information.

```
[edit security pki]
user@host# set ca-profile csa ca-identity csa
user@host# set ca-profile csa revocation-check disable
```

7. Configure security zones.

```
[edit security zones security-zone trust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces st0.1
user@host# set interfaces ge-0/0/2.0

[edit security zones security-zone untrust]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces lo0.0
user@host# set interfaces ge-0/0/1.0

[edit security policies]
user@host# set default-policy permit-all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**
show security ike, **show security ipsec**, **show protocols bgp**, **show security pki**, **show security**

Copyright © 2016, Juniper Networks, Inc.

zones, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 4.4.4.1/24;
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 172.169.1.1/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 100.100.1.101/24;
      }
    }
  }
  st0 {
    unit 1 {
      family inet;
    }
  }
[edit]
user@host# show security ike
  proposal prop_ike {
    authentication-method rsa-signatures;
    dh-group group5;
    authentication-algorithm sha1;
    encryption-algorithm aes-256-cbc;
  }
  policy ph1_ike_policy {
    proposals prop_ike;
    certificate {
      local-certificate HubB_certificate;
    }
  }
  gateway HUB_GW {
    ike-policy ph1_ike_policy;
    dynamic {
      distinguished-name {
        wildcard DC=Common_component;
      }
      ike-user-type group-ike-id;
    }
    dead-peer-detection {
      probe-idle-tunnel;
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    local-identity distinguished-name;
    external-interface reth1;
    version v2-only;
  }
[edit]
user@host# show security ipsec
  proposal prop_ipsec {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-256-cbc;
  }
  policy ph2_ipsec_policy {
    perfect-forward-secrecy {
      keys group5;
    }
    proposals prop_ipsec;
  }
  vpn HUB_VPN {
    bind-interface st0.1;
    ike {
      gateway HUB_GW;
      ipsec-policy ph2_ipsec_policy;
    }
    traffic-selector ts1 {
      local-ip 172.0.0.0/8;
      remote-ip 50.0.0.0/8;
    }
    traffic-selector ts2 {
      local-ip 172.0.0.0/8;
      remote-ip 30.0.0.0/8;
    }
  }
[edit]
user@host# show protocols bgp
  group internal-peers {
    type internal;
    local-address 172.169.1.1;
    export [ inject_ts1_routes inject_ts2_routes inject_up_routes ];
    neighbor 172.169.1.2;
  }
user@host# show policy-options
policy-statement inject_ts1_routes {
  term cp_allow {
    from {
      protocol static;
      route-filter 30.1.2.0/24 orlonger;
      route-filter 30.1.1.0/24 orlonger;
    }
    then {
      next-hop self;
      accept;
    }
  }
}
policy-statement inject_ts2_routes {
```

Copyright © 2016, Juniper Networks, Inc.

```
    term mp_allow {
      from {
        protocol static;
        route-filter 50.1.1.0/24 orlonger;
        route-filter 50.1.2.0/24 orlonger;
      }
      then {
        next-hop self;
        accept;
      }
    }
}
policy-statement inject_up_routes {
    term up_allow {
      from {
        protocol static;
        route-filter 172.169.1.0/24 orlonger;
        route-filter 172.169.2.0/24 orlonger;
      }
      then {
        next-hop self;
        accept;
      }
    }
}
[edit]
user@host# show security pki
ca-profile csa {
    ca-identity csa;
    revocation-check {
      disable;
    }
}
[edit]
user@host# show security zones
    security-zone trust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
        }
      }
      interfaces {
        st0.1;
        ge-0/0/2.0;
      }
    }
    security-zone untrust {
      host-inbound-traffic {
        system-services {
          all;
        }
        protocols {
          all;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      interfaces {
        ge-0/0/1.0;
        lo0.0;
      }
    }
  [edit]
  user@host# show security policies
    default-policy {
      permit-all;
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the eNodeB (Sample Configuration)

Step-by-Step Procedure

The eNodeB configuration in this example is provided for reference. Detailed eNodeB configuration information is beyond the scope of this document. The eNodeB configuration must include the following information:

- Local certificate (X.509v3) and IKE identity information
- SRX Series IKE identity information and public IP address
- Phase 1 and Phase 2 proposals that match the configurations on the SRX Series hubs

Results

The eNodeB devices in this example use strongSwan open source software for IPsec-based VPN connections:

```
config setup
    plutostart=yes
    plutodebug=all
    charondebug="ike 4, cfg 4, chd 4, enc 1"
    charonstart=yes  #ikev2 deamon"
    nat_traversal=yes  #<======== need to enable even no nat_t

conn %default
    ikelifetime=60m
    keylife=45m
    rekeymargin=2m
    keyingtries=4
    mobike=no

conn Hub_A
    keyexchange=ikev2
    authby=pubkey
    ike=aes256-sha-modp1536
    esp=aes256-sha1-modp1536
    leftcert=/usr/local/etc/ipsec.d/certs/fight02Req.pem.Email.crt
    left=5.5.5.1 # self if
    leftsubnet=30.1.1.0/24 # left subnet
    leftid="CN=fight02, DC=Common_component, OU=Dept, O=Company, L=City,
ST=CA, C=US " # self id
    right=2.2.2.1 # peer if
```

Copyright © 2016, Juniper Networks, Inc.

```
    rightsubnet=80.1.1.0/24 # peer net for proxy id
    rightid="DC=Domain_component, CN=HubA_certificate, OU=Dept, O=Company,
L=City, ST=CA, C=US " # peer id
    auto=add
    leftfirewall=yes
    dpdaction=restart
    dpddelay=10
    dpdtimeout=120
    rekeyfuzz=10%
    reauth=no

conn Hub_B
    keyexchange=ikev2
    authby=pubkey
    ike=aes256-sha-modp1536
    esp=aes192-sha1-modp1536
    leftcert=/usr/local/etc/ipsec.d/certs/fight02Req.pem.Email.crt
    left=5.5.5.1 # self if
    leftsubnet=30.1.1.0/24 # self net for proxy id
    leftid="CN=fight02, DC=Common_component, OU=Dept, O=Company, L=City,
ST=CA, C=US " # self id
    right=4.4.4.1 # peer if
    rightsubnet=80.1.1.0/24 # peer net for proxy id
    rightid="DC=Domain_component, CN=HubB_certificate, OU=Dept, O=Company,
L=City, ST=CA, C=US " # peer id
    auto=add
    leftfirewall=yes
    dpdaction=restart
    dpddelay=10
    dpdtimeout=120
    rekeyfuzz=10%
    reauth=no
```

Verification

Confirm that the configuration is working properly.

- Verifying Tunnels on the AutoVPN Hubs on page 6642
- Verifying Traffic Selectors on page 6643
- Verifying ARI Routes on page 6643

Verifying Tunnels on the AutoVPN Hubs

Purpose Verify that tunnels are established between the AutoVPN hub and eNodeB devices.

Action From operational mode, enter the **show security ike security-associations** and **show security ipsec security-associations** commands on the hub.

```
user@host> show security ike security-associations
node0:
------------------------------------------------------------------------
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

276505706 UP   16d6e53f0866b5cc  ccd8ca944da7b63e  IKEv2          5.5.5.1
```


Copyright © 2016, Juniper Networks, Inc.

```
1350247532 UP  d5f0cb3a3b18cb92  91269f05527217a0  IKEv2          1.1.1.1

user@host> show security ipsec security-associations
node0:
--------------------------------------------------------------------------
  Total active tunnels: 2
  ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
  <77594626 ESP:aes-cbc-192/sha1 a82bbc3 3600/  64 - root 500  1.1.1.1
  >77594626 ESP:aes-cbc-192/sha1 c930a858 3600/  64 - root 500 1.1.1.1
  <69206018 ESP:aes-cbc-192/sha1 2b437fc 3600/  64 - root 500  5.5.5.1
  >69206018 ESP:aes-cbc-192/sha1 c6e02755 3600/  64 - root 500 5.5.5.1
```

Meaning The **show security ike security-associations** command lists all active IKE Phase 1 SAs. The **show security ipsec security-associations** command lists all active IKE Phase 2 SAs. The hub shows two active tunnels, one to each eNodeB device.

If no SAs are listed for IKE Phase 1, then there was a problem with Phase 1 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 1 proposal parameters must match on the hub and eNodeB devices.

If no SAs are listed for IKE Phase 2, then there was a problem with Phase 2 establishment. Check the IKE policy parameters and external interface settings in your configuration. Phase 2 proposal parameters must match on the hub and eNodeB devices.

Verifying Traffic Selectors

Purpose Verify the traffic selectors.

Action From operational mode, enter the **show security ipsec traffic-selector interface-name st0.1** command.

```
user@host> show security ipsec traffic-selector interface-name st0.1
node0:
--------------------------------------------------------------------------
 Source IP                        Destination IP                  Interface
Tunnel-id    IKE-ID
 80.1.1.0-80.1.1.255              30.1.1.0-30.1.1.255             st0.1
69206018     DC=Common_component, CN=enodebA, OU=Dept, O=Company, L=City, ST=CA,
 C=US
 80.1.1.0-80.1.1.255              50.1.1.0-50.1.1.255             st0.1
77594626     DC=Common_component, CN=enodebB, OU=Dept, O=Company, L=City, ST=CA,
 C=US
```

Meaning A traffic selector (also known as a proxy ID in IKEv1) is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses. Only traffic that conforms to a traffic selector is permitted through an SA. Traffic selectors are negotiated between the initiator and the responder (the SRX Series hub).

Verifying ARI Routes

Purpose Verify that the ARI routes are added to the routing table.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show route** command.

```
user@host> show route
inet.0: 23 destinations, 23 routes (22 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.0.0/16          *[Static/5] 02:57:57
                     > to 2.2.2.253 via reth1.0
2.2.2.0/24          *[Direct/0] 02:58:43
                     > via reth1.0
2.2.2.1/32          *[Local/0] 02:59:25
                       Local via reth1.0
5.5.0.0/16          *[Static/5] 02:57:57
                     > to 2.2.2.253 via reth1.0
10.0.0.0/8          *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
10.157.64.0/19      *[Direct/0] 21:54:52
                     > via fxp0.0
10.157.75.117/32    *[Local/0] 21:54:52
                       Local via fxp0.0
10.254.75.117/32    *[Direct/0] 21:54:52
                     > via lo0.0
30.1.1.0/24         *[Static/5] 02:28:10  [ARI route added based on TSi]
                     > via st0.1
50.1.1.0/24         *[Static/5] 02:28:26
                     > via st0.1
66.129.230.0/24     *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
66.129.236.0/24     *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
80.0.0.0/8          *[Direct/0] 02:57:57
                     > via reth0.0
80.1.1.1/32         *[Local/0] 02:57:57
                       Local via reth0.0
100.100.1.0/24      *[Direct/0] 02:57:57
                     > via lo0.0
100.100.1.100/32    *[Local/0] 02:57:57
                       Local via lo0.0
102.100.1.0/24      *[Static/5] 02:57:57
                     > to 2.2.2.253 via reth1.0
104.100.1.0/24      *[Static/5] 02:57:57
                     > to 2.2.2.253 via reth1.0
172.16.0.0/12       *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
192.168.0.0/16      *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
207.17.136.0/24     *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
207.17.137.227/32   *[Static/5] 21:54:52
                     > to 10.157.64.1 via fxp0.0
```

Meaning Auto route insertion (ARI) automatically inserts a static route for the remote network and hosts protected by a remote tunnel endpoint. A route is created based on the remote IP address configured in the traffic selector. In the case of traffic selectors, the configured remote address is inserted as a route in the routing instance associated with the st0 interface that is bound to the VPN.

 Copyright © 2016, Juniper Networks, Inc.

Static routes to the eNodeB destinations 30.1.1.0/24 and 50.1.1.0/24 are added to the routing table on the SRX Series hub. These routes are reachable through the st0.1 interface.

Related Documentation

- Understanding AutoVPN with Traffic Selectors on page 6607
- Understanding Traffic Selectors in Route-Based VPNs on page 6195
- Understanding Auto Route Insertion on page 6213

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 87

Configuring Group VPNs

- Managing IPsec SA By Configuring Group VPNs on page 6649
- Configuring Server-Member Communication on page 6693
- Configuring Group VPNv2 Server Clusters on page 6699

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 275

Managing IPsec SA By Configuring Group VPNs

- Group VPNv2 Overview on page 6649
- Understanding the GDOI Protocol on page 6650
- Understanding Group Servers and Members on page 6651
- Understanding Group VPNv2 Limitations on page 6652
- Group VPNv2 Configuration Overview on page 6652
- Understanding IKE Phase 1 Configuration for Group VPNv2 on page 6653
- Understanding IPsec SA Configuration for Group VPNv2 on page 6654
- Understanding Group Configuration on page 6655
- Understanding Traffic Steering on page 6655
- Example: Configuring a Group VPNv2 Server and Members on page 6657

Group VPNv2 Overview

An IPsec security association (SA) is a unidirectional agreement between virtual private network (VPN) participants that defines the rules to use for authentication and encryption algorithms, key exchange mechanisms, and secure communications. With many VPN implementations, the SA is a point-to-point tunnel between two security devices (see Figure 296). Group VPNv2 extends IPsec architecture to support SAs that are shared by a group of security devices (see Figure 297).

Figure 296: Point-to-Point SAs

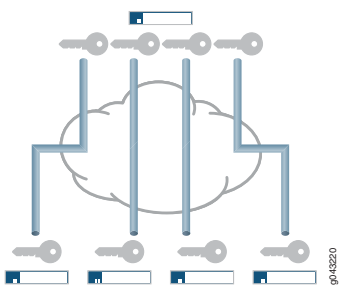

Copyright © 2016, Juniper Networks, Inc.

Figure 297: Shared SAs

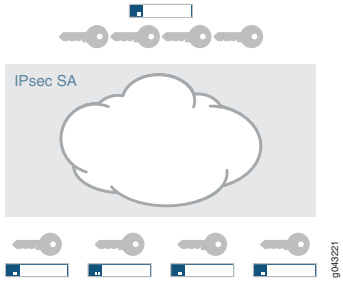

With Group VPNv2, any-to-any connectivity is achieved by preserving the original source and destination IP addresses in the outer header.

NOTE: Group VPNv2 is an enhanced version of the group VPN feature introduced in an earlier Junos OS release for SRX Series branch devices. Group VPNv2 supports RFC 6407, *The Group Domain of Interpretation (GDOI)*.

Related Documentation

- Understanding the GDOI Protocol on page 6650
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652

Understanding the GDOI Protocol

Group VPNv2 is based on RFC 6407, *The Group Domain of Interpretation* (GDOI). This RFC describes the protocol between group members and group servers to establish SAs among group members. GDOI messages create, maintain, or delete SAs for a group of devices. The GDOI protocol runs on UDP port 848.

The Internet Security Association and Key Management Protocol (ISAKMP) defines two negotiation phases to establish SAs for an IKE IPsec tunnel. Phase 1 allows two devices to establish an ISAKMP SA for other security protocols, such as GDOI.

With Group VPNv2, Phase 1 ISAKMP SA negotiation is performed between a group server and a group member. The server and member must use the same ISAKMP policy. GDOI exchanges between the server and member establish the SAs that are shared with other group members. A group member does not need to negotiate IPsec with other group members. GDOI exchanges must be protected by ISAKMP Phase 1 SAs.

There are two types of GDOI exchanges:

- The **groupkey-pull** exchange allows a member to request SAs and keys shared by the group from the server. Group members must register with a group server through a **groupkey-pull** exchange.

Copyright © 2016, Juniper Networks, Inc.

- The **groupkey-push** exchange is a single rekey message that allows the server to send group SAs and keys to members before existing group SAs expire. Rekey messages are unsolicited messages sent from the server to members.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Servers and Members on page 6651
- Understanding Group Key Operations on page 6694

Understanding Group Servers and Members

The center of Group VPNv2 is the group controller/key server (GCKS). A server cluster can be used to provide GCKS redundancy.

The GCKS or group server performs the following tasks:

- Controls group membership.
- Generates encryption keys.
- Sends new group SAs and keys to members. Group members encrypt traffic based on the group SAs and keys provided by the group server.

A group server can service multiple groups. A single security device can be a member of multiple groups.

Each group is represented by a group identifier, which is a number between 1 and 4,294,967,295. The group server and group members are linked together by the group identifier. There can be only one group identifier per group, and multiple groups cannot use the same group identifier.

The following is a high-level view of Group VPNv2 server and member actions:

1. The group server listens on UDP port 848 for members to register.
2. To register with the group server, the member first establishes an IKE SA with the server. A member device must provide correct IKE Phase 1 authentication to join the group. Preshared key authentication on a per-member basis is supported.
3. Upon successful authentication and registration, the member device retrieves group SAs and keys for the specified group identifier from the server with a GDOI **groupkey-pull** exchange.
4. The server adds the member to the membership for the group.
5. Group members exchange packets encrypted with group SA keys.

The server sends SA and key refreshes to group members with rekey (GDOI **groupkey-push**) messages. The server sends rekey messages before SAs expire to ensure that valid keys are available for encrypting traffic between group members.

A rekey message sent by the server requires an acknowledgement (ack) message from each group member. If the server does not receive an ack message from the member,

Copyright © 2016, Juniper Networks, Inc.

the rekey message is retransmitted at the configured **retransmission-period** (the default is 10 seconds). If there is no reply from the member after the configured **number-of-retransmission** (the default is 2 times), the member is removed from the server's registered members. The IKE SA between the server and member is also removed.

The server also sends rekey messages to provide new keys to members when the group SA has changed.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699
- Understanding the GDOI Protocol on page 6650
- Understanding Group VPNv2 Limitations on page 6652
- Group VPNv2 Configuration Overview on page 6652

Understanding Group VPNv2 Limitations

NOTE: Group VPNv2 servers only operate with Group VPNv2 members that support RFC 6407, *The Group Domain of Interpretation (GDOI)*.

The following are not supported in this release for Group VPNv2:

- SNMP.
- Deny policy from Cisco GET VPN server.
- PKI support for Phase 1 IKE authentication.
- Colocation of group server and member, where server and member functions coexist in the same physical device.
- Group members configured as chassis clusters.
- J-Web interface for configuration and monitoring.
- Multicast data traffic.

Group VPNv2 is not supported in deployments where IP addresses cannot be preserved—for example, across the Internet where NAT is used.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding the GDOI Protocol on page 6650
- Understanding Group Servers and Members on page 6651

Group VPNv2 Configuration Overview

This topic describes the main tasks for configuring Group VPNv2.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The group controller/key server (GCKS) manages Group VPNv2 security associations (SAs), and generates encryption keys and distributes them to group members. You can use a Group VPNv2 server cluster to provide GCKS redundancy. See "Understanding Group VPNv2 Server Clusters" on page 6699.

On the group server(s), configure the following:

1. IKE Phase 1 SA. See "Understanding IKE Phase 1 Configuration for Group VPNv2" on page 6653.
2. IPsec SA. See "Understanding IPsec SA Configuration for Group VPNv2" on page 6654.
3. VPN group information, including the group identifier, IKE gateways for group members, the maximum number of members in the group, and server-member communications. Group configuration includes a group policy that defines the traffic to which the SA and keys apply. Server cluster and antireplay time window can optionally be configured. See "Understanding Group Configuration" on page 6655 and "Understanding Traffic Steering" on page 6655.

On the group member, configure the following:

1. IKE Phase 1 SA. See "Understanding IKE Phase 1 Configuration for Group VPNv2" on page 6653.
2. IPsec SA. See "Understanding IPsec SA Configuration for Group VPNv2" on page 6654.
3. IPsec policy that defines the incoming zone (usually a protected LAN), outgoing zone (usually a WAN) and the VPN group to which the policy applies. Exclude or fail-open rules can also be specified. See "Understanding Traffic Steering" on page 6655.
4. Security policy to allow group VPN traffic between the zones specified in the IPsec policy.

NOTE: Group VPNv2 operation requires a working routing topology that allows client devices to reach their intended sites throughout the network.

Related Documentation

- Understanding Group Servers and Members on page 6651
- Understanding Server-Member Communication on page 6693

Understanding IKE Phase 1 Configuration for Group VPNv2

An IKE Phase 1 SA between a group server and a group member establishes a secure channel in which to negotiate IPsec SAs that are shared by a group. For standard IPsec VPNs on Juniper Networks security devices, Phase 1 SA configuration consists of specifying an IKE proposal, policy, and gateway.

Copyright © 2016, Juniper Networks, Inc.

For Group VPNv2, the IKE Phase 1 SA configuration is similar to the configuration for standard IPsec VPNs, but is performed at the [**edit security group-vpn server ike**] and [**edit security group-vpn member ike**] hierarchies.

In the IKE proposal configuration, you set the authentication method and the authentication and encryption algorithms that will be used to open a secure channel between participants. In the IKE policy configuration, you set the mode in which the Phase 1 channel will be negotiated, specify the type of key exchange to be used, and reference the Phase 1 proposal. In the IKE gateway configuration, you reference the Phase 1 policy.

The IKE proposal and policy configuration on the group server must match the IKE proposal and policy configuration on group members. On a group server, an IKE gateway is configured for each group member. On a group member, up to four server addresses can be specified in the IKE gateway configuration.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652
- Understanding IPsec SA Configuration for Group VPNv2 on page 6654

Understanding IPsec SA Configuration for Group VPNv2

After the server and member have established a secure and authenticated channel in Phase 1 negotiation, they proceed to establish the IPsec SAs that are shared by group members to secure data that is transmitted among members. While the IPsec SA configuration for Group VPNv2 is similar to the configuration for standard VPNs, a group member does not need to negotiate the SA with other group members.

IPsec configuration for Group VPNv2 consists of the following information:

- On the group server, an IPsec proposal is configured for the security protocol, authentication, and encryption algorithm to be used for the SA. The IPsec SA proposal is configured on the group server with the **proposal** configuration statement at the [**edit security group-vpn server ipsec**] hierarchy.
- On the group member, an Autokey IKE is configured that references the group identifier, the group server (configured with the **ike-gateway** configuration statement), and the interface used by the member to connect to group peers. The Autokey IKE is configured on the member with the **vpn** configuration statement at the [**edit security group-vpn member ipsec**] hierarchy.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652
- Understanding IKE Phase 1 Configuration for Group VPNv2 on page 6653

Copyright © 2016, Juniper Networks, Inc.

Understanding Group Configuration

The group is configured on the server with the **group** configuration statement at the [**edit security group-vpn server**] hierarchy.

The group information consists of the following information:

- Group identifier—A value that identifies the VPN group. The same group identifier must be configured on the group member.
- Each group member is configured with the **ike-gateway** configuration statement. There can be multiple instances of this configuration statement, one for each member of the group.
- Group policies—Policies that are to be downloaded to members. Group policies describe the traffic to which the SA and keys apply. See "Understanding Traffic Steering" on page 6655.
- Member threshold—The maximum number of members in the group. After the member threshold for a group is reached, a server stops responding to **groupkey-pull** initiations from new members. See "Understanding Group VPNv2 Server Clusters" on page 6699.
- Server-member communication—Optional configuration that allows the server to send **groupkey-push** rekey messages to members. See "Understanding Server-Member Communication" on page 6693.
- Server cluster—Optional configuration that supports group controller/key server (GCKS) redundancy. See "Understanding Group VPNv2 Server Clusters" on page 6699.
- Antireplay—Optional configuration that detects packet interception and replay. See "Understanding Antireplay" on page 6697.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652

Understanding Traffic Steering

The group server distributes IPsec security associations (SAs) and keys to members of a specified group. All members that belong to the same group share the same set of IPsec SAs. The SA that is installed on a specific group member is determined by the policy associated with the group SA and the IPsec policy that is configured on the group member.

- Group Policies Configured on Group Servers on page 6656
- IPsec Policies Configured on Group Members on page 6656
- Fail-Close on page 6656
- Exclude and Fail-Open Rules on page 6656
- Priorities of IPsec Policies and Rules on page 6657

Copyright © 2016, Juniper Networks, Inc.

Group Policies Configured on Group Servers

In a VPN group, each group SA and key that the server pushes to a member are associated with a group policy. The group policy describes the traffic on which the key should be used, including protocol, source address, source port, destination address, and destination port. On the server, the group policy is configured with the **match-policy** *policy-name* options at the [**edit security group-vpn server group** *name* **ipsec-sa** *name*] hierarchy level.

NOTE: Group policies that are identical (configured with the same source address, destination address, source port, destination port, and protocol values) cannot exist for a single group. An error is returned if you attempt to commit a configuration that contains identical group policies for a group. If this occurs, you must delete one of the identical group policies before you can commit the configuration.

IPsec Policies Configured on Group Members

On the group member, an IPsec policy consists of the following information:

- Incoming zone (**from-zone**) for group traffic.
- Outgoing zone (**to-zone**) for group traffic.
- The name of the group to which the IPsec policy applies. Only one Group VPNv2 name can be referenced by a specific from-zone/to-zone pair.

NOTE: The interface that is used by the group member to connect to the Group VPNv2 must belong to the outgoing zone. This interface is specified with the **group-vpn-external-interface** statement at the [**edit security group-vpn member ipsec vpn** *vpn-name*] hierarchy level.

On the group member, the IPsec policy is configured at the [**edit security ipsec-policy**] hierarchy level. Traffic that matches the IPsec policy is further checked against exclude and fail-open rules that are configured for the group.

Fail-Close

By default, traffic that does not match exclude or fail-open rules or group policies received from the group server is blocked; this is known as fail-close.

Exclude and Fail-Open Rules

On group members, the following types of rules can be configured for each group:

- Traffic that is excluded from VPN encryption. Examples of this type of traffic can include BGP or OSPF routing protocols. To exclude traffic from a group, use the **set security group-vpn member ipsec vpn** *vpn-name* **exclude rule** configuration. A maximum of 10 exclude rules can be configured.

Copyright © 2016, Juniper Networks, Inc.

- Traffic that is critical to the customer's operation and must be sent in cleartext (unencrypted) if the group member has not received a valid traffic encryption key (TEK) for the IPsec SA. Fail-open rules allow this traffic flow while all other traffic is blocked. Enable fail-open with the **set security group-vpn member ipsec vpn *vpn-name* fail-open rule** configuration. A maximum of 10 fail-open rules can be configured.

Priorities of IPsec Policies and Rules

IPsec policies and rules have the following priorities on the group member:

1. Exclude rules that define traffic to be excluded from VPN encryption.
2. Group policies that are downloaded from the group server.
3. Fail-open rules that define raffic that is sent in cleartext if there is no valid TEK for the SA.
4. Fail-close policy that blocks traffic. This is the default if traffic does not match exclude or fail-open rules or group policies.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652

Example: Configuring a Group VPNv2 Server and Members

This example shows how to configure a Group VPNv2 server to provide group controller/key server (GCKS) support to Group VPNv2 group members.

- Requirements on page 6657
- Overview on page 6658
- Configuration on page 6659
- Verification on page 6686

Requirements

The example uses the following hardware and software components:

- A supported SRX Series device or vSRX instance running Junos OS Release 15.1X49-D30 or later that supports Group VPNv2. This SRX Series device or vSRX instance operates as a Group VPNv2 server.
- Two supported SRX Series devices or vSRX instances running Junos OS Release 15.1X49-D30 or later that support Group VPNv2. These devices or instances operate as Group VPNv2 group members.
- Two supported MX Series devices running Junos OS Release 15.1R2 or later that support Group VPNv2. These devices operate as Group VPNv2 group members.

Copyright © 2016, Juniper Networks, Inc.

A hostname, a root administrator password, and management access must be configured on each device. We recommend that NTP also be configured on each device.

NOTE: Group VPNv2 operation requires a working routing topology that allows client devices to reach their intended sites throughout the network. This examples focuses on the Group VPNv2 configuration; the routing configuration is not described.

Overview

In this example, the Group VPNv2 network consists of a server and four members. Two of the members are SRX Series devices or vSRX instances while the other two members are MX Series devices. The shared group VPN SAs secure traffic between group members.

The group VPN SAs must be protected by a Phase 1 SA. Therefore, the group VPN configuration must include configuring IKE Phase 1 negotiations on both the group server and the group members.

The same group identifier must be configured on both the group server and the group members. In this example, the group name is GROUP_ID-0001 and the group identifier is 1. The group policy configured on the server specifies that the SA and key are applied to traffic between subnetworks in the 172.16.0.0/12 range.

On SRX or vSRX group members, an IPsec policy is configured for the group with the LAN zone as the from-zone (incoming traffic) and the WAN zone as the to-zone (outgoing traffic). A security policy is also needed to allow traffic between the LAN and WAN zones.

Topology

Figure 298 shows the Juniper Networks devices to be configured for this example.

Copyright © 2016, Juniper Networks, Inc.

Figure 298: Group VPNv2 Server with SRX or vSRX and MX Series Members

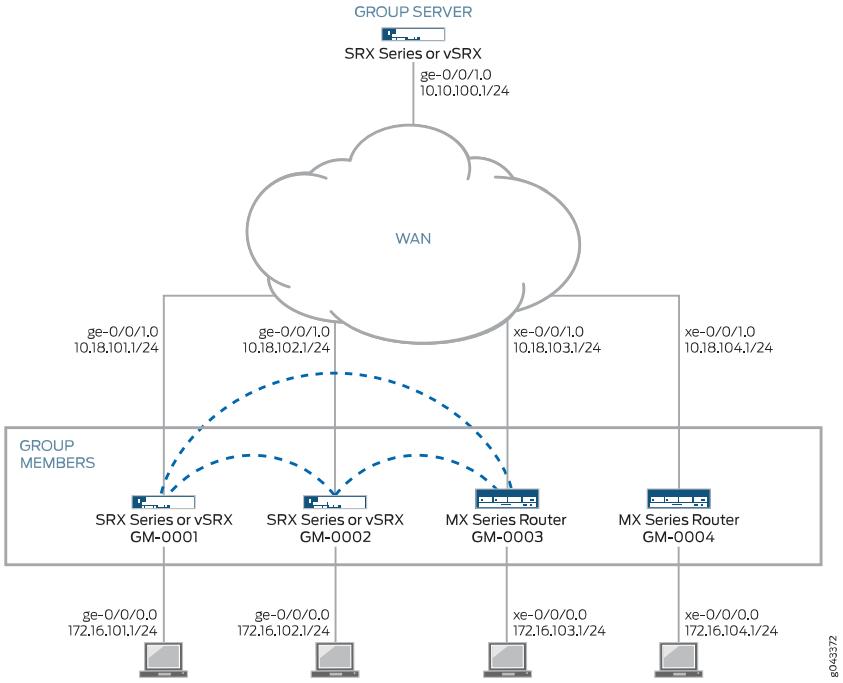

Configuration

- Configuring the Group Server on page 6659
- Configuring Group Member GM-0001 (SRX Series Device or vSRX Instance) on page 6665
- Configuring Group Member GM-0002 (SRX Series Device or vSRX Instance) on page 6671
- Configuring Group Member GM-0003 (MX Series Device) on page 6677
- Configuring Group Member GM-0004 (MX Series Device) on page 6681

Configuring the Group Server

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/1 unit 0 family inet address 10.10.100.1/24
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then reject
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
```

Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces ge-0/0/1.0
set routing-options static route 10.18.101.0/24 next-hop 10.10.100.254
set routing-options static route 10.18.102.0/24 next-hop 10.10.100.254
set routing-options static route 10.18.103.0/24 next-hop 10.10.100.254
set routing-options static route 10.18.104.0/24 next-hop 10.10.100.254
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy GMs mode main
set security group-vpn server ike policy GMs proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy GMs pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway GM-0001 ike-policy GMs
set security group-vpn server ike gateway GM-0001 address 10.18.101.1
set security group-vpn server ike gateway GM-0001 local-address 10.10.100.1
set security group-vpn server ike gateway GM-0002 ike-policy GMs
set security group-vpn server ike gateway GM-0002 address 10.18.102.1
set security group-vpn server ike gateway GM-0002 local-address 10.10.100.1
set security group-vpn server ike gateway GM-0003 ike-policy GMs
set security group-vpn server ike gateway GM-0003 address 10.18.103.1
set security group-vpn server ike gateway GM-0003 local-address 10.10.100.1
set security group-vpn server ike gateway GM-0004 ike-policy GMs
set security group-vpn server ike gateway GM-0004 address 10.18.104.1
set security group-vpn server ike gateway GM-0004 local-address 10.10.100.1
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0001
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0002
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0003
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0004
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0005
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
```


Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 server:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family inet address 10.10.100.1/24

[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/1.0

[edit security policies]
user@host# set global policy 1000 match source-address any
user@host# set global policy 1000 match destination-address any
user@host# set global policy 1000 match application any
user@host# set global policy 1000 match from-zone any
user@host# set global policy 1000 match to-zone any
user@host# set global policy 1000 then reject
user@host# set global policy 1000 then log session-init
user@host# set global policy 1000 then count
user@host# set default-policy deny-all
```

2. Configure the static routes.

```
[edit routing-options]
user@host# set static route 10.18.101.0/24 next-hop 10.10.100.254
user@host# set static route 10.18.102.0/24 next-hop 10.10.100.254
user@host# set static route 10.18.103.0/24 next-hop 10.10.100.254
user@host# set static route 10.18.104.0/24 next-hop 10.10.100.254
```

3. Configure the IKE proposal, policy, and gateways.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set authentication-algorithm sha-256
user@host# set dh-group group14
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy GMs]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway GM-0001]
user@host# set ike-policy GMs
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set address 10.18.101.1
user@host# set local-address 10.10.100.1

[edit security group-vpn server ike gateway GM-0002]
user@host# set ike-policy GMs
user@host# set address 10.18.102.1
user@host# set local-address 10.10.100.1

[edit security group-vpn server ike gateway GM-0003]
user@host# set ike-policy GMs
user@host# set address 10.18.103.1
user@host# set local-address 10.10.100.1

[edit security group-vpn server ike gateway GM-0004]
user@host# set ike-policy GMs
user@host# set address 10.18.104.1
user@host# set local-address 10.10.100.1
```

4. Configure the IPsec proposal.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600 VPN Group
```

5. Configure the group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set ike-gateway GM-0001
user@host# set ike-gateway GM-0002
user@host# set ike-gateway GM-0003
user@host# set ike-gateway GM-0004
user@host# set anti-replay-time-window 1000
```

6. Configure server-to-member communications.

```
[edit security group-vpn server group GROUP_ID-0001
  server-member-communication]
user@host# set communication-type unicast
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 7200
user@host# set sig-hash-algorithm sha-256
```

7. Configure the group policy to be downloaded to the group members.

```
[edit security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001]
user@host# set proposal AES256-SHA256-L3600
user@host# set match-policy 1 source 172.16.0.0/12
user@host# set match-policy 1 destination 172.16.0.0/12
user@host# set match-policy 1 protocol 0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
   unit 0 {
      family inet {
         address 10.10.100.1/24;
      }
   }
}
[edit]
user@host# show routing-options
static {
   route 10.18.101.0/24 next-hop 10.10.100.254;
   route 10.18.102.0/24 next-hop 10.10.100.254;
   route 10.18.103.0/24 next-hop 10.10.100.254;
   route 10.18.104.0/24 next-hop 10.10.100.254;
}
[edit]
user@host# show security
group-vpn {
   server {
      ike {
         proposal PSK-SHA256-DH14-AES256 {
            authentication-method pre-shared-keys;
            authentication-algorithm sha-256;
            dh-group group14;
            encryption-algorithm aes-256-cbc;
         }
         policy GMs {
            mode main;
            proposals PSK-SHA256-DH14-AES256;
            pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
         }
         gateway GM-0001 {
            ike-policy GMs;
            address 10.18.101.1;
            local-address 10.10.100.1;
         }
         gateway GM-0002 {
            ike-policy GMs;
            address 10.18.102.1;
            local-address 10.10.100.1;
         }
         gateway GM-0003 {
            ike-policy GMs;
            address 10.18.103.1;
            local-address 10.10.100.1;
         }
         gateway GM-0004 {
            ike-policy GMs;
            address 10.18.104.1;
            local-address 10.10.100.1;
         }
      }
      ipsec {
         proposal AES256-SHA256-L3600 {
```

Copyright © 2016, Juniper Networks, Inc.

```
            authentication-algorithm hmac-sha-256-128;
            encryption-algorithm aes-256-cbc;
            lifetime-seconds 3600;
          }
        }
        group GROUP_ID-0001 {
          group-id 1;
          member-threshold 2000;
          ike-gateway GM-0001;
          ike-gateway GM-0002;
          ike-gateway GM-0003;
          ike-gateway GM-0004;
          anti-replay-time-window 1000;
          server-member-communication {
            communication-type unicast;
            lifetime-seconds 7200;
            encryption-algorithm aes-256-cbc;
            sig-hash-algorithm sha-256;
          }
          ipsec-sa GROUP_ID-0001 {
            proposal AES256-SHA256-L3600;
            match-policy 1 {
              source 172.16.0.0/12;
              destination 172.16.0.0/12;
              protocol 0;
            }
          }
        }
      }
    }
    policies {
      global {
        policy 1000 {
          match {
            source-address any;
            destination-address any;
            application any;
            from-zone any;
            to-zone any;
          }
          then {
            reject;
            log {
              session-init;
            }
            count;
          }
        }
      }
      default-policy {
        deny-all;
      }
    }
    zones {
      security-zone GROUPVPN {
        host-inbound-traffic {
```


Copyright © 2016, Juniper Networks, Inc.

```
        system-services {
            ike;
            ssh;
            ping;
        }
    }
    interfaces {
        ge-0/0/1.0;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Group Member GM-0001 (SRX Series Device or vSRX Instance)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_LAN
set interfaces ge-0/0/0 unit 0 family inet address 172.16.101.1/24
set interfaces ge-0/0/1 unit 0 description To_KeySrv
set interfaces ge-0/0/1 unit 0 family inet address 10.18.101.1/24
set security zones security-zone LAN host-inbound-traffic system-services ike
set security zones security-zone LAN host-inbound-traffic system-services ssh
set security zones security-zone LAN host-inbound-traffic system-services ping
set security zones security-zone LAN interfaces ge-0/0/0.0
set security zones security-zone WAN host-inbound-traffic system-services ike
set security zones security-zone WAN host-inbound-traffic system-services ssh
set security zones security-zone WAN host-inbound-traffic system-services ping
set security zones security-zone WAN interfaces ge-0/0/1.0
set security address-book global address 172.16.0.0/12 172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match application any
set security policies from-zone LAN to-zone WAN policy 1 then permit
set security policies from-zone LAN to-zone WAN policy 1 then log session-init
set security policies from-zone WAN to-zone LAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match application any
set security policies from-zone WAN to-zone LAN policy 1 then permit
set security policies from-zone WAN to-zone LAN policy 1 then log session-init
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then reject
set security policies global policy 1000 then log session-init
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies global policy 1000 then count
set security policies default-policy deny-all
set routing-options static route 10.18.102.0/24 next-hop 10.18.101.254
set routing-options static route 10.18.103.0/24 next-hop 10.18.101.254
set routing-options static route 10.18.104.0/24 next-hop 10.18.101.254
set routing-options static route 172.16.101.0/24 next-hop 10.18.101.254
set routing-options static route 172.16.102.0/24 next-hop 10.18.101.254
set routing-options static route 172.16.103.0/24 next-hop 10.18.101.254
set routing-options static route 172.16.104.0/24 next-hop 10.18.101.254
set routing-options static route 10.10.100.0/24 next-hop 10.18.101.254
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy KeySrv mode main
set security group-vpn member ike policy KeySrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy KeySrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway KeySrv ike-policy KeySrv
set security group-vpn member ike gateway KeySrv server-address 10.10.100.1
set security group-vpn member ike gateway KeySrv local-address 10.18.101.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway KeySrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group-vpn-external-interface
  ge-0/0/1.0
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 recovery-probe
set security ipsec-policy from-zone LAN to-zone WAN ipsec-group-vpn GROUP_ID-0001
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_LAN
user@host# set ge-0/0/0 unit 0 family inet address 172.16.101.1/24
user@host# set ge-0/0/1 unit 0 description To_KeySrv
user@host# set ge-0/0/1 unit 0 family inet address 10.18.101.1/24

[edit security zones security-zone LAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0

[edit security]
user@host# set address-book global address 172.16.0.0/12 172.16.0.0/12

[edit security zones security-zone WAN]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/1.0

[edit security policies from-zone LAN to-zone WAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set then log session-init

[edit security policies from-zone WAN to-zone LAN
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set then log session-init

[edit security policies]
user@host# set global policy 1000 match source-address any
user@host# set global policy 1000 match destination-address any
user@host# set global policy 1000 match application any
user@host# set global policy 1000 match from-zone any
user@host# set global policy 1000 match to-zone any
user@host# set global policy 1000 match then reject
user@host# set global policy 1000 match then log session-init
user@host# set global policy 1000 match then count
user@host# set default-policy deny-all
```

2. Configure the static routes.

```
[edit routing-options]
user@host# set static route 10.18.102.0/24 next-hop 10.18.101.254
user@host# set static route 10.18.103.0/24 next-hop 10.18.101.254
user@host# set static route 10.18.104.0/24 next-hop 10.18.101.254
user@host# set static route 172.16.101.0/24 next-hop 10.18.101.254
user@host# set static route 172.16.102.0/24 next-hop 10.18.101.254
user@host# set static route 172.16.103.0/24 next-hop 10.18.101.254
user@host# set static route 172.16.104.0/24 next-hop 10.18.101.254
user@host# set static route 10.10.100.0/24 next-hop 10.18.101.254
```

3. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set authentication-algorithm sha-256
user@host# set dh-group group14
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy KeySrv ]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit security group-vpn member ike gateway KeySrv]
user@host# set ike-policy KeySrv
user@host# set server-address 10.10.100.1
user@host# set local-address 10.18.101.1
```

4. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway KeySrv
user@host# set group-vpn-external-interface ge-0/0/1.0
user@host# set group 1
user@host# set recovery-probe
```

5. Configure the IPsec policy.

```
[edit security ipsec-policy from-zone LAN to-zone WAN]
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_LAN;
    family inet {
      address 172.16.101.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_KeySrv;
    family inet {
      address 10.18.101.1/24;
    }
  }
}
[edit]
user@host# show routing-options
static {
  route 10.18.102.0/24 next-hop 10.18.101.254;
  route 10.18.103.0/24 next-hop 10.18.101.254;
  route 10.18.104.0/24 next-hop 10.18.101.254;
  route 172.16.101.0/24 next-hop 10.18.101.254;
  route 172.16.102.0/24 next-hop 10.18.101.254;
  route 172.16.103.0/24 next-hop 10.18.101.254;
  route 172.16.104.0/24 next-hop 10.18.101.254;
  route 10.10.100.0/24 next-hop 10.18.101.254;
}
[edit]
user@host# show security
address-book {
  global {
```


Copyright © 2016, Juniper Networks, Inc.

```
        address 172.16.0.0/12 172.16.0.0/12;
    }
}
group-vpn {
    member {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                dh-group group14;
                authentication-algorithm sha-256;
                encryption-algorithm aes-256-cbc;
            }
            policy KeySrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            gateway KeySrv {
                ike-policy KeySrv;
                server-address 10.10.100.1;
                local-address 10.18.101.1;
            }
        }
        ipsec {
            vpn GROUP_ID-0001 {
                ike-gateway KeySrv;
                group-vpn-external-interface ge-0/0/1.0;
                group 1;
                recovery-probe;
            }
        }
    }
}
ipsec-policy {
    from-zone LAN to-zone WAN {
        ipsec-group-vpn GROUP_ID-0001;
    }
}
policies {
    from-zone LAN to-zone WAN {
        policy 1 {
            match {
                source-address 172.16.0.0/12;
                destination-address 172.16.0.0/12;
                application any;
            }
            then {
                permit;
                log {
                    session-init;
                }
            }
        }
    }
    from-zone WAN to-zone LAN {
        policy 1 {
```

Copyright © 2016, Juniper Networks, Inc.

```
        match {
          source-address 172.16.0.0/12;
          destination-address 172.16.0.0/12;
          application any;
        }
        then {
          permit;
          log {
            session-init;
          }
        }
      }
    }
    global {
      policy 1000 {
        match {
          source-address any;
          destination-address any;
          application any;
          from-zone any;
          to-zone any;
        }
        then {
          reject;
          log {
            session-init;
          }
          count;
        }
      }
    }
    default-policy {
      deny-all;
    }
  }
  zones {
    security-zone LAN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
      interfaces {
        ge-0/0/0.0;
      }
    }
    security-zone WAN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
        interfaces {
            ge-0/0/1.0;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Group Member GM-0002 (SRX Series Device or vSRX Instance)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_LAN
set interfaces ge-0/0/0 unit 0 family inet address 172.16.102.1/24
set interfaces ge-0/0/1 unit 0 description To_KeySrv
set interfaces ge-0/0/1 unit 0 family inet address 10.18.102.1/24
set security zones security-zone LAN host-inbound-traffic system-services ike
set security zones security-zone LAN host-inbound-traffic system-services ssh
set security zones security-zone LAN host-inbound-traffic system-services ping
set security zones security-zone LAN interfaces ge-0/0/0.0
set security zones security-zone WAN host-inbound-traffic system-services ike
set security zones security-zone WAN host-inbound-traffic system-services ssh
set security zones security-zone WAN host-inbound-traffic system-services ping
set security zones security-zone WAN interfaces ge-0/0/1.0
set security address-book global address 172.16.0.0/12 172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match application any
set security policies from-zone LAN to-zone WAN policy 1 then permit
set security policies from-zone LAN to-zone WAN policy 1 then log session-init
set security policies from-zone WAN to-zone LAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match application any
set security policies from-zone WAN to-zone LAN policy 1 then permit
set security policies from-zone WAN to-zone LAN policy 1 then log session-init
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then reject
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set routing-options static route 10.18.101.0/24 next-hop 10.18.102.254
set routing-options static route 10.18.103.0/24 next-hop 10.18.102.254
set routing-options static route 10.18.104.0/24 next-hop 10.18.102.254
set routing-options static route 172.16.101.0/24 next-hop 10.18.102.254
```

Copyright © 2016, Juniper Networks, Inc.

```
set routing-options static route 172.16.102.0/24 next-hop 10.18.102.254
set routing-options static route 172.16.103.0/24 next-hop 10.18.102.254
set routing-options static route 172.16.104.0/24 next-hop 10.18.102.254
set routing-options static route 10.10.100.0/24 next-hop 10.18.102.254
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy KeySrv mode main
set security group-vpn member ike policy KeySrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy KeySrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway KeySrv ike-policy KeySrv
set security group-vpn member ike gateway KeySrv server-address 10.10.100.1
set security group-vpn member ike gateway KeySrv local-address 10.18.102.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway KeySrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group-vpn-external-interface
  ge-0/0/1.0
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 recovery-probe
set security ipsec-policy from-zone LAN to-zone WAN ipsec-group-vpn GROUP_ID-0001
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_LAN
user@host# set ge-0/0/0 unit 0 family inet address 172.16.102.1/24
user@host# set ge-0/0/1 unit 0 description To_KeySrv
user@host# set ge-0/0/1 unit 0 family inet address 10.18.101.1/24

[edit security zones security-zone LAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0

[edit security zones security-zone WAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/1.0

[edit security]
user@host# set address-book global address 172.16.0.0/12 172.16.0.0/12
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies from-zone LAN to-zone WAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set then log session-init

[edit security policies from-zone WAN to-zone LAN
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set then log session-init

[edit security policies]
user@host# set global policy 1000 match source-address any
user@host# set global policy 1000 match destination-address any
user@host# set global policy 1000 match application any
user@host# set global policy 1000 match from-zone any
user@host# set global policy 1000 match to-zone any
user@host# set global policy 1000 match then reject
user@host# set global policy 1000 match then log session-init
user@host# set global policy 1000 match then count
user@host# set default-policy deny-all
```

2. Configure the static routes.

```
[edit routing-options]
user@host# set static route 10.18.101.0/24 next-hop 10.18.102.254
user@host# set static route 10.18.103.0/24 next-hop 10.18.102.254
user@host# set static route 10.18.104.0/24 next-hop 10.18.102.254
user@host# set static route 172.16.101.0/24 next-hop 10.18.102.254
user@host# set static route 172.16.102.0/24 next-hop 10.18.102.254
user@host# set static route 172.16.103.0/24 next-hop 10.18.102.254
user@host# set static route 172.16.104.0/24 next-hop 10.18.102.254
user@host# set static route 10.10.100.0/24 next-hop 10.18.102.254
```

3. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set authentication-algorithm sha-256
user@host# set dh-group group14
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy KeySrv ]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway KeySrv]
user@host# set ike-policy KeySrv
user@host# set server-address 10.10.100.1
user@host# set local-address 10.18.102.1
```

4. Configure the IPsec SA.

Copyright © 2016, Juniper Networks, Inc.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway KeySrv
user@host# set group-vpn-external-interface ge-0/0/1.0
user@host# set group 1
user@host# set recovery-probe
```

5. Configure the IPsec policy.

```
[edit security ipsec-policy from-zone LAN to-zone WAN]
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_LAN;
    family inet {
      address 172.16.102.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_KeySrv;
    family inet {
      address 10.18.102.1/24;
    }
  }
}
[edit]
user@host# show routing-options
static {
  route 10.18.101.0/24 next-hop 10.18.102.254;
  route 10.18.103.0/24 next-hop 10.18.102.254;
  route 10.18.104.0/24 next-hop 10.18.102.254;
  route 172.16.101.0/24 next-hop 10.18.102.254;
  route 172.16.102.0/24 next-hop 10.18.102.254;
  route 172.16.103.0/24 next-hop 10.18.102.254;
  route 172.16.104.0/24 next-hop 10.18.102.254;
  route 10.10.100.0/24 next-hop 10.18.102.254;
}
[edit]
user@host# show security
address-book {
  global {
    address 172.16.0.0/12 172.16.0.0/12;
  }
}
group-vpn {
  member {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
```


Copyright © 2016, Juniper Networks, Inc.

```
            authentication-method pre-shared-keys;
            dh-group group14;
            authentication-algorithm sha-256;
            encryption-algorithm aes-256-cbc;
          }
          policy KeySrv {
            mode main;
            proposals PSK-SHA256-DH14-AES256;
            pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
          }
          gateway KeySrv {
            ike-policy KeySrv;
            server-address 10.10.100.1;
            local-address 10.18.102.1;
          }
        }
        ipsec {
          vpn GROUP_ID-0001 {
            ike-gateway KeySrv;
            group-vpn-external-interface ge-0/0/1.0;
            group 1;
            recovery-probe;
          }
        }
      }
    }
    policies {
      from-zone LAN to-zone WAN {
        policy 1 {
          match {
            source-address 172.16.0.0/12;
            destination-address 172.16.0.0/12;
            application any;
          }
          then {
            permit;
            log {
              session-init;
            }
          }
        }
      }
      from-zone WAN to-zone LAN {
        policy 1 {
          match {
            source-address 172.16.0.0/12;
            destination-address 172.16.0.0/12;
            application any;
          }
          then {
            permit;
            log {
              session-init;
            }
          }
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
    global {
      policy 1000 {
        match {
          source-address any;
          destination-address any;
          application any;
          from-zone any;
          to-zone any;
        }
        then {
          reject;
          log {
            session-init;
          }
          count;
        }
      }
    }
    default-policy {
      deny-all;
    }
  }
  zones {
    security-zone LAN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
      interfaces {
        ge-0/0/0.0;
      }
    }
    security-zone WAN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
      interfaces {
        ge-0/0/1.0;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring Group Member GM-0003 (MX Series Device)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet service output service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet address 10.18.103.1/24
set interfaces xe-0/0/2 unit 0 family inet address 172.16.103.1/24
set interfaces ms-0/2/0 unit 0 family inet
set routing-options static route 10.18.101.0/24 next-hop 10.18.103.254
set routing-options static route 10.18.102.0/24 next-hop 10.18.103.254
set routing-options static route 10.18.104.0/24 next-hop 10.18.103.254
set routing-options static route 172.16.101.0/24 next-hop 10.18.103.254
set routing-options static route 172.16.102.0/24 next-hop 10.18.103.254
set routing-options static route 172.16.103.0/24 next-hop 10.18.103.254
set routing-options static route 172.16.104.0/24 next-hop 10.18.103.254
set routing-options static route 10.10.100.0/24 next-hop 10.18.103.254
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy KeySrv mode main
set security group-vpn member ike policy KeySrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy KeySrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway KeySrv ike-policy KeySrv
set security group-vpn member ike gateway KeySrv server-address 10.10.100.1
set security group-vpn member ike gateway KeySrv local-address 10.18.103.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway KeySrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 match-direction output
set security group-vpn member ipsec vpn GROUP_ID-0001 tunnel-mtu 1400
set security group-vpn member ipsec vpn GROUP_ID-0001 df-bit clear
set services service-set GROUP_ID-0001 interface-service service-interface ms-0/2/0.0
set services service-set GROUP_ID-0001 ipsec-group-vpn GROUP_ID-0001
set firewall family inet service-filter GroupVPN-KS term inbound-ks from
  destination-address 10.10.100.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.10.100.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.10.100.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  source-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  destination-address 172.16.0.0/12
```

Copyright © 2016, Juniper Networks, Inc.

```
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 then service
```

Step-by-Step Procedure

To configure the Group VPNv2 member:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet service output service-set
  GROUP_ID-0001 service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet address 10.18.103.1/24
user@host# set xe-0/0/2 unit 0 family inet address 172.16.103.1/24
user@host# set ms-0/2/0 unit 0 family inet
```

2. Configure routing.

```
[edit routing-options]
user@host# set static route 10.18.101.0/24 next-hop 10.18.103.254
user@host# set static route 10.18.102.0/24 next-hop 10.18.103.254
user@host# set static route 10.18.104.0/24 next-hop 10.18.103.254
user@host# set static route 172.16.101.0/24 next-hop 10.18.103.254
user@host# set static route 172.16.102.0/24 next-hop 10.18.103.254
user@host# set static route 172.16.103.0/24 next-hop 10.18.103.254
user@host# set static route 172.16.104.0/24 next-hop 10.18.103.254
user@host# set static route 10.10.100.0/24 next-hop 10.18.103.254
```

3. Configure IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256 ]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy KeySrv ]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway KeySrv]
user@host# set ike-policy KeySrv
user@host# set server-address 10.10.100.1
user@host# set local-address 10.18.103.1
```

4. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway KeySrv
user@host# set group 1
user@host# set match-direction output
user@host# set tunnel-mtu 1400
user@host# set df-bit clear
```

5. Configure the service filter.

```
[edit firewall family inet service-filter GroupVPN-KS]
user@host# set term inbound-ks from destination-address 10.10.100.1/32
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set term inbound-ks from source-address 10.10.100.1/32
user@host# set term inbound-ks then skip
user@host# set term outbound-ks from destination-address 10.10.100.1/32
user@host# set term outbound-ks then skip
user@host# set term GROUP_ID-0001 from source-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 from destination-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 then service
```

6. Configure the service set.

```
[edit services service-set GROUP_ID-0001]
user@host# set interface-service service-interface ms-0/2/0.0
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security**, **show services**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
xe-0/0/1 {
  unit 0 {
    family inet {
      service {
        input {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
        output {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
      }
      address 10.18.103.1/24;
    }
  }
}
xe-0/0/2 {
  unit 0 {
    family inet {
      address 172.16.103.1/24;
    }
  }
}
ms-0/2/0 {
  unit 0 {
    family inet;
  }
}
[edit]
user@host# show routing-options
static {
  route 10.18.101.0/24 next-hop 10.18.103.254;
  route 10.18.102.0/24 next-hop 10.18.103.254;
  route 10.18.104.0/24 next-hop 10.18.103.254;
  route 172.16.101.0/24 next-hop 10.18.103.254;
  route 172.16.102.0/24 next-hop 10.18.103.254;
```

Copyright © 2016, Juniper Networks, Inc.

```
    route 172.16.103.0/24 next-hop 10.18.103.254;
    route 172.16.104.0/24 next-hop 10.18.103.254;
}
[edit]
user@host# show security
group-vpn {
    member {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                dh-group group14;
                authentication-algorithm sha-256;
                encryption-algorithm aes-256-cbc;
            }
            policy KeySrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            gateway KeySrv {
                ike-policy KeySrv;
                local-address 10.18.103.1;
                server-address 10.10.101.1;
            }
        }
        ipsec {
            vpn GROUP_ID-0001 {
                ike-gateway KeySrv
                group 1;
                match-direction output;
                tunnel-mtu 1400;
                df-bit clear;
            }
        }
    }
}
[edit]
user@host# show services
service-set GROUP_ID-0001 {
    interface-service {
        service-interface ms-0/2/0.0;
    }
    ipsec-group-vpn GROUP_ID-0001;
}
[edit]
user@host# show firewall
family inet {
    service-filter GroupVPN-KS {
        term inbound-ks {
            from {
                destination-address {
                    10.10.100.1/32;
                }
                source-address {
                    10.10.100.1/32;
                }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        then skip;
      }
      term outbound-ks {
        from {
          destination-address {
            10.10.100.1/32;
          }
        }
        then skip;
      }
      term GROUP_ID-0001 {
        from {
          source-address {
            172.16.0.0/12;
          }
          destination-address {
            172.16.0.0/12;
          }
        }
        then service;
      }
    }
  }
```

Configuring Group Member GM-0004 (MX Series Device)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet service output service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet address 10.18.104.1/24
set interfaces xe-0/0/2 unit 0 family inet address 172.16.104.1/24
set interfaces ms-0/2/0 unit 0 family inet
set routing-options static route 10.18.101.0/24 next-hop 10.18.104.254
set routing-options static route 10.18.102.0/24 next-hop 10.18.104.254
set routing-options static route 10.18.103.0/24 next-hop 10.18.104.254
set routing-options static route 172.16.101.0/24 next-hop 10.18.104.254
set routing-options static route 172.16.102.0/24 next-hop 10.18.104.254
set routing-options static route 172.16.103.0/24 next-hop 10.18.104.254
set routing-options static route 172.16.104.0/24 next-hop 10.18.104.254
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy SubSrv mode main
```

Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn member ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway SubSrv ike-policy SubSrv
set security group-vpn member ike gateway SubSrv server-address 10.17.101.1
set security group-vpn member ike gateway SubSrv server-address 10.17.102.1
set security group-vpn member ike gateway SubSrv server-address 10.17.103.1
set security group-vpn member ike gateway SubSrv server-address 10.17.104.1
set security group-vpn member ike gateway SubSrv local-address 10.18.104.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway SubSrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 match-direction output
set security group-vpn member ipsec vpn GROUP_ID-0001 tunnel-mtu 1400
set security group-vpn member ipsec vpn GROUP_ID-0001 df-bit clear
set services service-set GROUP_ID-0001 interface-service service-interface ms-0/2/0.0
set services service-set GROUP_ID-0001 ipsec-group-vpn GROUP_ID-0001
set firewall family inet service-filter GroupVPN-KS term inbound-ks from
  destination-address 10.10.100.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.10.100.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.101.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.102.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.103.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.104.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  source-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  destination-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 then service
```

Step-by-Step Procedure

To configure the Group VPNv2 member:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet service output service-set
  GROUP_ID-0001 service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet address 10.18.104.1/24
user@host# set xe-0/0/2 unit 0 family inet address 172.16.104.1/24
user@host# set ms-0/2/0 unit 0 family inet
```

2. Configure routing.

```
[edit routing-options]
user@host# set static route 10.18.101.0/24 next-hop 10.18.104.254
user@host# set static route 10.18.102.0/24 next-hop 10.18.104.254
user@host# set static route 10.18.103.0/24 next-hop 10.18.104.254
user@host# set static route 172.16.101.0/24 next-hop 10.18.104.254
user@host# set static route 172.16.102.0/24 next-hop 10.18.104.254
user@host# set static route 172.16.103.0/24 next-hop 10.18.104.254
user@host# set static route 172.16.104.0/24 next-hop 10.18.104.254
```


Copyright © 2016, Juniper Networks, Inc.

3. Configure IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256 ]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy KeySrv ]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway KeySrv]
user@host# set ike-policy KeySrv
user@host# set server-address 10.10.100.1
user@host# set local-address 10.18.104.1
```

4. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway KeySrv
user@host# set group 1
user@host# set match-direction output
user@host# set tunnel-mtu 1400
user@host# set df-bit clear
```

5. Configure the service filter.

```
[edit firewall family inet service-filter GroupVPN-KS]
user@host# set term inbound-ks from destination-address 10.10.101.1/32
user@host# set term inbound-ks from source-address 10.10.101.1/32
user@host# set term inbound-ks then skip
user@host# set term outbound-ks from destination-address 10.17.101.1/32
user@host# set term outbound-ks from destination-address 10.17.102.1/32
user@host# set term outbound-ks from destination-address 10.17.103.1/32
user@host# set term outbound-ks from destination-address 10.17.104.1/32
user@host# set term outbound-ks then skip
user@host# set term GROUP_ID-0001 from source-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 from destination-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 then service
```

6. Configure the service set.

```
[edit services service-set GROUP_ID-0001]
user@host# set interface-service service-interface ms-0/2/0.0
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show security**, **show services**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
xe-0/0/1 {
   unit 0 {
```

Copyright © 2016, Juniper Networks, Inc.

```
      family inet {
        service {
          input {
            service-set GROUP_ID-0001 service-filter GroupVPN-KS;
          }
          output {
            service-set GROUP_ID-0001 service-filter GroupVPN-KS;
          }
        }
        address 10.18.104.1/24;
      }
    }
}
xe-0/0/2 {
    unit 0 {
      family inet {
        address 172.16.104.1/24;
      }
    }
}
ms-0/2/0 {
    unit 0 {
      family inet;
    }
}
[edit]
user@host# show routing-options
static {
    route 10.18.101.0/24 next-hop 10.18.104.254;
    route 10.18.102.0/24 next-hop 10.18.104.254;
    route 10.18.103.0/24 next-hop 10.18.104.254;
    route 172.16.101.0/24 next-hop 10.18.104.254;
    route 172.16.102.0/24 next-hop 10.18.104.254;
    route 172.16.103.0/24 next-hop 10.18.104.254;
    route 172.16.104.0/24 next-hop 10.18.104.254;
}
[edit]
user@host# show security
group-vpn {
    member {
      ike {
        proposal PSK-SHA256-DH14-AES256 {
          authentication-method pre-shared-keys;
          dh-group group14;
          authentication-algorithm sha-256;
          encryption-algorithm aes-256-cbc;
        }
        policy KeySrv {
          mode main;
          proposals PSK-SHA256-DH14-AES256;
          pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
        }
        gateway KeySrv {
          ike-policy KeySrv;
          local-address 10.18.104.1;
          server-address 10.17.101.1;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
      ipsec {
        vpn GROUP_ID-0001 {
          ike-gateway KeySrv
          group 1;
          match-direction output;
          tunnel-mtu 1400;
          df-bit clear;
        }
      }
    }
}
[edit]
user@host# show services
service-set GROUP_ID-0001 {
  interface-service {
    service-interface ms-0/2/0.0;
  }
  ipsec-group-vpn GROUP_ID-0001;
}
[edit]
user@host# show firewall
family inet {
  service-filter GroupVPN-KS {
    term inbound-ks {
      from {
        destination-address {
          10.10.100.1/32;
        }
        source-address {
          10.10.100.1/32;
        }
      }
      then skip;
    }
    term outbound-ks {
      from {
        destination-address {
          10.17.101.1/32;
          10.17.102.1/32;
          10.17.103.1/32;
          10.17.104.1/32;
        }
      }
      then skip;
    }
    term GROUP_ID-0001 {
      from {
        source-address {
          172.16.0.0/12;
        }
        destination-address {
          172.16.0.0/12;
        }
      }
```

```
        then service;
      }
    }
  }
```

Verification

Confirm that the configuration is working properly.

- Verifying Group Member Registration on page 6686
- Verifying That Group Keys Are Distributed on page 6687
- Verifying Group VPN SAs on the Group Server on page 6687
- Verifying Group VPN SAs on Group Members on page 6687
- Verifying IPsec SAs on the Group Server on page 6688
- Verifying IPsec SAs on the Group Members on page 6689
- Verifying Group Policies (SRX or vSRX Group Members Only) on page 6691

Verifying Group Member Registration

Purpose Verify that group members are registered on the server.

Action From operational mode, enter the **show security group-vpn server registered-members** and **show security group-vpn server registered-members detail** commands on the server.

```
user@host> show security group-vpn server registered-members
Group: GROUP_ID-0001, Group Id: 1
  Total number of registered members: 2
  Member Gateway                   Member IP       Last Update              Vsys

  GM-0001                          10.18.101.1     Thu Nov 19 2015 16:31:09 root

  GM-0003                          10.18.103.1     Thu Nov 19 2015 16:29:47 root
```

```
user@host> show security group-vpn server registered-members detail
GGroup: GROUP_ID-0001, Group Id: 1
  Total number of registered members: 2

  Member gateway: GM-0001, Member IP: 10.18.101.1, Vsys: root
  Last Update: Thu Nov 19 2015 16:31:09
  Stats:
      Pull Succeeded                : 2
      Pull Failed                   : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0

  Member gateway: GM-0003, Member IP: 10.18.103.1, Vsys: root
  Last Update: Thu Nov 19 2015 16:29:47
  Stats:
      Pull Succeeded                : 1
      Pull Failed                   : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```


Copyright © 2016, Juniper Networks, Inc.

Verifying That Group Keys Are Distributed

Purpose Verify that group keys are distributed to members.

Action From operational mode, enter the **show security group-vpn server statistics** command on the group server.

```
user@host> show security group-vpn server statistics
Group: GROUP_ID-0001, Group Id: 1
  Stats:
      Pull Succeeded                : 4
      Pull Failed                   : 0
      Pull Exceed Member Threshold  : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```

Verifying Group VPN SAs on the Group Server

Purpose Verify Group VPN SAs on the group server.

Action From operational mode, enter the **show security group-vpn server kek security-associations** and **show security group-vpn server kek security-associations detail** commands on the group server.

```
user@host> show security group-vpn server kek security-associations
Index   Life:sec  Initiator cookie  Responder cookie  GroupId
738879  1206      a471513492db1e13  24045792a4b3dd64  1
```

```
user@host> show security group-vpn server kek security-associations detail
Index 738879, Group Name: GROUP_ID-0001, Group Id: 1
Initiator cookie: a471513492db1e13, Responder cookie: 24045792a4b3dd64
Authentication method: RSA
Lifetime: Expires in 1204 seconds, Activated
Rekey in 694 seconds
  Algorithms:
   Sig-hash             : sha256
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Server Member Communication: Unicast
  Retransmission Period: 10, Number of Retransmissions: 2
  Group Key Push sequence number: 0

PUSH negotiations in progress: 0
```

Verifying Group VPN SAs on Group Members

Purpose Verify Group VPN SAs on the group members.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security group-vpn member kek security-associations** and **show security group-vpn member kek security-associations detail** commands on the SRX or vSRX group member.

```
user@host> show security group-vpn member kek security-associations
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
5455810 10.10.100.1     1093      a471513492db1e13  24045792a4b3dd64  1

user@host> show security group-vpn member kek security-associations detail
  Index 5455810, Group Id: 1
  Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.101.1, GDOI Server: 10.10.100.1
  Initiator cookie: a471513492db1e13, Responder cookie: 24045792a4b3dd64
  Lifetime: Expires in 1090 seconds
  Group Key Push Sequence number: 0

  Algorithms:
   Sig-hash             : hmac-sha256-128
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Stats:
      Push received           :   0
      Delete received         :   0
```

From operational mode, enter the **show security group-vpn member kek security-associations** and **show security group-vpn member kek security-associations detail** commands on the MX Series group member.

```
user@host> show security group-vpn member kek security-associations
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
488598  10.10.100.1     963       a471513492db1e13  24045792a4b3dd64  1

user@host> show security group-vpn member kek security-associations detail
  Index 488598, Group Id: 1
  Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.103.1, GDOI Server: 10.10.100.1
  Initiator cookie: a471513492db1e13, Responder cookie: 24045792a4b3dd64
  Lifetime: Expires in 961 seconds
  Group Key Push Sequence number: 0

  Algorithms:
   Sig-hash             : hmac-sha256-128
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Stats:
      Push received           :   0
      Delete received         :   0
```

Verifying IPsec SAs on the Group Server

Purpose Verify IPsec SAs on the group server.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security group-vpn server ipsec security-associations** and **show security group-vpn server ipsec security-associations detail** commands on the group server.

```
user@host> show security group-vpn server ipsec security-associations
Group: GROUP_ID-0001, Group Id: 1
  Total IPsec SAs: 1
  IPsec SA          Algorithm        SPI              Lifetime
  GROUP_ID-0001     ESP:aes-256/sha256 1c548e4e       1156

user@host> show security group-vpn server ipsec security-associations detail
Group: GROUP_ID-0001, Group Id: 1
Total IPsec SAs: 1
  IPsec SA: GROUP_ID-0001
    Protocol: ESP, Authentication: sha256, Encryption: aes-256
    Anti-replay: D3P enabled
    SPI: 1c548e4e
    Lifetime: Expires in 1152 seconds, Activated
    Rekey in 642 seconds
    Policy Name: 1
      Source: 172.16.0.0/12
      Destination: 172.16.0.0/12
      Source Port: 0
      Destination Port: 0
      Protocol: 0
```

Verifying IPsec SAs on the Group Members

Purpose Verify IPsec SAs on the group members.

Action From operational mode, enter the **show security group-vpn member ipsec security-associations** and **show security group-vpn member ipsec security-associations detail** commands on the SRX or vSRX group member.

```
user@host> show security group-vpn member ipsec security-associations
  Total active tunnels: 1
  ID    Server           Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>49152 10.10.100.1    848   ESP:aes-256/sha256-128 1c548e4e 1073/ unlim 1 root

user@host> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.101.1, GDOI Server: 10.10.100.1
  Group Id: 1
  Routing Instance: default
  Recovery Probe: Enabled
  DF-bit: clear
  Stats:
      Pull Succeeded             :   4
      Pull Failed                :   3
      Pull Timeout               :   3
      Pull Aborted               :   0
      Push Succeeded             :   6
      Push Failed                :   0
      Server Failover            :   0
      Delete Received            :   0
      Exceed Maximum Keys(4)     :   0
      Exceed Maximum Policies(10):   0
      Unsupported Algo           :   0
  Flags:
```

Copyright © 2016, Juniper Networks, Inc.

```
    Rekey Needed:   no

  List of policies received from server:
  Tunnel-id: 49152
    Source IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)
    Destination IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)

    Direction: bi-directional, SPI: 1c548e4e
    Protocol: ESP, Authentication: sha256-128, Encryption: aes-256
    Hard lifetime: Expires in 1070 seconds, Activated
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 931 seconds
    Mode: Tunnel, Type: Group VPN, State: installed
    Anti-replay service: D3P enabled
```

From operational mode, enter the **show security group-vpn member ipsec security-associations** and **show security group-vpn member ipsec security-associations detail** commands on the MX Series group member.

```
user@host> show security group-vpn member ipsec security-associations
  Total active tunnels: 1
  ID    Server           Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>10001 10.10.100.1    848   ESP:aes-256/sha256-128 1c548e4e 947/ unlim 1 root

user@host> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.103.1, GDOI Server: 10.10.100.1
  Group Id: 1
  Rule Match Direction: output,  Tunnel-MTU: 1400
  Routing Instance: default
  DF-bit: clear
  Stats:
      Pull Succeeded            :   2
      Pull Failed               :   0
      Pull Timeout              :   1
      Pull Aborted              :   0
      Push Succeeded            :   2
      Push Failed               :   0
      Server Failover           :   0
      Delete Received           :   0
      Exceed Maximum Keys(4)    :   0
      Exceed Maximum Policies(1):   0
      Unsupported Algo          :   0
  Flags:
      Rekey Needed:   no

    List of policies received from server:
    Tunnel-id: 10001
      Source IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)
      Destination IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)

      Direction: bi-directional, SPI: 1c548e4e
      Protocol: ESP, Authentication: sha256-128, Encryption: aes-256
      Hard lifetime: Expires in 945 seconds, Activated
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 840 seconds
      Mode: Tunnel, Type: Group VPN, State: installed
      Anti-replay service: D3P enabled
```


Copyright © 2016, Juniper Networks, Inc.

Verifying Group Policies (SRX or vSRX Group Members Only)

Purpose Verify group policies on SRX or vSRX group members.

Action From operational mode, enter the **show security group-vpn member policy** command on the group member.

```
user@host> show security group-vpn member policy
Group VPN Name: GROUP_ID-0001, Group Id: 1
From-zone: LAN, To-zone: WAN
 Tunnel-id: 49152, Policy type: Secure
  Source      : IP <172.16.0.0 - 172.31.255.255>, Port <0 - 65535>, Protocol <0>

  Destination : IP <172.16.0.0 - 172.31.255.255>, Port <0 - 65535>, Protocol <0>


 Tunnel-id: 63488, Policy type: Fail-close
  Source      : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
  Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
```

Related Documentation

- Group VPNv2 Configuration Overview on page 6652
- Configuring Group VPNs in Group VPNv2 on Routing Device

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 276

Configuring Server-Member Communication

- Understanding Server-Member Communication on page 6693
- Understanding Group Key Operations on page 6694
- Example: Configuring Server-Member Communication for Unicast Rekey Messages on page 6695
- Understanding the Recovery Probe Process on page 6696
- Understanding Antireplay on page 6697

Understanding Server-Member Communication

Server-member communication allows the server to send GDOI **groupkey-push** (rekey) messages to members. If server-member communication is not configured for the group, members can send GDOI **groupkey-pull** messages to register and reregister with the server, but the server is not able to send **groupkey-push** messages to members.

Server-member communication is configured for the group by using the **server-member-communication** configuration statement at the [**edit security group-vpn server**] hierarchy. The following options can be defined:

- Authentication algorithm (sha-256 or sha-384) used to authenticate the member to the server. There is no default algorithm.
- Encryption algorithm used for communications between the server and member. You can specify aes-128-cbc, aes-192-cbc, or aes-256-cbc. There is no default algorithm.
- Unicast communication type for rekey messages sent to group members. See "Understanding Group Key Operations" on page 6694.
- Lifetime for the key encryption key (KEK). The default is 3600 seconds.
- Number of times the group server retransmits **groupkey-push** messages to a group member without a response (the default is 2 times) and the period of time between retransmissions (the default is 10 seconds).

If server-member communication for a group is not configured, the membership list displayed by the **show security group-vpn server registered-members** command shows group members who have registered with the server; members can be active or not. When

Copyright © 2016, Juniper Networks, Inc.

server-member communication for a group is configured, the group membership list is cleared. For unicast communication type, the **show security group-vpn server registered-members** command shows only active members.

Related Documentation

- Understanding Group Key Operations on page 6694
- Understanding Group Configuration on page 6655
- Example: Configuring Server-Member Communication for Unicast Rekey Messages on page 6695

Understanding Group Key Operations

This topic contains the following sections:

- Group Keys on page 6694
- Rekey Messages on page 6694
- Member Registration on page 6695

Group Keys

The group server maintains a database to track the relationship among VPN groups, group members, and group keys. There are two kinds of group keys that the server downloads to members:

- Key Encryption Key (KEK)—Used to encrypt SA rekey (GDOI **groupkey-push**) exchanges. One KEK is supported per group.
- Traffic Encryption Key (TEK)—Used to encrypt and decrypt IPsec data traffic between group members.

The key associated with an SA is accepted by a group member only if there is a matching policy configured on the member. An accepted key is installed for the group, whereas a rejected key is discarded.

Rekey Messages

If the group is configured for server-member communications (see "Understanding Server-Member Communication" on page 6693), the server sends SA and key refreshes to group members with rekey (GDOI **groupkey-push**) messages. Rekey messages are sent before SAs expire; this ensures that valid keys are available for encrypting traffic between group members.

The server also sends rekey messages to provide new keys to members when there is a change in group membership or the group SA has changed (for example, a group policy is added or deleted).

Server-member communications options must be configured on the server to allow the server to send rekey messages to group members.

The group server sends one copy of the unicast rekey message to each group member. Upon receipt of the rekey message, members must send an acknowledgment (ACK) to

Copyright © 2016, Juniper Networks, Inc.

the server. If the server does not receive an ACK from a member (including retransmission of rekey messages), the server considers the member to be inactive and removes it from the membership list. The server stops sending rekey messages to the member.

The **number-of-retransmission** and **retransmission-period** configuration statements for server-member communications control the resending of rekey messages by the server when no ACK is received from a member.

The interval at which the server sends rekey messages is based on the value of the **lifetime-seconds** configuration statement at the [**edit security group-vpn server group** ***group-name***] hierarchy. New keys are generated before the expiration of the KEK and TEK keys.

The **lifetime-seconds** for the KEK is configured as part of the server-member communications; the default is 3600 seconds. The **lifetime-seconds** for the TEK is configured for the IPsec proposal; the default is 3600 seconds.

Member Registration

If a group member does not receive a new SA key from the server before the current key expires, the member must reregister with the server and obtain updated keys with a GDOI **groupkey-pull** exchange.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding the GDOI Protocol on page 6650
- Understanding Group Servers and Members on page 6651
- Group VPNv2 Configuration Overview on page 6652

Example: Configuring Server-Member Communication for Unicast Rekey Messages

This example shows how to enable the server to send unicast rekey messages to group members to ensure that valid keys are available for encrypting traffic between group members.

- Requirements on page 6695
- Overview on page 6696
- Configuration on page 6696
- Verification on page 6696

Requirements

Before you begin:

- Configure the group server and members for IKE Phase 1 negotiation.
- Configure the group server and members for IPsec SA.
- Configure the group **g1** on the group server.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you specify the following server-member communication parameters for group **g1**:

- The server sends unicast rekey messages to group members.
- aes-128-cbc is used to encrypt traffic between the server and members.
- sha-256 is used for member authentication.

Default values are used for KEK lifetime and retransmissions.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure server-member communication:

1. Set the communications type.

   ```
   [edit security group-vpn server group g1 server-member-communication]
   user@host# set communications-type unicast
   ```

2. Set the encryption algorithm.

   ```
   [edit security group-vpn server group g1 server-member-communication]
   user@host# set encryption-algorithm aes-128-cbc
   ```

3. Set the member authentication.

   ```
   [edit security group-vpn server group g1 server-member-communication]
   user@host# set sig-hash-algorithm sha-256
   ```

Verification

To verify the configuration is working properly, enter the **show security group-vpn server group g1 server-member-communication** command.

Related Documentation

- Group VPNv2 Configuration Overview on page 6652
- Understanding Server-Member Communication on page 6693
- Understanding Group Key Operations on page 6694
- Understanding Group Configuration on page 6655

Understanding the Recovery Probe Process

Two situations could indicate that a group member is out of synchronization with the group server and other group members:

- The group member receives an Encapsulating Security Payload (ESP) packet with an unrecognized Security Parameter Index (SPI).

Copyright © 2016, Juniper Networks, Inc.

- There is outgoing IPsec traffic but no incoming IPsec traffic on the group member.

When either situation is detected, a recovery probe process can be triggered on the group member. The recovery probe process initiates GDOI **groupkey-pull** exchanges at specific intervals to update the member's SA from the group server. If there is a DoS attack of bad SPI packets or if the sender itself is out of synchronization, the out-of-synchronization indication on the group member might be a false alarm. To avoid overloading the system, the **groupkey-pull** initiation is retried at intervals of 10, 20, 40, 80, 160, and 320 seconds.

The recovery probe process is disabled by default. To enable the recovery probe process, configure **recovery-probe** at the [**edit security group-vpn member ipsec vpn** *vpn-name*] hierarchy level.

Related Documentation

- Understanding Group Servers and Members on page 6651
- Understanding the GDOI Protocol on page 6650

Understanding Antireplay

Antireplay is an IPsec feature that can detect when a packet is intercepted and then replayed by attackers. Antireplay is disabled by default for a group.

Each IPsec packet contains a timestamp. The group member checks whether the packet's timestamp falls within the configured **anti-replay-time-window** value. A packet is dropped if the timestamp exceeds the value.

Juniper recommends that NTP be configured on all Group VPNv2 devices.

NOTE: Group members that are running on vSRX instances on a host machine where the hypervisor is running under a heavy load may experience issues that can be corrected by reconfiguring the **anti-replay-time-window** value. If data that matches the IPsec policy on the group member is not being transferred, check the **show security group-vpn member ipsec statistics** output for D3P errors. Make sure that NTP is operating correctly. If there are errors, adjust the **anti-replay-time-window** value.

Related Documentation

- Understanding Group Servers and Members on page 6651
- Understanding Group Configuration on page 6655

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 277

Configuring Group VPNv2 Server Clusters

- Understanding Group VPNv2 Server Clusters on page 6699
- Understanding Group VPNv2 Server Cluster Limitations on page 6703
- Understanding Group VPNv2 Server Cluster Messages on page 6704
- Understanding Configuration Changes with Server Clusters on page 6706
- Migrating a Standalone Group Server to a Server Cluster on page 6708
- Example: Configuring a Group VPNv2 Server Cluster and Members on page 6709

Understanding Group VPNv2 Server Clusters

In the Group Domain of Interpretation (GDOI) protocol, the group controller/key server (GCKS) manages Group VPN security associations (SAs), and generates encryption keys and distributes them to group members. Group members encrypt traffic based on the group SAs and keys provided by the GCKS. If the GCKS fails, group members cannot register or obtain keys. A Group VPNv2 server cluster provides GCKS redundancy so there is no single point of failure for the entire group VPN network. Group VPNv2 server clusters can also provide load balancing, scaling, and link redundancy.

NOTE: All servers in a Group VPNv2 server cluster must be supported SRX Series devices or vSRX instances. Group VPNv2 server clusters are a Juniper Networks proprietary solution and have no interoperability with other vendor's GCKS.

- Root-Server and Sub-Servers on page 6699
- Group Member Registration with Server Clusters on page 6701
- Dead Peer Detection on page 6702
- Load Balancing on page 6702

Root-Server and Sub-Servers

A Group VPNv2 server cluster consists of one root-server with up to four connected sub-servers. All servers in the cluster share the same SA and encryption keys that are distributed to Group VPNv2 members. Servers in the cluster can be located at different sites, as shown in Figure 299.

Copyright © 2016, Juniper Networks, Inc.

Figure 299: Group VPNv2 Server Cluster

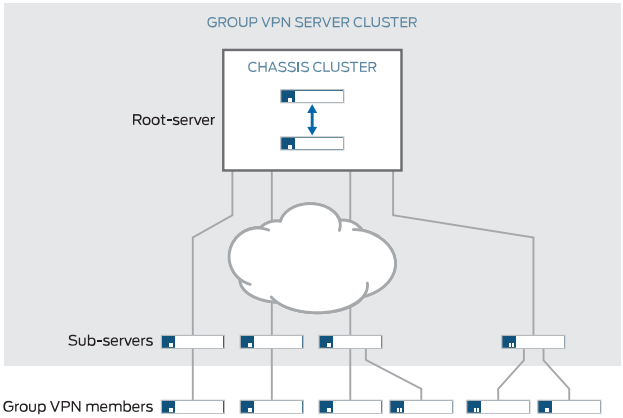

Messages between servers in the cluster are encrypted and authenticated by IKE SAs. The root-server is responsible for generating and distributing encryption keys to sub-servers; because of this responsibility, we recommend that the root-server be configured as a chassis cluster. Sub-servers are single devices and cannot be chassis clusters. Sub-servers must be able to connect to the root-server, although direct links between sub-servers are not necessary.

NOTE: If a sub-server loses its connection to the root-server, no further connection to the sub-server from group members are allowed and SAs are deleted. Therefore, we recommend that you use a different link to connect each sub-server to the root-server.

Group VPNv2 server clusters are configured with the **server-cluster** statements at the [**edit security group-vpn server** *group-name*] hierarchy level. The following values must be configured for each server in a cluster:

- The server role—Specify either **root-server** or **sub-server**. A given server can be part of multiple Group VPNv2 server clusters, but it must have the same server role in all clusters. A server cannot be configured with the root-server role in one group and the sub-server role in another group.

  NOTE: You must ensure that there is only one root-server at any time for a Group VPNv2 server cluster.

- IKE gateway—Specify the name of an IKE gateway configured at the [**edit security group-vpn server ike**] hierarchy level. For a root-server, the IKE gateway must be a sub-server in the cluster; up to four sub-servers can be specified. For sub-servers, the IKE gateway must be the root-server.

Copyright © 2016, Juniper Networks, Inc.

NOTE: The root-server and sub-servers must be configured with **dead-peer-detection always-send** and cannot be configured for a dynamic (unspecified) IP address. Group members are not configured with dead peer detection.

The Group VPNv2 configuration must be the same on each sub-server in a given group.

Each sub-server in the Group VPNv2 server cluster operates as a normal GCKS for registering and deleting members. Upon successful member registration, the registering server is responsible for sending updates to the member. For a given group, you can configure the maximum number of Group VPNv2 members that can be accepted by each sub-server; this number must be the same on all sub-servers in the cluster. A sub-server stops responding to registration requests by new members when it reaches the configured maximum number of Group VPNv2 members. See "Load Balancing" on page 6702.

Group Member Registration with Server Clusters

Group members can register with any server in the Group VPNv2 server cluster for a given group, however we recommend that members only connect to sub-servers and not the root-server. Up to four server addresses can be configured on each group member. The server addresses configured on group members can be different. In the example shown below, group member A is configured for sub-servers 1 through 4, while member B is configured for sub-servers 4 and 3:

	Group member A:	Group member B:
Server addresses:	Sub-server 1	Sub-server 4
	Sub-server 2	Sub-server 3
	Sub-server 3	
	Sub-server 4	

The order that the server addresses is configured on a member is important. A group member attempts to register with the first configured server. If registration with a configured server is not successful, the group member tries to register with the next configured server.

Each server in a Group VPNv2 server cluster operates as a normal GCKS for registering and deleting members. Upon successful registration, the registering server is responsible for sending updates to the member via **groupkey-push** exchanges. For a given group, you can configure the maximum number of group members that can be accepted by each server, however this number must be the same on all servers in the cluster for a given group. Upon reaching the configured maximum number of group members, a server stops responding to registration requests by new members. See "Load Balancing" on page 6702 for additional information.

Copyright © 2016, Juniper Networks, Inc.

Dead Peer Detection

To verify the availability of peer servers in a Group VPNv2 server cluster, each server in the cluster must be configured to send dead peer detection (DPD) requests regardless of whether there is outgoing IPsec traffic to the peer. This is configured with the **dead-peer-detection always-send** statement at the [**edit security group-vpn server ike gateway *gateway-name***] hierarchy level.

An active server in a Group VPNv2 server cluster sends DPD probes to the IKE gateway(s) configured in the server cluster. DPD is not configured for a group because multiple groups can share the same peer server IKE gateway configuration. When DPD detects that a server is down, the IKE SA with that server is deleted. All groups mark the server as inactive and DPD to the server is stopped.

NOTE: **DPD is not configured for the IKE gateway on group members.**

When DPD marks the root-server as inactive, the sub-servers stop responding to new group member requests however existing SAs for current group members remain active. An inactive sub-server does not send deletes to group members because the SAs could be still valid and group members can continue using existing SAs.

If an IKE SA expires while a peer server is still active, DPD triggers IKE SA negotiation. Because both root-servers and sub-servers can trigger IKE SAs through DPD, simultaneous negotiation might result in multiple IKE SAs. No impact on server-cluster functionality is expected in this case.

Load Balancing

Load balancing in the Group VPNv2 server cluster can be achieved by configuring the right **member-threshold** value for the group. When the number of members registered on a server exceeds the **member-threshold** value, subsequent member registration on that server is rejected. The member registration fails over to the next server configured on the group member until it reaches a server whose **member-threshold** is not yet reached.

There are two restrictions on configuring the **member-threshold**:

- For a given group, the same **member-threshold** value must be configured on the root-server and all sub-servers in a group server cluster. If the total number of members in the group exceeds the configured **member-threshold** value, then a **groupkey-pull** registration initiated by a new member is rejected (the server does not send a response).
- A server can support members in multiple groups. Each server has a maximum number of group members that it can support. If a server reaches the maximum number of members it can support, then a **groupkey-pull** registration initiated by a new member is rejected even if the **member-threshold** value of a specific group has not been reached.

There is no member synchronization among servers in the cluster. The root-server does not have information about the number of registered members on sub-servers. Each sub-server can only show its own registered members.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group Key Operations on page 6694
- Understanding Group VPNv2 Server Cluster Messages on page 6704
- Example: Configuring a Group VPNv2 Server Cluster and Members on page 6709

Understanding Group VPNv2 Server Cluster Limitations

Note the following caveats when configuring Group VPNv2 server clusters:

- Certificate authentication is not supported for server authentication; only preshared keys can be configured.
- There is no configuration synchronization between servers in the Group VPNv2 server cluster.
- When enabling a Group VPNv2 server cluster, configuration must be done on the root-server first and then on the sub-servers. Until the configuration is manually synchronized among the servers, traffic loss can be expected during the configuration change.
- In certain corner cases, the SAs on Group VPNv2 members can be out of sync. Group VPN members can synchronize SAs by getting a new key through a **groupkey-pull** exchange. You can manually clear SAs on a Group VPNv2 member with the **clear security group-vpn member ipsec security-associations** or **clear security group-vpn member group** commands to help speed recovery.
- The Group VPNv2 server cluster does not support ISSU.
- If the last **groupkey-pull** message is lost during a Group VPNv2 member's registration, a server might consider the member to be a registered member even though the member might fail over to the next server in the server cluster. In this case, the same member might appear to be registered on multiple servers. If the total member-threshold on all servers equals the total number of deployed members, subsequent group members might fail to register.

Note the following caveats for chassis cluster operations on the root-server:

- No statistics are preserved.
- No negotiation data or state is saved. If a root-server chassis cluster failover occurs during a **groupkey-pull** or **groupkey-push** negotiation, the negotiation is not restarted after the failover.
- If both chassis cluster nodes of a root-server go down during a rekey of an encryption key, some Group VPNv2 members might receive the new key while other members do not. Traffic might be impacted. Manually clearing SAs on a Group VPNv2 member with the **clear security group-vpn member ipsec security-associations** or **clear security**

Copyright © 2016, Juniper Networks, Inc.

group-vpn member group commands might help speed up recovery when the root-server becomes reachable.

- In a large-scale environment, RG0 failover on the root-server might take time. If the DPD interval and threshold on a sub-server are configured with small values, it can result in the sub-server marking the root-server as inactive during an RG0 failover. Traffic might be impacted. We recommend that you configure the IKE gateway for the sub-server with a DPD **interval** * **threshold** value larger than 150 seconds.

Related Documentation

- Understanding Group VPNv2 Server Clusters on page 6699

Understanding Group VPNv2 Server Cluster Messages

All messages between servers in a Group VPNv2 server cluster are encrypted and authenticated by an IKE security association (SA). Each sub-server initiates an IKE SA with the root-server; this IKE SA must be established before messages can be exchanged between the servers.

This section describes the messages exchanged between the root-server and sub-servers.

- Cluster Exchanges on page 6704
- Cluster-Init Exchanges on page 6704
- Cluster-Update Messages on page 6705

Cluster Exchanges

Figure 300 shows the basic messages exchanged between the Group VPNv2 server cluster and Group VPNv2 members.

Figure 300: Group VPNv2 Server Cluster Messages

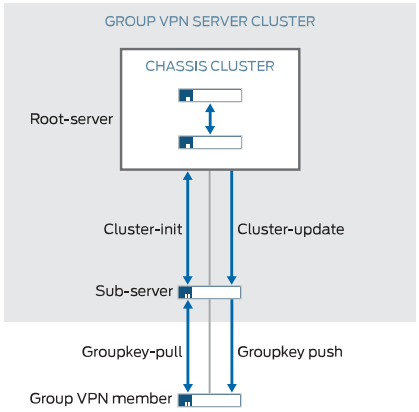

Cluster-Init Exchanges

A sub-server launches a cluster initialization (**cluster-init**) exchange with the root-server to obtain SA and encryption key information. The root-server responds by sending current SA information to the sub-server through the **cluster-init** exchange.

Copyright © 2016, Juniper Networks, Inc.

Sub-servers can then respond to registration requests from Group VPNv2 members through a **groupkey-pull** exchange. The **groupkey-pull** exchange allows a Group VPNv2 member to request SAs and keys shared by the group from a sub-server.

Sub-servers start a **cluster-init** exchange with the root-server when:

- The root-server is considered inactive. This is the initial assumed state of the root-server. If there is no IKE SA between the root-server and the sub-server, the sub-server initiates an IKE SA with the root-server. After a successful **cluster-init** exchange, the sub-server obtains information on SAs and marks the root-server as active.
- The soft lifetime of the SA has expired.
- A **cluster-update** message is received to delete all SAs.
- There are group configuration changes.

If the **cluster-init** exchange fails, the sub-server retries the exchange with the root-server every 5 seconds.

Cluster-Update Messages

The **groupkey-push** exchange is a single rekey message that allows a group controller/key server (GCKS) to send group SAs and keys to members before existing group SAs expire and to update group membership. Rekey messages are unsolicited messages sent from the GCKS to members

Upon generating new encryption keys for an SA, the root-server sends SA updates to all active sub-servers through a **cluster-update** message. After receiving a **cluster-update** from the root-server, the sub-server installs the new SA and sends the new SA information through a **groupkey-push** to its registered group members.

A **cluster-update** message sent from the root-server requires an acknowledgement from the sub-server. If there is no acknowledgement received from a sub-server, the root-server retransmits the **cluster-update** at the configured retransmission period (the default is 10 seconds). The root-server does not retransmit if dead peer detection (DPD) indicates that the sub-server is unavailable. If a sub-server fails to update SA information after receiving a **cluster-update**, it does not send an acknowledgement and the root-server retransmits the **cluster-update** message.

If the soft lifetime of an SA expires before a new SA is received from the root-server, the sub-server sends a **cluster-init** message to the root-server to get all SAs and does not send a **groupkey-push** message to its members until it has a new update. If the hard lifetime of an SA expires on the sub-server before it receives a new SA, the sub-server marks the root-server inactive, deletes all registered group members, and continues to send **cluster-init** messages to the root-server.

A **cluster-update** message can be sent to delete an SA or a group member; this can be the result of a **clear** command or a configuration change. If a sub-server receives a **cluster-update** message to delete an SA, it sends a **groupkey-push** delete message to its group members and deletes the corresponding SA. If all SAs for a group are deleted, the sub-server initiates a **cluster-init** exchange with the root-server. If all registered members are deleted, the sub-server deletes all locally registered members.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding Group VPNv2 Server Clusters on page 6699

Understanding Configuration Changes with Server Clusters

Group VPNv2 server clusters behave differently from standalone Group VPNv2 servers when there are configuration changes that result in new encryption keys and changes to security associations (SAs). The root-server sends SA updates or deletions to sub-servers through **cluster-update** messages. The sub-servers then send **groupkey-push** messages to members. Sub-servers cannot send delete messages to group members without first receiving delete messages from the root-server.

NOTE: All configuration changes must be made on the root-server first and then on sub-servers to ensure that group members receive updates or deletions as expected. Until configuration is synchronized between the servers in the Group VPNv2 server cluster, traffic loss can be expected.

Table 569 describes the effects of various configuration changes on Group VPNv2 servers.

Table 569: Effects of Configuration Changes on Group VPNv2 Servers

Configuration Change	Standalone Group VPNv2 Server Action	Group VPNv2 Server Cluster Action	
		Root-server	Sub-server
Change IKE proposal, policy, or gateway	Delete the IKE SA for the affected gateway. For IKE proposal, policy, or gateway deletions, delete the registered members for the affected gateway.		
Change IPsec proposal	Changes take effect after the traffic encryption key (TEK) rekey.		
Group changes:			
Delete group name	Send “delete all” to group members. Delete all IKE SAs in the group. Delete all keys in the group immediately. Delete all registered members in the group.	Send “delete all” to sub-servers. Delete all keys in the group immediately. Mark all peers inactive. Delete sub-server IKE SAs. Delete all member IKE SAs.	Delete all member IKE SAs. Delete all keys in the group immediately. Delete all registered members in the group. Mark peer inactive. Delete peer server IKE SAs.
Change ID	Send “delete all” to all members. Delete all IKE SAs in the group. Delete all keys in the group immediately. Delete all registered members in the group. Generate new keys according to the configuration.	Send ”delete all” to sub-servers. Delete all member IKE SAs in the group. Delete all keys in the group immediately. Mark all peers inactive. Delete all peer server IKE SAs. Generate new keys according to the configuration.	Delete all member IKE SAs in the group. Delete all keys in the group immediately. Delete all registered members in the group. Mark peer inactive. Delete peer server IKE SAs. Initiate new **cluster-init** exchange.
Add or delete IKE gateway	No changes for additions. For deletions, delete the IKE SA and registered members for the affected gateway.		

Copyright © 2016, Juniper Networks, Inc.

Table 569: Effects of Configuration Changes on Group VPNv2 Servers *(continued)*

Configuration Change	Standalone Group VPNv2 Server Action	Group VPNv2 Server Cluster Action	
		Root-server	Sub-server
Add or change anti-replay time window	New value takes effect after the TEK rekey.		
Add or change no anti-replay	New value takes effect after the TEK rekey.		
Server-member communication changes:			
Add	Delete all registered members. Generate key encryption key (KEK) SA.	Generate KEK SA. Send new KEK SA to sub-server. Delete all member IKE SAs.	Delete all registered members.
Change	New value takes effect after KEK rekey.		
Delete	Send delete to delete all KEK SAs. Delete KEK SA.	Send delete to sub-servers. Delete KEK SA. Delete all member IKE SAs.	Delete KEK SA.
IPsec SA:			
Add	Generate new TEK SA. Update the new TEK SA on members.	Generate new TEK SA. Send new TEK SA to sub-servers.	No action.
Change	New value takes effect after TEK rekey. If the match-policy changes, the current TEK is removed immediately and delete groupkey-push is sent because members need to be explicitly notified that this configuration is removed.	If the match-policy changes, send delete to sub-servers. Delete TEK immediately.	If the match-policy changes, delete TEK immediately.
Delete	Delete TEK immediately. Send delete to delete this TEK SA.	Send delete to sub-servers. Delete TEK immediately.	Delete TEK immediately.

Table 570 describes the effects of changing Group VPNv2 server cluster configuration.

NOTE: You must ensure that there is only one root-server in a server cluster at any time.

Copyright © 2016, Juniper Networks, Inc.

Table 570: Effects of Group VPNv2 Server Cluster Configuration Changes

Server Cluster Configuration Change	Group VPNv2 Server Cluster	
	Root-server	Sub-server
IKE proposal, policy, or gateway (cluster peer)	For additions, there is no change. For changes or deletions, delete the IKE SA for the affected peer.	
Server cluster:		
Add	None.	Send “delete all” to group members. Delete all member IKE SAs in the group. Delete all TEKs and KEKs immediately in the group. Delete all registered members in the group. Send **cluster-init** to root-server.
Change role NOTE: You must ensure that there is only one root-server in a server cluster at any time.	Send “delete all” to sub-servers. Delete all member IKE SAs in the group. Delete all TEKs and KEKs immediately in the group. Mark all peers inactive. Delete all peer server IKE SAs. Send **cluster-init** to root-server.	Rekey TEK. Rekey KEK. Send new keys to sub-servers. Send new keys to members.
Add peer	None.	
Delete peer	Mark peer inactive. Clear peer IKE SA.	Mark peer inactive. Clear KEK. Clear TEK. Clear peer IKE SA.
Change retransmission period	None.	
Delete server cluster	Send “delete all” to sub-servers. Delete all TEKs and KEKs immediately in the group. Mark all peers inactive. Delete all peer server IKE SAs. Generate new TEKs and KEKs according to the configuration.	Delete all member IKE SAs in the group. Delete all TEKs and KEKs immediately in the group. Delete all registered members in the group. Mark peer inactive. Delete peer server IKE SAs. Generate new TEK and KEK according to the configuration.

Related Documentation

- Understanding Group VPNv2 Server Cluster Messages on page 6704
- Understanding Group Key Operations on page 6694

Migrating a Standalone Group Server to a Server Cluster

This section describes how to migrate a standalone Group VPNv2 server to a Group VPNv2 server cluster.

Copyright © 2016, Juniper Networks, Inc.

To migrate a standalone Group VPNv2 server to a root-server:

NOTE: We highly recommend that the root-server be a chassis cluster.

1. Upgrade the standalone Group VPNv2 server to a chassis cluster. See *SRX Series Chassis Cluster Configuration Overview*.

NOTE: A reboot is required during the upgrade of a standalone SRX Series device to a chassis cluster node. Traffic loss is expected.

2. On the chassis cluster, add the Group VPNv2 server cluster root-server configuration. The configured server role for the cluster must be **root-server**.

   There should be no traffic loss among existing group members during the configuration change.

To add a sub-server to the Group VPNv2 server cluster:

1. On the root-server, configure both a Group VPNv2 server IKE gateway and a server cluster IKE gateway for the sub-server. SAs and existing member traffic should not be impacted.
2. On the sub-server, configure the server cluster. Remember that the Group VPNv2 configuration must be the same on each server in the cluster, with the exception of the Group VPNv2 server IKE gateways, the server role in the cluster, and the server cluster IKE gateway configurations. On the sub-server, the configured server role in the cluster must be **sub-server**. Configure a Group VPNv2 server IKE gateway and a server cluster IKE gateway for the root-server.

To delete a sub-server from the Group VPNv2 server cluster:

1. On the root-server, delete both the Group VPNv2 server IKE gateway and the server cluster IKE gateway configurations for the sub-server. SAs and existing member traffic should not be impacted.
2. Power off the sub-server.

Related Documentation

- Understanding Group VPNv2 Server Clusters on page 6699

Example: Configuring a Group VPNv2 Server Cluster and Members

This example shows how to configure a Group VPNv2 server cluster to provide group controller/key server (GCKS) redundancy and scaling to Group VPNv2 group members.

- Requirements on page 6710
- Overview on page 6710

Copyright © 2016, Juniper Networks, Inc.

- Configuration on page 6712
- Verification on page 6765

Requirements

The example uses the following hardware and software components:

- Eight supported SRX Series devices or vSRX instances running Junos OS Release 15.1X49-D30 or later that support Group VPNv2:
  - Two devices or instances are configured to operate as a chassis cluster. The chassis cluster operates as the root-server in the Group VPNv2 server cluster. The devices or instances must have the same software version and licenses.

    NOTE: The root-server is responsible for generating and distributing encryption keys to sub-servers in the group VPN server cluster; because of this responsibility, we recommend that the root-server be a chassis cluster.

  - Four other devices or instances operate as sub-servers in the Group VPNv2 server cluster.
  - Two other devices or instances operate as Group VPNv2 group members.
- Two supported MX Series devices running Junos OS Release 15.1R2 or later that support Group VPNv2. These devices operate as Group VPNv2 group members.

A hostname, a root administrator password, and management access must be configured on each SRX Series device or vSRX instance. We recommend that NTP also be configured on each device.

NOTE: The configurations in this example focus on what is needed for Group VPNv2 operation, based on the topology shown in Figure 301. Some configurations, such as interface, routing, or chassis cluster setups, are not included here. For example, Group VPNv2 operation requires a working routing topology that allows client devices to reach their intended sites throughout the network; this example does not cover the configuration of static or dynamic routing.

Overview

In this example, the Group VPNv2 network consists of a server cluster and four members. The server cluster consists of a root-server and four sub-servers. Two of the members are SRX Series devices or vSRX instances while the other two members are MX Series devices.

The group VPN SAs must be protected by a Phase 1 SA. Therefore, the group VPN configuration must include configuring IKE Phase 1 negotiations on the root-server, the sub-servers, and the group members. IKE configurations are described as follows.

Copyright © 2016, Juniper Networks, Inc.

On the root-server:

- The IKE policy **SubSrv** is used to establish Phase 1 SAs with each sub-server.
- An IKE gateway is configured with dead peer detection (DPD) for each sub-server.
- The server cluster role is **root-server** and each sub-server is configured as an IKE gateway for the server cluster.

NOTE: The root-server should be configured to support chassis cluster operation. In the example, redundant Ethernet interfaces on the root-server connect to each of the sub-servers in the server cluster; the entire chassis cluster configuration is not shown.

On each sub-server:

- Two IKE policies are configured: **RootSrv** is used to establish a Phase 1 SA with the root-server, and **GMs** is used to establish Phase 1 SAs with each group member.

NOTE: Preshared keys are used to secure the Phase 1 SAs between the root-server and the sub-servers and between the sub-servers and the group members. Ensure that the preshared keys used are strong keys. On the sub-servers, the preshared key configured for the IKE policy **RootSrv** must match the preshared key configured on the root-server, and the preshared key configured for the IKE policy **GMs** must match the preshared key configured on the group members.

- An IKE gateway is configured with DPD for the root-server. In addition, an IKE gateway is configured for each group member.
- The server cluster role is **sub-server** and the root-server is configured as the IKE gateway for the server cluster.

On each group member:

- The IKE policy **SubSrv** is used to establish Phase 1 SAs with the sub-servers.
- The IKE gateway configuration includes the addresses for the sub-servers.

On SRX Series devices or vSRX group members, an IPsec policy is configured for the group with the LAN zone as the from-zone (incoming traffic) and the WAN zone as the to-zone (outgoing traffic). A security policy is also needed to allow traffic between the LAN and WAN zones.

The same group identifier must be configured on both the group server and the group members. In this example, the group name is GROUP_ID-0001 and the group identifier is 1. The group policy configured on the server specifies that the SA and key are applied to traffic between subnetworks in the 172.16.0.0/12 range.

Copyright © 2016, Juniper Networks, Inc.

Topology

Figure 301 shows the Juniper Networks devices to be configured for this example.

Figure 301: Group VPNv2 Server Cluster with SRX or vSRX and MX Series Members

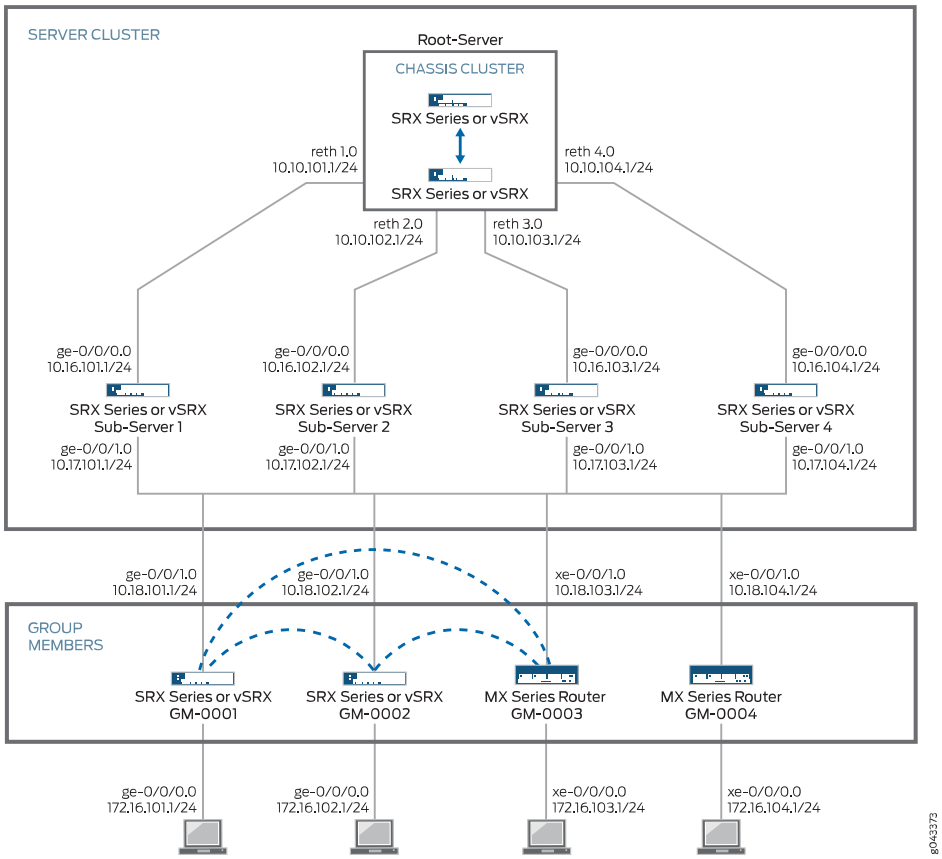

Configuration

Configuring the Root-Server

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth1 unit 0 description To_SubSrv01
set interfaces reth1 unit 0 family inet address 10.10.101.1/24
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth2 unit 0 description To_SubSrv02
set interfaces reth2 unit 0 family inet address 10.10.102.1/24
set interfaces reth3 redundant-ether-options redundancy-group 1
```


Copyright © 2016, Juniper Networks, Inc.

```
set interfaces reth3 unit 0 description To_SubSrv03
set interfaces reth3 unit 0 family inet address 10.10.103.1/24
set interfaces reth4 redundant-ether-options redundancy-group 1
set interfaces reth4 unit 0 description To_SubSrv04
set interfaces reth4 unit 0 family inet address 10.10.104.1/24
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces reth1.0
set security zones security-zone GROUPVPN interfaces reth2.0
set security zones security-zone GROUPVPN interfaces reth3.0
set security zones security-zone GROUPVPN interfaces reth4.0
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set chassis cluster reth-count 5
set chassis cluster redundancy-group 1 node 0 priority 254
set chassis cluster redundancy-group 1 node 1 priority 1
set chassis cluster redundancy-group 0 node 0 priority 254
set chassis cluster redundancy-group 0 node 1 priority 1
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy SubSrv mode main
set security group-vpn server ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway SubSrv01 ike-policy SubSrv
set security group-vpn server ike gateway SubSrv01 address 10.16.101.1
set security group-vpn server ike gateway SubSrv01 dead-peer-detection always-send
set security group-vpn server ike gateway SubSrv01 local-address 10.10.101.1
set security group-vpn server ike gateway SubSrv02 ike-policy SubSrv
set security group-vpn server ike gateway SubSrv02 address 10.16.102.1
set security group-vpn server ike gateway SubSrv02 dead-peer-detection always-send
set security group-vpn server ike gateway SubSrv02 local-address 10.10.102.1
set security group-vpn server ike gateway SubSrv03 ike-policy SubSrv
set security group-vpn server ike gateway SubSrv03 address 10.16.103.1
set security group-vpn server ike gateway SubSrv03 dead-peer-detection always-send
set security group-vpn server ike gateway SubSrv03 local-address 10.10.103.1
set security group-vpn server ike gateway SubSrv04 ike-policy SubSrv
set security group-vpn server ike gateway SubSrv04 address 10.16.104.1
set security group-vpn server ike gateway SubSrv04 dead-peer-detection always-send
set security group-vpn server ike gateway SubSrv04 local-address 10.10.104.1
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
```

Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 server-cluster server-role root-server
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway SubSrv01
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway SubSrv02
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway SubSrv03
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway SubSrv04
set security group-vpn server group GROUP_ID-0001 server-cluster retransmission-period
  10
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the root-server:

1. Configure security zones and security policies.

```
[edit interfaces]
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth1 unit 0 description To_SubSrv01
user@host# set reth1 unit 0 family inet address 10.10.101.1/24
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth2 unit 0 description To_SubSrv02
user@host# set reth2 unit 0 family inet address 10.10.102.1/24
user@host# set reth3 redundant-ether-options redundancy-group 1
user@host# set reth3 unit 0 description To_SubSrv03
user@host# set reth3 unit 0 family inet address 10.10.103.1/24
user@host# set reth4 redundant-ether-options redundancy-group 1
user@host# set reth4 unit 0 description To_SubSrv04
user@host# set reth4 unit 0 family inet address 10.10.104.1/24

[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interfaces reth1.0
user@host# set interfaces reth2.0
user@host# set interfaces reth3.0
user@host# set interfaces reth4.0

[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit security policies]
user@host# set default-policy deny-all
```

2. Configure the chassis cluster.

```
[edit chassis cluster]
user@host# set reth-count 5
user@host# set redundancy-group 1 node 0 priority 254
user@host# set redundancy-group 1 node 1 priority 1
user@host# set redundancy-group 0 node 0 priority 254
user@host# set redundancy-group 0 node 1 priority 1
```

3. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy SubSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway SubSrv01]
user@host# set ike-policy SubSrv
user@host# set address 10.16.101.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.10.101.1

[edit security group-vpn server ike gateway SubSrv02]
user@host# set ike-policy SubSrv
user@host# set address 10.16.102.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.10.102.1

[edit security group-vpn server ike gateway SubSrv03]
user@host# set ike-policy SubSrv
user@host# set address 10.16.103.1
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set dead-peer-detection always-send
user@host# set local-address 10.10.103.1

[edit security group-vpn server ike gateway SubSrv04]
user@host# set ike-policy SubSrv
user@host# set address 10.16.104.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.10.104.1
```

4. Configure the IPsec SA.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600
```

5. Configure the VPN group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set server-cluster server-role root-server
user@host# set server-cluster ike-gateway SubSrv01
user@host# set server-cluster ike-gateway SubSrv02
user@host# set server-cluster ike-gateway SubSrv03
user@host# set server-cluster ike-gateway SubSrv04
user@host# set server-cluster retransmission-period 10
user@host# set anti-replay-time-window 1000
user@host# set server-member-communication communication-type unicast
user@host# set server-member-communication encryption-algorithm aes-256-cbc
user@host# set server-member-communication lifetime-seconds 7200
user@host# set server-member-communication sig-hash-algorithm sha-256
```

6. Configure the group policy.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 source 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 destination 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 protocol 0
user@host# set ipsec-sa GROUP_ID-0001 proposal AES256-SHA256-L3600
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show chassis cluster**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
reth1 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    description To_SubSrv01;
    family inet {
      address 10.10.101.1/24;
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
}
reth2 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    description To_SubSrv02;
    family inet {
      address 10.10.102.1/24;
    }
  }
}
reth3 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    description To_SubSrv03;
    family inet {
      address 10.10.103.1/24;
    }
  }
}
reth4 {
  redundant-ether-options {
    redundancy-group 1;
  }
  unit 0 {
    description To_SubSrv04;
    family inet {
      address 10.10.104.1/24;
    }
  }
}
[edit]
user@host# show chassis cluster
reth-count 5;
redundancy-group 1 {
  node 0 priority 254;
  node 1 priority 1;
}
redundancy-group 0 {
  node 0 priority 254;
  node 1 priority 1;
}
[edit]
user@host# show security
group-vpn {
  server {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
        authentication-method pre-shared-keys;
        authentication-algorithm sha-256;
        dh-group group14;
        encryption-algorithm aes-256-cbc;
      }
```

```
    policy SubSrv {
      mode main;
      proposals PSK-SHA256-DH14-AES256;
      pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
    }
    gateway SubSrv01 {
      ike-policy SubSrv;
      address 10.16.101.1;
      dead-peer-detection always-send;
      local-address 10.10.101.1;
    }
    gateway SubSrv02 {
      ike-policy SubSrv;
      address 10.16.102.1;
      dead-peer-detection always-send;
      local-address 10.10.102.1;
    }
    gateway SubSrv03 {
      ike-policy SubSrv;
      address 10.16.103.1;
      dead-peer-detection always-send;
      local-address 10.10.103.1;
    }
    gateway SubSrv04 {
      ike-policy SubSrv;
      address 10.16.104.1;
      dead-peer-detection always-send;
      local-address 10.10.104.1;
    }
  }
  ipsec {
    proposal AES256-SHA256-L3600 {
      authentication-algorithm hmac-sha-256-128;
      encryption-algorithm aes-256-cbc;
      lifetime-seconds 3600;
    }
  }
  group GROUP_ID-0001 {
    group-id 1;
    member-threshold 2000;
    server-cluster {
      server-role root-server;
      ike-gateway SubSrv01;
      ike-gateway SubSrv02;
      ike-gateway SubSrv03;
      ike-gateway SubSrv04;
      retransmission-period 10;
    }
    anti-replay-time-window 1000;
    server-member-communication {
      communication-type unicast;
      lifetime-seconds 7200;
      encryption-algorithm aes-256-cbc;
      sig-hash-algorithm sha-256;
    }
    ipsec-sa GROUP_ID-0001 {
```


Copyright © 2016, Juniper Networks, Inc.

```
            proposal AES256-SHA256-L3600;
            match-policy 1 {
              source 172.16.0.0/12;
              destination 172.16.0.0/12;
              protocol 0;
            }
          }
        }
      }
    }
    policies {
      global {
        policy 1000 {
          match {
            source-address any;
            destination-address any;
            application any;
            from-zone any;
            to-zone any;
          }
          then {
            deny;
            log {
              session-init;
            }
            count;
          }
        }
      }
      default-policy {
        deny-all;
      }
    }
    zones {
      security-zone GROUPVPN {
        host-inbound-traffic {
          system-services {
            ike;
            ssh;
            ping;
          }
        }
        interfaces {
          reth1.0;
          reth2.0;
          reth3.0;
          reth4.0;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring Sub-Server 1

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_RootSrv
set interfaces ge-0/0/0 unit 0 family inet address 10.16.101.1/24
set interfaces ge-0/0/1 unit 0 description To_WAN
set interfaces ge-0/0/1 unit 0 family inet address 10.17.101.1/24
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces ge-0/0/0.0
set security zones security-zone GROUPVPN interfaces ge-0/0/1.0
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy RootSrv mode main
set security group-vpn server ike policy RootSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy RootSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike policy GMs mode main
set security group-vpn server ike policy GMs proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy GMs pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway RootSrv ike-policy RootSrv
set security group-vpn server ike gateway RootSrv address 10.10.101.1
set security group-vpn server ike gateway RootSrv dead-peer-detection always-send
set security group-vpn server ike gateway RootSrv local-address 10.16.101.1
set security group-vpn server ike gateway GM-0001 ike-policy GMs
set security group-vpn server ike gateway GM-0001 address 10.18.101.1
set security group-vpn server ike gateway GM-0001 local-address 10.17.101.1
set security group-vpn server ike gateway GM-0002 ike-policy GMs
set security group-vpn server ike gateway GM-0002 address 10.18.102.1
set security group-vpn server ike gateway GM-0002 local-address 10.17.101.1
set security group-vpn server ike gateway GM-0003 ike-policy GMs
set security group-vpn server ike gateway GM-0003 address 10.18.103.1
set security group-vpn server ike gateway GM-0003 local-address 10.17.101.1
set security group-vpn server ike gateway GM-0004 ike-policy GMs
set security group-vpn server ike gateway GM-0004 address 10.18.104.1
set security group-vpn server ike gateway GM-0004 local-address 10.17.101.1
```


Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 server-cluster server-role sub-server
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway RootSrv
set security group-vpn server group GROUP_ID-0001 server-cluster retransmission-period
  10
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0001
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0002
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0003
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0004
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the sub-server in the Group VPNv2 server cluster:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_RootSrv
user@host# set ge-0/0/0 unit 0 family inet address 10.16.101.1/24
user@host# set ge-0/0/1 unit 0 description To_WAN
user@host# set ge-0/0/1 unit 0 family inet address 10.17.101.1/24


[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0
user@host# set interfaces ge-0/0/1.0
```

Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit security policies]
user@host# set default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy RootSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike policy GMs]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway RootSrv]
user@host# set ike-policy RootSrv
user@host# set address 10.10.101.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.16.101.1

[edit security group-vpn server ike gateway GM-0001]
user@host# set ike-policy GMs
user@host# set address 10.18.101.1
user@host# set local-address 10.17.101.1

[edit security group-vpn server ike gateway GM-0002]
user@host# set ike-policy GMs
user@host# set address 10.18.102.1
user@host# set local-address 10.17.101.1

[edit security group-vpn server ike gateway GM-0003]
user@host# set ike-policy GMs
user@host# set address 10.18.103.1
user@host# set local-address 10.17.101.1

[edit security group-vpn server ike gateway GM-0004]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ike-policy GMs
user@host# set address 10.18.104.1
user@host# set local-address 10.17.101.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600
```

4. Configure the VPN group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set server-cluster server-role sub-server
user@host# set server-cluster ike-gateway RootSrv
user@host# set server-cluster retransmission-period 10
user@host# set ike-gateway GM-0001
user@host# set ike-gateway GM-0002
user@host# set ike-gateway GM-0003
user@host# set ike-gateway GM-0004
user@host# set anti-replay-time-window 1000
user@host# set server-member-communication communication-type unicast
user@host# set server-member-communication encryption-algorithm aes-256-cbc
user@host# set server-member-communication lifetime-seconds 7200
user@host# set server-member-communication sig-hash-algorithm sha-256
user@host# set ipsec-sa GROUP_ID-0001 proposal AES256-SHA256-L3600
```

5. Configure the group policy.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 source 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 destination 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 protocol 0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_RootSrv;
    family inet {
      address 10.16.101.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_WAN;
    family inet {
      address 10.17.101.1/24;
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
}
[edit]
user@host# show security
group-vpn {
    server {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                authentication-algorithm sha-256;
                dh-group group14;
                encryption-algorithm aes-256-cbc;
            }
            policy RootSrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            policy GMs {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            gateway RootSrv {
                ike-policy RootSrv;
                address 10.10.101.1;
                dead-peer-detection always-send;
                local-address 10.16.101.1;
            }
            gateway GM-0001 {
                ike-policy GMs;
                address 10.18.101.1;
                local-address 10.17.101.1;
            }
            gateway GM-0002 {
                ike-policy GMs;
                address 10.18.102.1;
                local-address 10.17.101.1;
            }
            gateway GM-0003 {
                ike-policy GMs;
                address 10.18.103.1;
                local-address 10.17.101.1;
            }
            gateway GM-0004 {
                ike-policy GMs;
                address 10.18.104.1;
                local-address 10.17.101.1;
            }
        }
        ipsec {
            proposal AES256-SHA256-L3600 {
                authentication-algorithm hmac-sha-256-128;
                encryption-algorithm aes-256-cbc;
                lifetime-seconds 3600;
            }
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
        group GROUP_ID-0001 {
            group-id 1;
            member-threshold 2000;
            server-cluster {
                server-role sub-server;
                ike-gateway RootSrv;
                retransmission-period 10;
            }
            ike-gateway GM-0001;
            ike-gateway GM-0002;
            ike-gateway GM-0003;
            ike-gateway GM-0004;
            anti-replay-time-window 1000;
            server-member-communication {
                communication-type unicast;
                lifetime-seconds 7200;
                encryption-algorithm aes-256-cbc;
                sig-hash-algorithm sha-256;
            }
            ipsec-sa GROUP_ID-0001 {
                proposal AES256-SHA256-L3600;
                match-policy 1 {
                    source 172.16.0.0/12;
                    destination 172.16.0.0/12;
                    protocol 0;
                }
            }
        }
    }
}
policies {
    global {
        policy 1000 {
            match {
                source-address any;
                destination-address any;
                application any;
                from-zone any;
                to-zone any;
            }
            then {
                deny;
                log {
                    session-init;
                }
                count;
            }
        }
    }
    default-policy {
        deny-all;
    }
}
zones {
    security-zone GROUPVPN {
```

```
        host-inbound-traffic {
          system-services {
            ike;
            ssh;
            ping;
          }
        }
        interfaces {
          ge-0/0/0.0;
          ge-0/0/1.0;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Sub-Server 2

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_RootSrv
set interfaces ge-0/0/0 unit 0 family inet address 10.16.102.1/24
set interfaces ge-0/0/1 unit 0 description To_WAN
set interfaces ge-0/0/1 unit 0 family inet address 10.17.102.1/24
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces ge-0/0/0.0
set security zones security-zone GROUPVPN interfaces ge-0/0/1.0
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy RootSrv mode main
set security group-vpn server ike policy RootSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy RootSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike policy GMs mode main
set security group-vpn server ike policy GMs proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy GMs pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway RootSrv ike-policy RootSrv
```


Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server ike gateway RootSrv address 10.10.102.1
set security group-vpn server ike gateway RootSrv dead-peer-detection always-send
set security group-vpn server ike gateway RootSrv local-address 10.16.102.1
set security group-vpn server ike gateway GM-0001 ike-policy GMs
set security group-vpn server ike gateway GM-0001 address 10.18.101.1
set security group-vpn server ike gateway GM-0001 local-address 10.17.102.1
set security group-vpn server ike gateway GM-0002 ike-policy GMs
set security group-vpn server ike gateway GM-0002 address 10.18.102.1
set security group-vpn server ike gateway GM-0002 local-address 10.17.102.1
set security group-vpn server ike gateway GM-0003 ike-policy GMs
set security group-vpn server ike gateway GM-0003 address 10.18.103.1
set security group-vpn server ike gateway GM-0003 local-address 10.17.102.1
set security group-vpn server ike gateway GM-0004 ike-policy GMs
set security group-vpn server ike gateway GM-0004 address 10.18.104.1
set security group-vpn server ike gateway GM-0004 local-address 10.17.102.1
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 server-cluster server-role sub-server
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway RootSrv
set security group-vpn server group GROUP_ID-0001 server-cluster retransmission-period
  10
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0001
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0002
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0003
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0004
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the sub-server in the Group VPNv2 server cluster:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_RootSrv
user@host# set ge-0/0/0 unit 0 family inet address 10.16.102.1/24
user@host# set ge-0/0/1 unit 0 description To_WAN
user@host# set ge-0/0/1 unit 0 family inet address 10.17.102.1/24

[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0
user@host# set interfaces ge-0/0/1.0

[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit security policies]
user@host# set default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy RootSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike policy GMs]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway RootSrv]
user@host# set ike-policy RootSrv
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set address 10.10.102.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.16.102.1

[edit security group-vpn server ike gateway GM-0001]
user@host# set ike-policy GMs
user@host# set address 10.18.101.1
user@host# set local-address 10.17.102.1

[edit security group-vpn server ike gateway GM-0002]
user@host# set ike-policy GMs
user@host# set address 10.18.102.1
user@host# set local-address 10.17.102.1

[edit security group-vpn server ike gateway GM-0003]
user@host# set ike-policy GMs
user@host# set address 10.18.103.1
user@host# set local-address 10.17.102.1

[edit security group-vpn server ike gateway GM-0004]
user@host# set ike-policy GMs
user@host# set address 10.18.104.1
user@host# set local-address 10.17.102.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600
```

4. Configure the VPN group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set server-cluster server-role sub-server
user@host# set server-cluster ike-gateway RootSrv
user@host# set server-cluster retransmission-period 10
user@host# set ike-gateway GM-0001
user@host# set ike-gateway GM-0002
user@host# set ike-gateway GM-0003
user@host# set ike-gateway GM-0004
user@host# set anti-replay-time-window 1000
user@host# set server-member-communication communication-type unicast
user@host# set server-member-communication encryption-algorithm aes-256-cbc
user@host# set server-member-communication lifetime-seconds 7200
user@host# set server-member-communication sig-hash-algorithm sha-256
```

5. Configure the group policy.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 source 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 destination 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 protocol 0
user@host# set ipsec-sa GROUP_ID-0001 proposal AES256-SHA256-L3600
```

Copyright © 2016, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        description To_RootSrv;
        family inet {
            address 10.16.102.1/24;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        description To_WAN;
        family inet {
            address 10.17.102.1/24;
        }
    }
}
[edit]
user@host# show security
group-vpn {
    server {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                authentication-algorithm sha-256;
                dh-group group14;
                encryption-algorithm aes-256-cbc;
            }
            policy RootSrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            policy GMs {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            gateway RootSrv {
                ike-policy RootSrv;
                address 10.10.102.1;
                dead-peer-detection always-send;
                local-address 10.16.102.1;
            }
            gateway GM-0001 {
                ike-policy GMs;
                address 10.18.101.1;
                local-address 10.17.102.1;
            }
            gateway GM-0002 {
```


Copyright © 2016, Juniper Networks, Inc.

```
            ike-policy GMs;
            address 10.18.102.1;
            local-address 10.17.102.1;
          }
          gateway GM-0003 {
            ike-policy GMs;
            address 10.18.103.1;
            local-address 10.17.102.1;
          }
          gateway GM-0004 {
            ike-policy GMs;
            address 10.18.104.1;
            local-address 10.17.102.1;
          }
        }
        ipsec {
          proposal AES256-SHA256-L3600 {
            authentication-algorithm hmac-sha-256-128;
            encryption-algorithm aes-256-cbc;
            lifetime-seconds 3600;
          }
        }
        group GROUP_ID-0001 {
          group-id 1;
          member-threshold 2000;
          server-cluster {
            server-role sub-server;
            ike-gateway RootSrv;
            retransmission-period 10;
          }
          ike-gateway GM-0001;
          ike-gateway GM-0002;
          ike-gateway GM-0003;
          ike-gateway GM-0004;
          anti-replay-time-window 1000;
          server-member-communication {
            communication-type unicast;
            lifetime-seconds 7200;
            encryption-algorithm aes-256-cbc;
            sig-hash-algorithm sha-256;
          }
          ipsec-sa GROUP_ID-0001 {
            proposal AES256-SHA256-L3600;
            match-policy 1 {
              source 172.16.0.0/12;
              destination 172.16.0.0/12;
              protocol 0;
            }
          }
        }
      }
    }
    policies {
      global {
        policy 1000 {
          match {
```

Copyright © 2016, Juniper Networks, Inc.

```
            source-address any;
            destination-address any;
            application any;
            from-zone any;
            to-zone any;
          }
          then {
            deny;
            log {
              session-init;
            }
            count;
          }
        }
      }
      default-policy {
        deny-all;
      }
    }
    zones {
      security-zone GROUPVPN {
        host-inbound-traffic {
          system-services {
            ike;
            ssh;
            ping;
          }
        }
        interfaces {
          ge-0/0/0.0;
          ge-0/0/1.0;
        }
      }
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Sub-Server 3

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_RootSrv
set interfaces ge-0/0/0 unit 0 family inet address 10.16.103.1/24
set interfaces ge-0/0/1 unit 0 description To_WAN
set interfaces ge-0/0/1 unit 0 family inet address 10.17.103.1/24
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces ge-0/0/0.0
set security zones security-zone GROUPVPN interfaces ge-0/0/1.0
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
```


Copyright © 2016, Juniper Networks, Inc.

```
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy RootSrv mode main
set security group-vpn server ike policy RootSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy RootSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike policy GMs mode main
set security group-vpn server ike policy GMs proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy GMs pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway RootSrv ike-policy RootSrv
set security group-vpn server ike gateway RootSrv address 10.10.103.1
set security group-vpn server ike gateway RootSrv dead-peer-detection always-send
set security group-vpn server ike gateway RootSrv local-address 10.16.103.1
set security group-vpn server ike gateway GM-0001 ike-policy GMs
set security group-vpn server ike gateway GM-0001 address 10.18.101.1
set security group-vpn server ike gateway GM-0001 local-address 10.17.103.1
set security group-vpn server ike gateway GM-0002 ike-policy GMs
set security group-vpn server ike gateway GM-0002 address 10.18.102.1
set security group-vpn server ike gateway GM-0002 local-address 10.17.103.1
set security group-vpn server ike gateway GM-0003 ike-policy GMs
set security group-vpn server ike gateway GM-0003 address 10.18.103.1
set security group-vpn server ike gateway GM-0003 local-address 10.17.103.1
set security group-vpn server ike gateway GM-0004 ike-policy GMs
set security group-vpn server ike gateway GM-0004 address 10.18.104.1
set security group-vpn server ike gateway GM-0004 local-address 10.17.103.1
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 server-cluster server-role sub-server
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway RootSrv
set security group-vpn server group GROUP_ID-0001 server-cluster retransmission-period
  10
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0001
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0002
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0003
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0004
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
```

Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the sub-server in the Group VPNv2 server cluster:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_RootSrv
user@host# set ge-0/0/0 unit 0 family inet address 10.16.103.1/24
user@host# set ge-0/0/1 unit 0 description To_WAN
user@host# set ge-0/0/1 unit 0 family inet address 10.17.103.1/24

[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0
user@host# set interfaces ge-0/0/1.0

[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit security policies]
user@host# set default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy RootSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike policy GMs]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway RootSrv]
user@host# set ike-policy RootSrv
user@host# set address 10.10.103.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.16.103.1

[edit security group-vpn server ike gateway GM-0001]
user@host# set ike-policy GMs
user@host# set address 10.18.101.1
user@host# set local-address 10.17.103.1

[edit security group-vpn server ike gateway GM-0002]
user@host# set ike-policy GMs
user@host# set address 10.18.102.1
user@host# set local-address 10.17.103.1

[edit security group-vpn server ike gateway GM-0003]
user@host# set ike-policy GMs
user@host# set address 10.18.103.1
user@host# set local-address 10.17.103.1

[edit security group-vpn server ike gateway GM-0004]
user@host# set ike-policy GMs
user@host# set address 10.18.104.1
user@host# set local-address 10.17.103.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600
```

4. Configure the VPN group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set server-cluster server-role sub-server
user@host# set server-cluster ike-gateway RootSrv
user@host# set server-cluster retransmission-period 10
user@host# set ike-gateway GM-0001
user@host# set ike-gateway GM-0002
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ike-gateway GM-0003
user@host# set ike-gateway GM-0004
user@host# set anti-replay-time-window 1000
user@host# set server-member-communication communication-type unicast
user@host# set server-member-communication encryption-algorithm aes-256-cbc
user@host# set server-member-communication lifetime-seconds 7200
user@host# set server-member-communication sig-hash-algorithm sha-256
```

5. Configure the group policy.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 source 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 destination 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 protocol 0
user@host# set ipsec-sa GROUP_ID-0001 proposal AES256-SHA256-L3600
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        description To_RootSrv;
        family inet {
            address 10.16.103.1/24;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        description To_WAN;
        family inet {
            address 10.17.103.1/24;
        }
    }
}
[edit]
user@host# show security
group-vpn {
    server {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                authentication-algorithm sha-256;
                dh-group group14;
                encryption-algorithm aes-256-cbc;
            }
            policy RootSrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            policy GMs {
                mode main;
```

 Copyright © 2016, Juniper Networks, Inc.

```
        proposals PSK-SHA256-DH14-AES256;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
      }
      gateway RootSrv {
        ike-policy RootSrv;
        address 10.10.103.1;
        dead-peer-detection always-send;
        local-address 10.16.103.1;
      }
      gateway GM-0001 {
        ike-policy GMs;
        address 10.18.101.1;
        local-address 10.17.103.1;
      }
      gateway GM-0002 {
        ike-policy GMs;
        address 10.18.102.1;
        local-address 10.17.103.1;
      }
      gateway GM-0003 {
        ike-policy GMs;
        address 10.18.103.1;
        local-address 10.17.103.1;
      }
      gateway GM-0004 {
        ike-policy GMs;
        address 10.18.104.1;
        local-address 10.17.103.1;
      }
    }
    ipsec {
      proposal AES256-SHA256-L3600 {
        authentication-algorithm hmac-sha-256-128;
        encryption-algorithm aes-256-cbc;
        lifetime-seconds 3600;
      }
    }
    group GROUP_ID-0001 {
      group-id 1;
      member-threshold 2000;
      server-cluster {
        server-role sub-server;
        ike-gateway RootSrv;
        retransmission-period 10;
      }
      ike-gateway GM-0001;
      ike-gateway GM-0002;
      ike-gateway GM-0003;
      ike-gateway GM-0004;
      anti-replay-time-window 1000;
      server-member-communication {
        communication-type unicast;
        lifetime-seconds 7200;
        encryption-algorithm aes-256-cbc;
        sig-hash-algorithm sha-256;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
        ipsec-sa GROUP_ID-0001 {
          proposal AES256-SHA256-L3600;
          match-policy 1 {
            source 172.16.0.0/12;
            destination 172.16.0.0/12;
            protocol 0;
          }
        }
      }
    }
  }
  policies {
    global {
      policy 1000 {
        match {
          source-address any;
          destination-address any;
          application any;
          from-zone any;
          to-zone any;
        }
        then {
          deny;
          log {
            session-init;
          }
          count;
        }
      }
    }
    default-policy {
      deny-all;
    }
  }
  zones {
    security-zone GROUPVPN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
      interfaces {
        ge-0/0/0.0;
        ge-0/0/1.0;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2016, Juniper Networks, Inc.

Configuring Sub-Server 4

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_RootSrv
set interfaces ge-0/0/0 unit 0 family inet address 10.16.104.1/24
set interfaces ge-0/0/1 unit 0 description To_WAN
set interfaces ge-0/0/1 unit 0 family inet address 10.17.104.1/24
set security zones security-zone GROUPVPN host-inbound-traffic system-services ike
set security zones security-zone GROUPVPN host-inbound-traffic system-services ssh
set security zones security-zone GROUPVPN host-inbound-traffic system-services ping
set security zones security-zone GROUPVPN interfaces ge-0/0/0.0
set security zones security-zone GROUPVPN interfaces ge-0/0/1.0
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256 dh-group group14
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn server ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn server ike policy RootSrv mode main
set security group-vpn server ike policy RootSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy RootSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike policy GMs mode main
set security group-vpn server ike policy GMs proposals PSK-SHA256-DH14-AES256
set security group-vpn server ike policy GMs pre-shared-key ascii-text "$ABC123"
set security group-vpn server ike gateway RootSrv ike-policy RootSrv
set security group-vpn server ike gateway RootSrv address 10.10.104.1
set security group-vpn server ike gateway RootSrv dead-peer-detection always-send
set security group-vpn server ike gateway RootSrv local-address 10.16.104.1
set security group-vpn server ike gateway GM-0001 ike-policy GMs
set security group-vpn server ike gateway GM-0001 address 10.18.101.1
set security group-vpn server ike gateway GM-0001 local-address 10.17.104.1
set security group-vpn server ike gateway GM-0002 ike-policy GMs
set security group-vpn server ike gateway GM-0002 address 10.18.102.1
set security group-vpn server ike gateway GM-0002 local-address 10.17.104.1
set security group-vpn server ike gateway GM-0003 ike-policy GMs
set security group-vpn server ike gateway GM-0003 address 10.18.103.1
set security group-vpn server ike gateway GM-0003 local-address 10.17.104.1
set security group-vpn server ike gateway GM-0004 ike-policy GMs
set security group-vpn server ike gateway GM-0004 address 10.18.104.1
set security group-vpn server ike gateway GM-0004 local-address 10.17.104.1
```

Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn server ipsec proposal AES256-SHA256-L3600
  authentication-algorithm hmac-sha-256-128
set security group-vpn server ipsec proposal AES256-SHA256-L3600 encryption-algorithm
  aes-256-cbc
set security group-vpn server ipsec proposal AES256-SHA256-L3600 lifetime-seconds
  3600
set security group-vpn server group GROUP_ID-0001 group-id 1
set security group-vpn server group GROUP_ID-0001 member-threshold 2000
set security group-vpn server group GROUP_ID-0001 server-cluster server-role sub-server
set security group-vpn server group GROUP_ID-0001 server-cluster ike-gateway RootSrv
set security group-vpn server group GROUP_ID-0001 server-cluster retransmission-period
  10
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0001
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0002
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0003
set security group-vpn server group GROUP_ID-0001 ike-gateway GM-0004
set security group-vpn server group GROUP_ID-0001 anti-replay-time-window 1000
set security group-vpn server group GROUP_ID-0001 server-member-communication
  communication-type unicast
set security group-vpn server group GROUP_ID-0001 server-member-communication
  encryption-algorithm aes-256-cbc
set security group-vpn server group GROUP_ID-0001 server-member-communication
  lifetime-seconds 7200
set security group-vpn server group GROUP_ID-0001 server-member-communication
  sig-hash-algorithm sha-256
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001 proposal
  AES256-SHA256-L3600
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 source 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 destination 172.16.0.0/12
set security group-vpn server group GROUP_ID-0001 ipsec-sa GROUP_ID-0001
  match-policy 1 protocol 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the sub-server in the Group VPNv2 server cluster:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_RootSrv
user@host# set ge-0/0/0 unit 0 family inet address 10.16.104.1/24
user@host# set ge-0/0/1 unit 0 description To_WAN
user@host# set ge-0/0/1 unit 0 family inet address 10.17.104.1/24


[edit security zones security-zone GROUPVPN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0
user@host# set interfaces ge-0/0/1.0
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit security policies]
user@host# set default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn server ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn server ike policy RootSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike policy GMs]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn server ike gateway RootSrv]
user@host# set ike-policy RootSrv
user@host# set address 10.10.104.1
user@host# set dead-peer-detection always-send
user@host# set local-address 10.16.104.1

[edit security group-vpn server ike gateway GM-0001]
user@host# set ike-policy GMs
user@host# set address 10.18.101.1
user@host# set local-address 10.17.104.1

[edit security group-vpn server ike gateway GM-0002]
user@host# set ike-policy GMs
user@host# set address 10.18.102.1
user@host# set local-address 10.17.104.1

[edit security group-vpn server ike gateway GM-0003]
user@host# set ike-policy GMs
user@host# set address 10.18.103.1
user@host# set local-address 10.17.104.1

[edit security group-vpn server ike gateway GM-0004]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set ike-policy GMs
user@host# set address 10.18.104.1
user@host# set local-address 10.17.104.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn server ipsec proposal AES256-SHA256-L3600]
user@host# set authentication-algorithm hmac-sha-256-128
user@host# set encryption-algorithm aes-256-cbc
user@host# set lifetime-seconds 3600
```

4. Configure the VPN group.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set group-id 1
user@host# set member-threshold 2000
user@host# set server-cluster server-role sub-server
user@host# set server-cluster ike-gateway RootSrv
user@host# set server-cluster retransmission-period 10
user@host# set ike-gateway GM-0001
user@host# set ike-gateway GM-0002
user@host# set ike-gateway GM-0003
user@host# set ike-gateway GM-0004
user@host# set anti-replay-time-window 1000
user@host# set server-member-communication communication-type unicast
user@host# set server-member-communication encryption-algorithm aes-256-cbc
user@host# set server-member-communication lifetime-seconds 7200
user@host# set server-member-communication sig-hash-algorithm sha-256
```

5. Configure the group policy.

```
[edit security group-vpn server group GROUP_ID-0001]
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 source 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 destination 172.16.0.0/12
user@host# set ipsec-sa GROUP_ID-0001 match-policy 1 protocol 0
user@host# set ipsec-sa GROUP_ID-0001 proposal AES256-SHA256-L3600
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_RootSrv;
    family inet {
      address 10.16.104.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_WAN;
    family inet {
      address 10.17.104.1/24;
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
}
[edit]
user@host# show security
group-vpn {
    server {
        ike {
            proposal PSK-SHA256-DH14-AES256 {
                authentication-method pre-shared-keys;
                authentication-algorithm sha-256;
                dh-group group14;
                encryption-algorithm aes-256-cbc;
            }
            policy RootSrv {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            policy GMs {
                mode main;
                proposals PSK-SHA256-DH14-AES256;
                pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
            }
            gateway RootSrv {
                ike-policy RootSrv;
                address 10.10.104.1;
                dead-peer-detection always-send;
                local-address 10.16.104.1;
            }
            gateway GM-0001 {
                ike-policy GMs;
                address 10.18.101.1;
                local-address 10.17.104.1;
            }
            gateway GM-0002 {
                ike-policy GMs;
                address 10.18.102.1;
                local-address 10.17.104.1;
            }
            gateway GM-0003 {
                ike-policy GMs;
                address 10.18.103.1;
                local-address 10.17.104.1;
            }
            gateway GM-0004 {
                ike-policy GMs;
                address 10.18.104.1;
                local-address 10.17.104.1;
            }
        }
        ipsec {
            proposal AES256-SHA256-L3600 {
                authentication-algorithm hmac-sha-256-128;
                encryption-algorithm aes-256-cbc;
                lifetime-seconds 3600;
            }
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        group GROUP_ID-0001 {
          group-id 1;
          member-threshold 2000;
          server-cluster {
            server-role sub-server;
            ike-gateway RootSrv;
            retransmission-period 10;
          }
          ike-gateway GM-0001;
          ike-gateway GM-0002;
          ike-gateway GM-0003;
          ike-gateway GM-0004;
          anti-replay-time-window 1000;
          server-member-communication {
            communication-type unicast;
            lifetime-seconds 7200;
            encryption-algorithm aes-256-cbc;
            sig-hash-algorithm sha-256;
          }
          ipsec-sa GROUP_ID-0001 {
            proposal AES256-SHA256-L3600;
            match-policy 1 {
              source 172.16.0.0/12;
              destination 172.16.0.0/12;
              protocol 0;
            }
          }
        }
      }
    }
    policies {
      global {
        policy 1000 {
          match {
            source-address any;
            destination-address any;
            application any;
            from-zone any;
            to-zone any;
          }
          then {
            deny;
            log {
              session-init;
            }
            count;
          }
        }
      }
      default-policy {
        deny-all;
      }
    }
    zones {
      security-zone GROUPVPN {
```


Copyright © 2016, Juniper Networks, Inc.

```
    host-inbound-traffic {
      system-services {
        ike;
        ssh;
        ping;
      }
    }
    interfaces {
      ge-0/0/0.0;
      ge-0/0/1.0;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring GM-0001 (SRX Series Device or vSRX Instance)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_LAN
set interfaces ge-0/0/0 unit 0 family inet address 172.16.101.1/24
set interfaces ge-0/0/1 unit 0 description To_SubSrv
set interfaces ge-0/0/1 unit 0 family inet address 10.18.101.1/24
set security zones security-zone LAN host-inbound-traffic system-services ike
set security zones security-zone LAN host-inbound-traffic system-services ssh
set security zones security-zone LAN host-inbound-traffic system-services ping
set security zones security-zone LAN interfaces ge-0/0/0.0
set security zones security-zone WAN host-inbound-traffic system-services ike
set security zones security-zone WAN host-inbound-traffic system-services ssh
set security zones security-zone WAN host-inbound-traffic system-services ping
set security zones security-zone WAN interfaces ge-0/0/1.0
set security address-book global address 172.16.0.0/12 172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match application any
set security policies from-zone LAN to-zone WAN policy 1 then permit
set security policies from-zone LAN to-zone WAN policy 1 then log session-init
set security policies from-zone WAN to-zone LAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match application any
set security policies from-zone WAN to-zone LAN policy 1 then permit
set security policies from-zone WAN to-zone LAN policy 1 then log session-init
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
```

Copyright © 2016, Juniper Networks, Inc.

```
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy SubSrv mode main
set security group-vpn member ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway SubSrv ike-policy SubSrv
set security group-vpn member ike gateway SubSrv server-address 10.17.101.1
set security group-vpn member ike gateway SubSrv server-address 10.17.102.1
set security group-vpn member ike gateway SubSrv server-address 10.17.103.1
set security group-vpn member ike gateway SubSrv server-address 10.17.104.1
set security group-vpn member ike gateway SubSrv local-address 10.18.101.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway SubSrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group-vpn-external-interface
  ge-0/0/1.0
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 recovery-probe
set security ipsec-policy from-zone LAN to-zone WAN ipsec-group-vpn GROUP_ID-0001
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_LAN
user@host# set ge-0/0/0 unit 0 family inet address 172.16.101.1/24
user@host# set ge-0/0/1 unit 0 description To_SubSrv
user@host# set ge-0/0/1 unit 0 family inet address 10.18.101.1/24

[edit security zones security-zone LAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0

[edit security zones security-zone WAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/1.0

[edit security]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set address-book global address 172.16.0.0/12 172.16.0.0/12

[edit security policies from-zone LAN to-zone WAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set policy 1 then log session-init

[edit security policies from-zone WAN to-zone LAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set policy 1 then log session-init

[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count

[edit]
user@host# set security policies default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy SubSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway SubSrv]
user@host# set ike-policy SubSrv
user@host# set server-address 10.17.101.1
user@host# set server-address 10.17.102.1
user@host# set server-address 10.17.103.1
user@host# set server-address 10.17.104.1
user@host# set local-address 10.18.101.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway SubSrv
user@host# set group-vpn-external-interface ge-0/0/1.0
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set group 1
user@host# set recovery-probe
```

4. Configure the IPsec policy.

```
[edit security ipsec-policy from-zone LAN to-zone WAN]
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_LAN;
    family inet {
      address 172.16.101.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_SubSrv;
    family inet {
      address 10.18.101.1/24;
    }
  }
}
[edit]
user@host# show security
address-book {
  global {
    address 172.16.0.0/12 172.16.0.0/12;
  }
}
group-vpn {
  member {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
        authentication-method pre-shared-keys;
        dh-group group14;
        authentication-algorithm sha-256;
        encryption-algorithm aes-256-cbc;
      }
      policy SubSrv {
        mode main;
        proposals PSK-SHA256-DH14-AES256;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
      }
      gateway SubSrv {
        ike-policy SubSrv;
        server-address [ 10.17.101.1 10.17.102.1 10.17.103.1 10.17.104.1 ];
        local-address 10.18.101.1;
      }
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    ipsec {
      vpn GROUP_ID-0001 {
        ike-gateway SubSrv;
        group-vpn-external-interface ge-0/0/1.0;
        group 1;
        recovery-probe;
      }
    }
  }
}
ipsec-policy {
  from-zone LAN to-zone WAN {
    ipsec-group-vpn GROUP_ID-0001;
  }
}
policies {
  from-zone LAN to-zone WAN {
    policy 1 {
      match {
        source-address 172.16.0.0/12;
        destination-address 172.16.0.0/12;
        application any;
      }
      then {
        permit;
        log {
          session-init;
        }
      }
    }
  }
  from-zone WAN to-zone LAN {
    policy 1 {
      match {
        source-address 172.16.0.0/12;
        destination-address 172.16.0.0/12;
        application any;
      }
      then {
        permit;
        log {
          session-init;
        }
      }
    }
  }
  global {
    policy 1000 {
      match {
        source-address any;
        destination-address any;
        application any;
        from-zone any;
        to-zone any;
      }
```

Copyright © 2016, Juniper Networks, Inc.

```
            then {
                deny;
                log {
                    session-init;
                }
                count;
            }
        }
    }
    default-policy {
        deny-all;
    }
}
zones {
    security-zone LAN {
        host-inbound-traffic {
            system-services {
                ike;
                ssh;
                ping;
            }
        }
        interfaces {
            ge-0/0/0.0;
        }
    }
    security-zone WAN {
        host-inbound-traffic {
            system-services {
                ike;
                ssh;
                ping;
            }
        }
        interfaces {
            ge-0/0/1.0;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring GM-0002 (SRX Series Device or vSRX Instance)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 description To_LAN
set interfaces ge-0/0/0 unit 0 family inet address 172.16.102.1/24
set interfaces ge-0/0/1 unit 0 description To_SubSrv
set interfaces ge-0/0/1 unit 0 family inet address 10.18.102.1/24
set security zones security-zone LAN host-inbound-traffic system-services ike
set security zones security-zone LAN host-inbound-traffic system-services ssh
```


Copyright © 2016, Juniper Networks, Inc.

```
set security zones security-zone LAN host-inbound-traffic system-services ping
set security zones security-zone LAN interfaces ge-0/0/0.0
set security zones security-zone WAN host-inbound-traffic system-services ike
set security zones security-zone WAN host-inbound-traffic system-services ssh
set security zones security-zone WAN host-inbound-traffic system-services ping
set security zones security-zone WAN interfaces ge-0/0/1.0
set security address-book global address 172.16.0.0/12 172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone LAN to-zone WAN policy 1 match application any
set security policies from-zone LAN to-zone WAN policy 1 then permit
set security policies from-zone LAN to-zone WAN policy 1 then log session-init
set security policies from-zone WAN to-zone LAN policy 1 match source-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match destination-address
  172.16.0.0/12
set security policies from-zone WAN to-zone LAN policy 1 match application any
set security policies from-zone WAN to-zone LAN policy 1 then permit
set security policies from-zone WAN to-zone LAN policy 1 then log session-init
set security policies global policy 1000 match source-address any
set security policies global policy 1000 match destination-address any
set security policies global policy 1000 match application any
set security policies global policy 1000 match from-zone any
set security policies global policy 1000 match to-zone any
set security policies global policy 1000 then deny
set security policies global policy 1000 then log session-init
set security policies global policy 1000 then count
set security policies default-policy deny-all
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy SubSrv mode main
set security group-vpn member ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway SubSrv ike-policy SubSrv
set security group-vpn member ike gateway SubSrv server-address 10.17.101.1
set security group-vpn member ike gateway SubSrv server-address 10.17.102.1
set security group-vpn member ike gateway SubSrv server-address 10.17.103.1
set security group-vpn member ike gateway SubSrv server-address 10.17.104.1
set security group-vpn member ike gateway SubSrv local-address 10.18.102.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway SubSrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group-vpn-external-interface
  ge-0/0/1.0
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 recovery-probe
set security ipsec-policy from-zone LAN to-zone WAN ipsec-group-vpn GROUP_ID-0001
```

Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure interfaces, security zones, and security policies.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 description To_LAN
user@host# set ge-0/0/0 unit 0 family inet address 172.16.102.1/24
user@host# set ge-0/0/1 unit 0 description To_SubSrv
user@host# set ge-0/0/1 unit 0 family inet address 10.18.102.1/24

[edit security zones security-zone LAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/0.0

[edit security zones security-zone WAN]
user@host# set host-inbound-traffic system-services ike
user@host# set host-inbound-traffic system-services ssh
user@host# set host-inbound-traffic system-services ping
user@host# set interfaces ge-0/0/1.0

[edit security]
user@host# set address-book global address 172.16.0.0/12 172.16.0.0/12

[edit security policies from-zone LAN to-zone WAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set policy 1 then log session-init

[edit security policies from-zone WAN to-zone LAN]
user@host# set policy 1 match source-address 172.16.0.0/12
user@host# set policy 1 match destination-address 172.16.0.0/12
user@host# set policy 1 match application any
user@host# set policy 1 then permit
user@host# set policy 1 then log session-init

[edit security policies global]
user@host# set policy 1000 match source-address any
user@host# set policy 1000 match destination-address any
user@host# set policy 1000 match application any
user@host# set policy 1000 match from-zone any
user@host# set policy 1000 match to-zone any
user@host# set policy 1000 then deny
user@host# set policy 1000 then log session-init
user@host# set policy 1000 then count
```


Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security policies default-policy deny-all
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy SubSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway SubSrv]
user@host# set ike-policy SubSrv
user@host# set server-address 10.17.101.1
user@host# set server-address 10.17.102.1
user@host# set server-address 10.17.103.1
user@host# set server-address 10.17.104.1
user@host# set local-address 10.18.102.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway SubSrv
user@host# set group-vpn-external-interface ge-0/0/1.0
user@host# set group 1
user@host# set recovery-probe
```

4. Configure the IPsec policy.

```
[edit security ipsec-policy from-zone LAN to-zone WAN]
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    description To_LAN;
    family inet {
      address 172.16.102.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    description To_SubSrv;
    family inet {
      address 10.18.102.1/24;
```

Copyright © 2016, Juniper Networks, Inc.

```
    }
  }
}
[edit]
user@host# show security
address-book {
  global {
    address 172.16.0.0/12 172.16.0.0/12;
  }
}
group-vpn {
  member {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
        authentication-method pre-shared-keys;
        dh-group group14;
        authentication-algorithm sha-256;
        encryption-algorithm aes-256-cbc;
      }
      policy SubSrv {
        mode main;
        proposals PSK-SHA256-DH14-AES256;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
      }
      gateway SubSrv {
        ike-policy SubSrv;
        server-address [ 10.17.101.1 10.17.102.1 10.17.103.1 10.17.104.1 ];
        local-address 10.18.102.1;
      }
    }
    ipsec {
      vpn GROUP_ID-0001 {
        ike-gateway SubSrv;
        group-vpn-external-interface ge-0/0/1.0;
        group 1;
        recovery-probe;
      }
    }
  }
}
ipsec-policy {
  from-zone LAN to-zone WAN {
    ipsec-group-vpn GROUP_ID-0001;
  }
}
policies {
  from-zone LAN to-zone WAN {
    policy 1 {
      match {
        source-address 172.16.0.0/12;
        destination-address 172.16.0.0/12;
        application any;
      }
      then {
        permit;
        log {
```


Copyright © 2016, Juniper Networks, Inc.

```
            session-init;
          }
        }
      }
    }
    from-zone WAN to-zone LAN {
      policy 1 {
        match {
          source-address 172.16.0.0/12;
          destination-address 172.16.0.0/12;
          application any;
        }
        then {
          permit;
          log {
            session-init;
          }
        }
      }
    }
    global {
      policy 1000 {
        match {
          source-address any;
          destination-address any;
          application any;
          from-zone any;
          to-zone any;
        }
        then {
          deny;
          log {
            session-init;
          }
          count;
        }
      }
    }
    default-policy {
      deny-all;
    }
  }
  zones {
    security-zone LAN {
      host-inbound-traffic {
        system-services {
          ike;
          ssh;
          ping;
        }
      }
      interfaces {
        ge-0/0/0.0;
      }
    }
    security-zone WAN {
```

```
        host-inbound-traffic {
            system-services {
                ike;
                ssh;
                ping;
            }
        }
        interfaces {
            ge-0/0/1.0;
        }
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring GM-0003 (MX Series Device)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
    service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet service output service-set GROUP_ID-0001
    service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet address 10.18.103.1/24
set interfaces xe-0/0/2 unit 0 family inet address 172.16.103.1/24
set interfaces ms-0/2/0 unit 0 family inet
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
    authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
    group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
    authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
    encryption-algorithm aes-256-cbc
set security group-vpn member ike policy SubSrv mode main
set security group-vpn member ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway SubSrv ike-policy SubSrv
set security group-vpn member ike gateway SubSrv server-address 10.17.101.1
set security group-vpn member ike gateway SubSrv server-address 10.17.102.1
set security group-vpn member ike gateway SubSrv server-address 10.17.103.1
set security group-vpn member ike gateway SubSrv server-address 10.17.104.1
set security group-vpn member ike gateway SubSrv local-address 10.18.103.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway SubSrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 match-direction output
set security group-vpn member ipsec vpn GROUP_ID-0001 tunnel-mtu 1400
set security group-vpn member ipsec vpn GROUP_ID-0001 df-bit clear
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
    10.17.101.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
    10.17.102.1/32
```


Copyright © 2016, Juniper Networks, Inc.

```
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.103.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.104.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.101.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.102.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.103.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.104.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  source-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  destination-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 then service
set services service-set GROUP_ID-0001 interface-service service-interface ms-0/2/0.0
set services service-set GROUP_ID-0001 ipsec-group-vpn GROUP_ID-0001
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure the interfaces.

```
[edit interfaces]
user@host# set xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet service output service-set
  GROUP_ID-0001 service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet address 10.18.103.1/24
user@host# set xe-0/0/2 unit 0 family inet address 172.16.103.1/24
user@host# set ms-0/2/0 unit 0 family inet
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy SubSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway SubSrv]
user@host# set ike-policy SubSrv
user@host# set server-address 10.17.101.1
user@host# set server-address 10.17.102.1
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set server-address 10.17.103.1
user@host# set server-address 10.17.104.1
user@host# set local-address 10.18.103.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway SubSrv
user@host# set group 1
user@host# set match-direction output
user@host# set tunnel-mtu 1400
user@host# set df-bit clear
```

4. Configure the service filter.

```
[edit firewall family inet service-filter GroupVPN-KS]
user@host# set term inbound-ks from source-address 10.17.101.1/32
user@host# set term inbound-ks from source-address 10.17.102.1/32
user@host# set term inbound-ks from source-address 10.17.103.1/32
user@host# set term inbound-ks from source-address 10.17.104.1/32
user@host# set term inbound-ks then skip
user@host# set term outbound-ks from destination-address 10.17.101.1/32
user@host# set term outbound-ks from destination-address 10.17.102.1/32
user@host# set term outbound-ks from destination-address 10.17.103.1/32
user@host# set term outbound-ks from destination-address 10.17.104.1/32
user@host# set term outbound-ks then skip
user@host# set term GROUP_ID-0001 from source-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 from destination-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 then service
```

5. Configure the service set.

```
[edit services service-set GROUP_ID-0001]
user@host# set interface-service service-interface ms-0/2/0.0
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security**, **show services**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
xe-0/0/1 {
  unit 0 {
    family inet {
      service {
        input {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
        output {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
      }
      address 10.18.103.1/24;
    }
  }
```

 Copyright © 2016, Juniper Networks, Inc.

```
}
xe-0/0/2 {
  unit 0 {
    family inet {
      address 172.16.103.1/24;
    }
  }
}
ms-0/2/0 {
  unit 0 {
    family inet;
  }
}
[edit]
user@host# show security
group-vpn {
  member {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
        authentication-method pre-shared-keys;
        dh-group group14;
        authentication-algorithm sha-256;
        encryption-algorithm aes-256-cbc;
      }
      policy SubSrv {
        mode main;
        proposals PSK-SHA256-DH14-AES256;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
      }
      gateway SubSrv {
        ike-policy SubSrv;
        server-address [ 10.17.101.1 10.17.102.1 10.17.103.1 10.17.104.1 ];
        local-address 10.18.103.1;
      }
    }
    ipsec {
      vpn GROUP_ID-0001 {
        ike-gateway SubSrv;
        group 1;
        match-direction output;
        tunnel-mtu 1400;
        df-bit clear;
      }
    }
  }
}
[edit]
user@host# show services
service-set GROUP_ID-0001 {
  interface-service {
    service-interface ms-0/2/0.0;
  }
  ipsec-group-vpn GROUP_ID-0001;
}
[edit]
user@host# show firewall
```

Copyright © 2016, Juniper Networks, Inc.

```
family inet {
  service-filter GroupVPN-KS {
    term inbound-ks {
      from {
        source-address {
          10.17.101.1/32;
          10.17.102.1/32;
          10.17.103.1/32;
          10.17.104.1/32;
        }
      }
      then skip;
    }
    term outbound-ks {
      from {
        destination-address {
          10.17.101.1/32;
          10.17.102.1/32;
          10.17.103.1/32;
          10.17.104.1/32;
        }
      }
      then skip;
    }
    term GROUP_ID-0001 {
      from {
        source-address {
          172.16.0.0/12;
        }
        destination-address {
          172.16.0.0/12;
        }
      }
      then service;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring GM-0004 (MX Series Device)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet service output service-set GROUP_ID-0001
  service-filter GroupVPN-KS
set interfaces xe-0/0/1 unit 0 family inet address 10.18.104.1/24
set interfaces xe-0/0/2 unit 0 family inet address 172.16.104.1/24
set interfaces ms-0/2/0 unit 0 family inet
```


Copyright © 2016, Juniper Networks, Inc.

```
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-method pre-shared-keys
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256 dh-group
  group14
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  authentication-algorithm sha-256
set security group-vpn member ike proposal PSK-SHA256-DH14-AES256
  encryption-algorithm aes-256-cbc
set security group-vpn member ike policy SubSrv mode main
set security group-vpn member ike policy SubSrv proposals PSK-SHA256-DH14-AES256
set security group-vpn member ike policy SubSrv pre-shared-key ascii-text "$ABC123"
set security group-vpn member ike gateway SubSrv ike-policy SubSrv
set security group-vpn member ike gateway SubSrv server-address 10.17.101.1
set security group-vpn member ike gateway SubSrv server-address 10.17.102.1
set security group-vpn member ike gateway SubSrv server-address 10.17.103.1
set security group-vpn member ike gateway SubSrv server-address 10.17.104.1
set security group-vpn member ike gateway SubSrv local-address 10.18.104.1
set security group-vpn member ipsec vpn GROUP_ID-0001 ike-gateway SubSrv
set security group-vpn member ipsec vpn GROUP_ID-0001 group 1
set security group-vpn member ipsec vpn GROUP_ID-0001 match-direction output
set security group-vpn member ipsec vpn GROUP_ID-0001 tunnel-mtu 1400
set security group-vpn member ipsec vpn GROUP_ID-0001 df-bit clear
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.101.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.102.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.103.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks from source-address
  10.17.104.1/32
set firewall family inet service-filter GroupVPN-KS term inbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.101.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.102.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.103.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks from
  destination-address 10.17.104.1/32
set firewall family inet service-filter GroupVPN-KS term outbound-ks then skip
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  source-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 from
  destination-address 172.16.0.0/12
set firewall family inet service-filter GroupVPN-KS term GROUP_ID-0001 then service
set services service-set GROUP_ID-0001 interface-service service-interface ms-0/2/0.0
set services service-set GROUP_ID-0001 ipsec-group-vpn GROUP_ID-0001
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the Group VPNv2 member:

1. Configure the interfaces.

Copyright © 2016, Juniper Networks, Inc.

```
[edit interfaces]
user@host# set xe-0/0/1 unit 0 family inet service input service-set GROUP_ID-0001
  service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet service output service-set
  GROUP_ID-0001 service-filter GroupVPN-KS
user@host# set xe-0/0/1 unit 0 family inet address 10.18.104.1/24
user@host# set xe-0/0/2 unit 0 family inet address 172.16.104.1/24
user@host# set ms-0/2/0 unit 0 family inet
```

2. Configure the IKE proposal, policy, and gateway.

```
[edit security group-vpn member ike proposal PSK-SHA256-DH14-AES256]
user@host# set authentication-method pre-shared-keys
user@host# set dh-group group14
user@host# set authentication-algorithm sha-256
user@host# set encryption-algorithm aes-256-cbc

[edit security group-vpn member ike policy SubSrv]
user@host# set mode main
user@host# set proposals PSK-SHA256-DH14-AES256
user@host# set pre-shared-key ascii-text "$ABC123"

[edit security group-vpn member ike gateway SubSrv]
user@host# set ike-policy SubSrv
user@host# set server-address 10.17.101.1
user@host# set server-address 10.17.102.1
user@host# set server-address 10.17.103.1
user@host# set server-address 10.17.104.1
user@host# set local-address 10.18.104.1
```

3. Configure the IPsec SA.

```
[edit security group-vpn member ipsec vpn GROUP_ID-0001]
user@host# set ike-gateway SubSrv
user@host# set group 1
user@host# set match-direction output
user@host# set tunnel-mtu 1400
user@host# set df-bit clear
```

4. Configure the service filter.

```
[edit firewall family inet service-filter GroupVPN-KS]
user@host# set term inbound-ks from source-address 10.17.101.1/32
user@host# set term inbound-ks from source-address 10.17.102.1/32
user@host# set term inbound-ks from source-address 10.17.103.1/32
user@host# set term inbound-ks from source-address 10.17.104.1/32
user@host# set term inbound-ks then skip
user@host# set term outbound-ks from destination-address 10.17.101.1/32
user@host# set term outbound-ks from destination-address 10.17.102.1/32
user@host# set term outbound-ks from destination-address 10.17.103.1/32
user@host# set term outbound-ks from destination-address 10.17.104.1/32
user@host# set term outbound-ks then skip
user@host# set term GROUP_ID-0001 from source-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 from destination-address 172.16.0.0/12
user@host# set term GROUP_ID-0001 then service
```

5. Configure the service set.

Copyright © 2016, Juniper Networks, Inc.

```
[edit services service-set GROUP_ID-0001]
user@host# set interface-service service-interface ms-0/2/0.0
user@host# set ipsec-group-vpn GROUP_ID-0001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show security**, **show services**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
xe-0/0/1 {
  unit 0 {
    family inet {
      service {
        input {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
        output {
          service-set GROUP_ID-0001 service-filter GroupVPN-KS;
        }
      }
      address 10.18.104.1/24;
    }
  }
}
xe-0/0/2 {
  unit 0 {
    family inet {
      address 172.16.104.1/24;
    }
  }
}
ms-0/2/0 {
  unit 0 {
    family inet;
  }
}
[edit]
user@host# show security
group-vpn {
  member {
    ike {
      proposal PSK-SHA256-DH14-AES256 {
        authentication-method pre-shared-keys;
        dh-group group14;
        authentication-algorithm sha-256;
        encryption-algorithm aes-256-cbc;
      }
      policy SubSrv {
        mode main;
        proposals PSK-SHA256-DH14-AES256;
        pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
      }
      gateway SubSrv {
```

Copyright © 2016, Juniper Networks, Inc.

```
            ike-policy SubSrv;
            server-address [ 10.17.101.1 10.17.102.1 10.17.103.1 10.17.104.1 ];
            local-address 10.18.104.1;
          }
        }
        ipsec {
          vpn GROUP_ID-0001 {
            ike-gateway SubSrv;
            group 1;
            match-direction output;
            tunnel-mtu 1400;
            df-bit clear;
          }
        }
      }
    }
    [edit]
    user@host# show services
    service-set GROUP_ID-0001 {
      interface-service {
        service-interface ms-0/2/0.0;
      }
      ipsec-group-vpn GROUP_ID-0001;
    }
    [edit]
    user@host# show firewall
    family inet {
      service-filter GroupVPN-KS {
        term inbound-ks {
          from {
            source-address {
              10.17.101.1/32;
              10.17.102.1/32;
              10.17.103.1/32;
              10.17.104.1/32;
            }
          }
          then skip;
        }
        term outbound-ks {
          from {
            destination-address {
              10.17.101.1/32;
              10.17.102.1/32;
              10.17.103.1/32;
              10.17.104.1/32;
            }
          }
          then skip;
        }
        term GROUP_ID-0001 {
          from {
            source-address {
              172.16.0.0/12;
            }
            destination-address {
```


Copyright © 2016, Juniper Networks, Inc.

```
            172.16.0.0/12;
          }
        }
        then service;
      }
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying Server Cluster Operation on page 6765
- Verifying That SAs Are Distributed to Members on page 6767
- Verifying IKE SAs on the Servers on page 6769
- Verifying IPsec SAs on the Servers and Group Members on page 6771
- Verifying IPsec Policies on Group Members on page 6773

Verifying Server Cluster Operation

Purpose Verify that devices in the server cluster recognize peer servers in the group. Ensure that the servers are active and roles in the cluster are properly assigned.

Action From operational mode, enter the **show security group-vpn server server-cluster**, **show security group-vpn server server-cluster detail**, and **show security group-vpn server statistics** commands on the root-server.

```
user@RootSrv> show security group-vpn server server-cluster
Group: GROUP_ID-0001, Group Id: 1
Role: Root-server, Version Number: 2,
  Peer Gateway                    Peer IP         Role                   Status

  SubSrv01                         10.16.101.1   Sub-server             Active
  SubSrv02                         10.16.102.1   Sub-server             Active
  SubSrv03                         10.16.103.1   Sub-server             Active
  SubSrv04                         10.16.104.1   Sub-server             Active
```

```
user@RootSrv> show security group-vpn server server-cluster detail
Group: GROUP_ID-0001, Group Id: 1
Role: Root-server, Version Number: 2

Peer gateway: SubSrv01
  Peer IP: 10.16.101.1, Local IP: 10.10.101.1, VR: default
  Role: Sub-server, Status: Active
  CLUSTER-INIT send:                 0
  CLUSTER-INIT recv:                 1
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
  CLUSTER-UPDATE send:               2
  CLUSTER-UPDATE recv:               0
```

```
  CLUSTER-UPDATE success:            2
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          677
  DPD send fail:                     0
  DPD ACK recv:                      677
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0

Peer gateway: SubSrv02
  Peer IP: 10.16.102.1, Local IP: 10.10.102.1, VR: default
  Role: Sub-server, Status: Active
  CLUSTER-INIT send:                 0
  CLUSTER-INIT recv:                 1
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
  CLUSTER-UPDATE send:               2
  CLUSTER-UPDATE recv:               0
  CLUSTER-UPDATE success:            2
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          676
  DPD send fail:                     0
  DPD ACK recv:                      676
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0
```

```
user@RootSrv> show security group-vpn server statistics
Group: GROUP_ID-0001, Group Id: 1
  Stats:
      Pull Succeeded                : 0
      Pull Failed                   : 0
      Pull Exceed Member Threshold  : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```

From operational mode, enter the **show security group-vpn server server-cluster**, **show security group-vpn server server-cluster detail**, and **show security group-vpn server statistics** commands on each sub-server.

```
user@SubSrv01> show security group-vpn server server-cluster
Group: GROUP_ID-0001, Group Id: 1
Role: Sub-server, Version Number: 2,
  Peer Gateway                    Peer IP         Role                    Status

  RootSrv                         10.10.101.1     Root-server             Active
```


Copyright © 2016, Juniper Networks, Inc.

```
user@SubSrv01> show security group-vpn server server-cluster detail
Group: GROUP_ID-0001, Group Id: 1
Role: Sub-server, Version Number: 2

Peer gateway: RootSrv
  Peer IP: 10.10.101.1, Local IP: 10.16.101.1, VR: default
  Role: Root-server, Status: Active
  CLUSTER-INIT send:                 1
  CLUSTER-INIT recv:                 0
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
  CLUSTER-UPDATE send:               0
  CLUSTER-UPDATE recv:               2
  CLUSTER-UPDATE success:            2
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          812
  DPD send fail:                     0
  DPD ACK recv:                      812
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0
```

```
user@SubSrv01> show security group-vpn server statistics
Group: GROUP_ID-0001, Group Id: 1
  Stats:
      Pull Succeeded               : 4
      Pull Failed                  : 0
      Pull Exceed Member Threshold : 0
      Push Sent                    : 8
      Push Acknowledged            : 8
      Push Unacknowledged          : 0
```

Verifying That SAs Are Distributed to Members

Purpose Verify that the sub-servers have received SAs for distribution to group members and the group members have received the SAs.

Action From operational mode, enter the **show security group-vpn server kek security-associations** and **show security group-vpn server kek security-associations detail** commands on the root-server.

```
user@RootSrv> show security group-vpn server kek security-associations
Index   Life:sec  Initiator cookie  Responder cookie  GroupId
738885  2888      5742c24020056c6a  d6d479543b56404c  1
```

```
user@RootSrv> show security group-vpn server kek security-associations detail
Index 738885, Group Name: GROUP_ID-0001, Group Id: 1
Initiator cookie: 5742c24020056c6a, Responder cookie: d6d479543b56404c
Authentication method: RSA
Lifetime: Expires in 2883 seconds, Activated
Rekey in 2373 seconds
```

Copyright © 2016, Juniper Networks, Inc.

```
  Algorithms:
   Sig-hash             : sha256
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Server Member Communication: Unicast
  Retransmission Period: 10, Number of Retransmissions: 2
  Group Key Push sequence number: 0

PUSH negotiations in progress: 0
```

From operational mode, enter the **show security group-vpn server kek security-associations** and **show security group-vpn server kek security-associations detail** commands on each sub-server.

```
user@SubSrv01> show security group-vpn server kek security-associations
Index   Life:sec  Initiator cookie  Responder cookie  GroupId
738885  1575      5742c24020056c6a  d6d479543b56404c  1

user@SubSrv01> show security group-vpn server kek security-associations detail
Index 738879, Group Name: GROUP_ID-0001, Group Id: 1
Initiator cookie: 114e4a214891e42f, Responder cookie: 4b2848d14372e5bd
Authentication method: RSA
Lifetime: Expires in 4186 seconds, Activated
Rekey in 3614 seconds
  Algorithms:
   Sig-hash             : sha256
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Server Member Communication: Unicast
  Retransmission Period: 10, Number of Retransmissions: 2
  Group Key Push sequence number: 0

PUSH negotiations in progress: 0
```

From operational mode, enter the **show security group-vpn member kek security-associations** and **show security group-vpn member kek security-associations detail** commands on each group member.

For SRX or vSRX group members:

```
user@GM-0001> show security group-vpn server kek security-associations
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
5455799 10.17.101.1     1466      5742c24020056c6a  d6d479543b56404c  1

user@GM-0001> show security group-vpn server kek security-associations detail
  Index 5455799, Group Id: 1
  Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.101.1, GDOI Server: 10.17.101.1
  Initiator cookie: 5742c24020056c6a, Responder cookie: d6d479543b56404c
  Lifetime: Expires in 1464 seconds
  Group Key Push Sequence number: 0

  Algorithms:
```


Copyright © 2016, Juniper Networks, Inc.

```
   Sig-hash              : hmac-sha256-128
   Encryption            : aes256-cbc
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Stats:
      Push received           :   0
      Delete received         :   0
```

For MX group members:

```
user@GM-0003> show security group-vpn member kek security-associations
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
5184329 10.17.101.1     1323      5742c24020056c6a  d6d479543b56404c  1

user@GM-0003> show security group-vpn member kek security-associations detail
  Index 5184329, Group Id: 1
  Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.103.1, GDOI Server: 10.17.101.1
  Initiator cookie: 5742c24020056c6a, Responder cookie: d6d479543b56404c
  Lifetime: Expires in 1321 seconds
  Group Key Push Sequence number: 0

  Algorithms:
   Sig-hash              : hmac-sha256-128
   Encryption            : aes256-cbc
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Stats:
      Push received           :   0
      Delete received         :   0
```

Verifying IKE SAs on the Servers

Purpose Display IKE security associations (SAs) on the servers.

Action From operational mode, enter the **show security group-vpn server ike security-associations** and **show security group-vpn server ike security-associations detail** commands on the root-server.

```
user@RootSrv> show security group-vpn server ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

738880  UP     2221001e980eb08b  5af00708f5da289c  Main           10.16.104.1
738881  UP     59e8c1d328b1d9fd  d63e823fb8be1f22  Main           10.16.101.1
738883  UP     9cb3a49c6771819e  8df3be8c9ddeb2a7  Main           10.16.102.1
738882  UP     9a8a75f05a1384c5  c6d58696c896b730  Main           10.16.103.1

user@RootSrv> show security group-vpn server ike security-associations detail
IKE peer 10.16.101.1, Index 738881, Gateway Name: SubSrv01
  Role: Responder, State: UP
  Initiator cookie: 59e8c1d328b1d9fd, Responder cookie: d63e823fb8be1f22
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.10.101.1:848, Remote: 10.16.101.1:848
```

Copyright © 2016, Juniper Networks, Inc.

```
  Lifetime: Expires in 21890 seconds
  Peer ike-id: 10.16.101.1
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
   Input  bytes  :              150112
   Output bytes  :              153472
   Input  packets:                1387
   Output packets:                1387
  Flags: IKE SA is created
IKE peer 10.16.102.1, Index 738883, Gateway Name: SubSrv02
  Role: Responder, State: UP
  Initiator cookie: 9cb3a49c6771819e, Responder cookie: 8df3be8c9ddeb2a7
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.10.102.1:848, Remote: 10.16.102.1:848
  Lifetime: Expires in 21899 seconds
  Peer ike-id: 10.16.102.1
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
   Input  bytes  :              149788
   Output bytes  :              153148
   Input  packets:                1384
   Output packets:                1384
  Flags: IKE SA is created
```

From operational mode, enter the **show security group-vpn server ike security-associations** and **show security group-vpn server ike security-associations detail** commands on each sub-server.

```
user@SubSrv01> show security group-vpn server ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

738878  UP     59e8c1d328b1d9fd  d63e823fb8be1f22  Main           10.10.101.1

user@SubSrv01> show security group-vpn server ike security-associations detail
IKE peer 10.10.101.1, Index 738878, Gateway Name: RootSrv
  Role: Initiator, State: UP
  Initiator cookie: 59e8c1d328b1d9fd, Responder cookie: d63e823fb8be1f22
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 10.16.101.1:848, Remote: 10.10.101.1:848
  Lifetime: Expires in 20589 seconds
  Peer ike-id: 10.10.101.1
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha256-128
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
```


Copyright © 2016, Juniper Networks, Inc.

```
   Input  bytes  :               181444
   Output bytes  :               178084
   Input  packets:                 1646
   Output packets:                 1646
  Flags: IKE SA is created
```

Verifying IPsec SAs on the Servers and Group Members

Purpose Display IPsec security associations (SAs) on the servers and group members.

Action From operational mode, enter the **show security group-vpn server ipsec security-associations** and **show security group-vpn server ipsec security-associations detail** commands on the root-server.

```
user@RootSrv> show security group-vpn server ipsec security-associations
Group: GROUP_ID-0001, Group Id: 1
  Total IPsec SAs: 1
  IPsec SA          Algorithm        SPI              Lifetime
  GROUP_ID-0001     ESP:aes-256/sha256 dddef414       2773
```

```
user@RootSrv> show security group-vpn server ipsec security-associations detail
Group: GROUP_ID-0001, Group Id: 1
Total IPsec SAs: 1
  IPsec SA: GROUP_ID-0001
    Protocol: ESP, Authentication: sha256, Encryption: aes-256
    Anti-replay: D3P enabled
    SPI: dddef414
    Lifetime: Expires in 1670 seconds, Activated
    Rekey in 1160 seconds
    Policy Name: 1
      Source: 172.16.0.0/12
      Destination: 172.16.0.0/12
      Source Port: 0
      Destination Port: 0
      Protocol: 0
```

From operational mode, enter the **show security group-vpn server ipsec security-associations** and **show security group-vpn server ipsec security-associations detail** commands on each sub-server.

```
user@SubSrv01> show security group-vpn server ipsec security-associations
Group: GROUP_ID-0001, Group Id: 1
  Total IPsec SAs: 1
  IPsec SA          Algorithm        SPI              Lifetime
  GROUP_ID-0001     ESP:aes-256/sha256 dddef414       1520
```

```
user@SubSrv01> show security group-vpn server ipsec security-associations detail
Group: GROUP_ID-0001, Group Id: 1
Total IPsec SAs: 1
  IPsec SA: GROUP_ID-0001
    Protocol: ESP, Authentication: sha256, Encryption: aes-256
    Anti-replay: D3P enabled
    SPI: dddef414
    Lifetime: Expires in 1518 seconds, Activated
    Rekey in 1230 seconds
    Policy Name: 1
      Source: 172.16.0.0/12
      Destination: 172.16.0.0/12
      Source Port: 0
```

Copyright © 2016, Juniper Networks, Inc.

```
      Destination Port: 0
      Protocol: 0
```

From operational mode, enter the **show security group-vpn member ipsec security-associations** and **show security group-vpn member ipsec security-associations detail** commands on each group member

For SRX or vSRX group members:

```
user@GM-0001> show security group-vpn member ipsec security-associations
  Total active tunnels: 1
  ID    Server           Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>49152 10.17.101.1    848   ESP:aes-256/sha256-128 dddef414 1412/ unlim 1 root

user@GM-0001> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.101.1, GDOI Server: 10.17.101.1
  Group Id: 1
  Routing Instance: default
  Recovery Probe: Enabled
  DF-bit: clear
  Stats:
      Pull Succeeded             :   1
      Pull Failed                :   0
      Pull Timeout               :   0
      Pull Aborted               :   0
      Push Succeeded             :   2
      Push Failed                :   0
      Server Failover            :   0
      Delete Received            :   0
      Exceed Maximum Keys(4)     :   0
      Exceed Maximum Policies(10):   0
      Unsupported Algo           :   0
  Flags:
      Rekey Needed:   no

    List of policies received from server:
    Tunnel-id: 49152
      Source IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)
      Destination IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)

      Direction: bi-directional, SPI: dddef414
      Protocol: ESP, Authentication: sha256-128, Encryption: aes-256
      Hard lifetime: Expires in 1409 seconds, Activated
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 1193 seconds
      Mode: Tunnel, Type: Group VPN, State: installed
      Anti-replay service: D3P enabled
```

For MX group members:

```
user@GM-0003> show security group-vpn member ipsec security-associations
  Total active tunnels: 1
  ID    Server           Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>10001 10.17.101.1    848   ESP:aes-256/sha256-128 dddef414 1308/ unlim 1 root

user@GM-0003> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: GROUP_ID-0001
  Local Gateway: 10.18.103.1, GDOI Server: 10.17.101.1
  Group Id: 1
  Rule Match Direction: output,  Tunnel-MTU: 1400
```


Copyright © 2016, Juniper Networks, Inc.

```
Routing Instance: default
DF-bit: clear
Stats:
    Pull Succeeded            :   1
    Pull Failed               :   0
    Pull Timeout              :   0
    Pull Aborted              :   0
    Push Succeeded            :   2
    Push Failed               :   0
    Server Failover           :   0
    Delete Received           :   0
    Exceed Maximum Keys(4)    :   0
    Exceed Maximum Policies(1):   0
    Unsupported Algo          :   0
Flags:
    Rekey Needed:   no

  List of policies received from server:
  Tunnel-id: 10001
    Source IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)
    Destination IP: ipv4_subnet(any:0,[0..7]=172.16.0.0/12)

    Direction: bi-directional, SPI: dddef414
    Protocol: ESP, Authentication: sha256-128, Encryption: aes-256
    Hard lifetime: Expires in 1305 seconds, Activated
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1087 seconds
    Mode: Tunnel, Type: Group VPN, State: installed
    Anti-replay service: D3P enabled
```

Verifying IPsec Policies on Group Members

Purpose Display the IPsec policy on an SRX or vSRX group member.

NOTE: This command is not available for MX Series group members.

Copyright © 2016, Juniper Networks, Inc.

Action From operational mode, enter the **show security group-vpn member policy** command on SRX or vSRX group members.

```
user@GM-0001> show security group-vpn member policy
Group VPN Name: GROUP_ID-0001, Group Id: 1
From-zone: LAN, To-zone: WAN
  Tunnel-id: 49152, Policy type: Secure
    Source      : IP <172.16.0.0 - 172.31.255.255>, Port <0 - 65535>, Protocol
<0>
    Destination : IP <172.16.0.0 - 172.31.255.255>, Port <0 - 65535>, Protocol
<0>

  Tunnel-id: 63488, Policy type: Fail-close
    Source      : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>

    Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
```

Related Documentation

- Group VPNv2 Configuration Overview on page 6652
- Understanding Group VPNv2 Server Clusters on page 6699
- Configuring Group VPNs in Group VPNv2 on Routing Device

Copyright © 2016, Juniper Networks, Inc.

PART 88

Monitoring and Improving VPN Traffic Performance

- Configuring VPN Monitoring Features on page 6777
- Improving IPsec VPN Traffic Performance on page 6789

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 278

Configuring VPN Monitoring Features

- Understanding VPN Alarms and Auditing on page 6777
- Example: Setting an Audible Alert as Notification of a Security Alarm on page 6779
- Example: Generating Security Alarms in Response to Potential Violations on page 6780
- Understanding VPN Monitoring and DPD on page 6782
- Understanding Dead Peer Detection on page 6783
- Understanding VPN Monitoring on page 6785
- Understanding Global SPI and VPN Monitoring Features on page 6785
- Example: Configuring Global SPI and VPN Monitoring Features on page 6786
- Understanding Tunnel Events on page 6786

Understanding VPN Alarms and Auditing

Configure the following command to enable security event logging during the initial set up of the device.

set security log cache

The administrators (audit, cryptographic, IDS and security) cannot modify the security event logging configuration if the above command is configured and each administrator role is configured to have a distinct, unique set of privileges apart from all other administrative roles.

Alarms are triggered by a VPN failure. A VPN alarm is generated when the system monitors any of the following audited events:

- **Authentication failures**—You can configure the device to generate a system alarm when the packet authentication failures reaches a specified number.
- **Encryption and decryption failures**—You can configure the device to generate a system alarm when encryption or decryption failures exceed a specified number.
- **IKE Phase 1 and IKE Phase 2 failures**—Internet Key Exchange (IKE) Phase 1 negotiations are used to establish IKE security associations (SAs). These SAs protect the IKE Phase 2 negotiations. You can configure the device to generate a system alarm when IKE Phase 1 or IKE Phase 2 failures exceed a specified number.

Copyright © 2016, Juniper Networks, Inc.

- **Self-test failures**—Self tests are tests that a device runs upon power on or reboot to verify whether security software is implemented correctly on your device.

  Self-tests ensure the correctness of cryptographic algorithms. The Junos-FIPS image performs self-tests automatically upon power-on, and continuously for key-pair generation. In either domestic or FIPS images, self-tests may be configured to be performed according to a defined schedule, upon demand or immediately after key generation.

  You can configure the device to generate a system alarm when a self-test failure occurs.

- **IDP flow policy attacks**—An intrusion detection and prevention (IDP) policy allows you to enforce various attack detection and prevention techniques on network traffic. You can configure the device to generate a system alarm when IDP flow policy violations occur.

- **Replay attacks**—A replay attack is a network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed. You can configure the device to generate a system alarm when a replay attack occurs.

The syslog messages are included in the following cases:

- Failed symmetric key generation
- Failed asymmetric key generation
- Failed manual key distribution
- Failed automated key distribution
- Failed key destruction
- Failed key handling and storage
- Failed data encryption or decryption
- Failed signature
- Failed key agreement
- Failed cryptographic hashing
- IKE failure
- Failed authentication of the received packets
- Decryption error due to invalid padding content
- Mismatch in the length specified in the alternative subject field of the certificate received from a remote VPN peer device.

Alarms are raised based on syslog messages. Every failure is logged, but an alarm is generated only when a threshold is reached.

To view the alarm information, run the **show security alarms** command. The violation count and the alarm do not persist across system reboots. After a reboot, the violation count resets to zero, and the alarm is cleared from the alarm queue.

Copyright © 2016, Juniper Networks, Inc.

After appropriate actions have been taken, you can clear the alarm. The alarm remains in the queue until you clear it (or until you reboot the device). To clear the alarm, run the **clear security alarms** command.

Related Documentation

- Example: Setting an Audible Alert as Notification of a Security Alarm on page 6779
- Example: Generating Security Alarms in Response to Potential Violations on page 6780

Example: Setting an Audible Alert as Notification of a Security Alarm

This example shows how to configure a device to generate a system alert beep when a new security event occurs. By default, alarms are not audible.

- Requirements on page 6779
- Overview on page 6779
- Configuration on page 6779
- Verification on page 6779

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set an audible beep to be generated in response to a security alarm.

Configuration

Step-by-Step Procedure

To set an audible alarm:

1. Enable security alarms.

   ```
   [edit]
   user@host# edit security alarms
   ```

2. Specify that you want to be notified of security alarms with an audible beep.

   ```
   [edit security alarms]
   user@host# set audible
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit security alarms]
   user@host# commit
   ```

Verification

To verify the configuration is working properly, enter the **show security alarms detail** command.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

Example: Generating Security Alarms in Response to Potential Violations

This example shows how to configure the device to generate a system alarm when a potential violation occurs. By default, no alarm is raised when a potential violation occurs.

- Requirements on page 6780
- Overview on page 6780
- Configuration on page 6780
- Verification on page 6782

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure an alarm to be raised when:

- The number of authentication failures exceeds 6.
- The cryptographic self-test fails.
- The non-cryptographic self-test fails.
- The key generation self-test fails.
- The number of encryption failures exceeds 10.
- The number of decryption failures exceeds 1.
- The number of IKE Phase 1 failures exceeds 10.
- The number of IKE Phase 2 failure exceeds 1.
- A replay attack occurs.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security alarms potential-violation authentication 6
set security alarms potential-violation cryptographic-self-test
set security alarms potential-violation non-cryptographic-self-test
set security alarms potential-violation key-generation-self-test
set security alarms potential-violation encryption-failures threshold 10
set security alarms potential-violation decryption-failures threshold 1
set security alarms potential-violation ike-phase1-failures threshold 10
set security alarms potential-violation ike-phase2-failures threshold 1
set security alarms potential-violation replay-attacks
```


Copyright © 2016, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure alarms in response to potential violations:

1. Enable security alarms.

   ```
   [edit]
   user@host# edit security alarms
   ```

2. Specify that an alarm should be raised when an authentication failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set authentication 6
   ```

3. Specify that an alarm should be raised when a cryptographic self-test failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set cryptographic-self-test
   ```

4. Specify that an alarm should be raised when a non-cryptographic self-test failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set non-cryptographic-self-test
   ```

5. Specify that an alarm should be raised when a key generation self-test failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set key-generation-self-test
   ```

6. Specify that an alarm should be raised when an encryption failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set encryption-failures threshold 10
   ```

7. Specify that an alarm should be raised when a decryption failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set decryption-failures threshold 1
   ```

8. Specify that an alarm should be raised when an IKE Phase 1 failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set ike-phase1-failures threshold 10
   ```

9. Specify that an alarm should be raised when an IKE Phase 2 failure occurs.

   ```
   [edit security alarms potential-violation]
   user@host# set ike-phase2-failures threshold 1
   ```

10. Specify that an alarm should be raised when a replay attack occurs.

    ```
    [edit security alarms potential-violation]
    user@host# set replay-attacks
    ```

Results

From configuration mode, confirm your configuration by entering the **show security alarms** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
```

Copyright © 2016, Juniper Networks, Inc.

```
user@host# show security alarms
potential-violation {
    authentication 6;
    cryptographic-self-test;
    decryption-failures {
        threshold 1;
    }
    encryption-failures {
        threshold 10;
    }
    ike-phase1-failures {
        threshold 10;
    }
    ike-phase2-failures {
        threshold 1;
    }
    key-generation-self-test;
    non-cryptographic-self-test;
    replay-attacks;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, from operational mode, enter the **show security alarms** command.

Related Documentation

- Understanding VPN Alarms and Auditing on page 6777
- Example: Setting an Audible Alert as Notification of a Security Alarm on page 6779

Understanding VPN Monitoring and DPD

VPN monitoring and dead peer detection (DPD) are features available on SRX Series devices to verify the availability of VPN peer devices. This section compares the operation and configuration of these features.

NOTE: The SRX Series device responds to DPD messages sent by VPN peers even if DPD is not configured on the device. You can configure the SRX Series device to initiate DPD messages to VPN peers. You can also configure DPD and VPN monitoring to operate simultaneously on the same SRX Series device, although the number of peers that can be monitored with either method is reduced.

VPN monitoring is a Junos OS mechanism that monitors only Phase 2 security associations (SAs). VPN monitoring is enabled on a per-VPN basis with the **vpn-monitor** statement at the [**edit security ipsec vpn *vpn-name***] hierarchy level. The destination IP and source interface must be specified. The **optimized** option enables the device to use traffic patterns as evidence of peer liveliness; ICMP requests are suppressed.

Copyright © 2016, Juniper Networks, Inc.

VPN monitoring options are configured with the **vpn-monitor-options** statement at the [**edit security ipsec**] hierarchy level. These options apply to all VPNs for which VPN monitoring is enabled. Options you can configure include the interval at which ICMP requests are sent to the peer (the default is 10 seconds) and the number of consecutive ICMP requests sent without receiving a response before the peer is considered unreachable (the default is 10 consecutive requests).

DPD is an implementation of RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*. It operates at the IKE level and monitors the peer based on both IKE and IPsec traffic activity.

DPD is configured on an individual IKE gateway with the **dead-peer-detection** statement at the [**edit security ike gateway *gateway-name***] hierarchy level. You can configure DPD modes of operation. The default (optimized) mode sends DPD messages to the peer if there is no incoming IKE or IPsec traffic within a configured interval after the local device sends outgoing packets to the peer. Other configurable options include the interval at which DPD messages are sent to the peer (the default is 10 seconds) and the number of consecutive DPD messages sent without receiving a response before the peer is considered unavailable (the default is five consecutive requests).

Related Documentation

- Understanding Dead Peer Detection on page 6783
- IPsec VPN Overview on page 6041

Understanding Dead Peer Detection

Dead peer detection (DPD) is a method that network devices use to verify the current existence and availability of other peer devices.

You can use DPD as an alternative to VPN monitoring. However, you cannot use both features simultaneously. VPN monitoring applies to an individual IPsec VPN, while DPD is configured only in an individual IKE gateway context.

A device performs DPD verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer. The device sends an R-U-THERE message only if it has not received any traffic from the peer during a specified DPD interval. If the device receives an R-U-THERE-ACK message from the peer during this interval, it considers the peer alive. If the device receives traffic on the tunnel from the peer, it resets its R-U-THERE message counter for that tunnel, thus starting a new interval. If the device does not receive an R-U-THERE-ACK message during the interval, it considers the peer dead. When the device changes the status of a peer device to be dead, the device removes the Phase 1 security association (SA) and all Phase 2 SAs for that peer.

The following DPD modes are supported on the SRX Series devices:

- Optimized—R-U-THERE messages are triggered if there is no incoming IKE or IPsec traffic within a configured interval after the device sends outgoing packets to the peer. This is the default mode.

Copyright © 2016, Juniper Networks, Inc.

- Probe idle tunnel—R-U-THERE messages are triggered if there is no incoming or outgoing IKE or IPsec traffic within a configured interval. R-U-THERE messages are sent periodically to the peer until there is traffic activity. This mode helps in early detection of a downed peer and makes the tunnel available for data traffic.
- Always send—R-U-THERE messages are sent at configured intervals regardless of traffic activity between the peers.

NOTE: **We recommend that the probe idle tunnel mode be used instead of the always-send mode.**

DPD timers are active as soon as the Phase 1 SA is established. The DPD behavior is the same for both IKEv1 and IKEv2 protocols.

You can configure the following DPD parameters:

- The interval parameter specifies the amount of time (expressed in seconds) the device waits for traffic from its peer before sending an R-U-THERE message. The default interval is 10 seconds, with a permissible range of 10 to 60 seconds.
- The threshold parameter specifies the maximum number of times to send the R-U-THERE message without a response from the peer before considering the peer dead. The default number of transmissions is five times, with a permissible range of 1 to 5 retries.

Note the following considerations before configuring DPD:

- When a DPD configuration is added to an existing gateway with active tunnels, R-U-THERE messages are started without clearing Phase 1 or Phase 2 SAs.
- When a DPD configuration is deleted from an existing gateway with active tunnels, R-U-THERE messages are stopped for the tunnels. IKE and IPsec SAs are not affected.
- Modifying any DPD configuration option such as the mode, interval, or threshold values updates the DPD operation without clearing Phase 1 or Phase 2 SAs.
- If the IKE gateway is configured with DPD and VPN but the option to establish tunnels immediately is not configured, DPD does not initiate Phase 1 negotiation.
- If the IKE gateway is configured with multiple peer IP addresses and DPD but Phase 1 SA fails to be established to the first peer IP address, a Phase 1 SA is attempted with the next peer IP address. DPD is active only after a Phase 1 SA is established.
- If the IKE gateway is configured with multiple peer IP addresses and DPD but DPD fails with the current peer's IP address, the Phase 1 and Phase 2 SAs are cleared and a failover to the next peer IP address is triggered.
- More than one Phase 1 or Phase 2 SA may exist with the same peer because of simultaneous negotiations. In this case, R-U-THERE messages are sent on all Phase 1 SAs. Failure to receive DPD responses for the configured number of consecutive times clears the Phase 1 SA and the associated Phase 2 SA (for IKEv2 only).

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Understanding VPN Monitoring and DPD on page 6782
- Example: Configuring a Policy-Based VPN with Both an Initiator and a Responder Behind a NAT Device on page 6272

Understanding VPN Monitoring

VPN monitoring uses ICMP echo requests (or pings) to determine if a VPN tunnel is up. When VPN monitoring is enabled, the security device sends pings through the VPN tunnel to the peer gateway or to a specified destination at the other end of the tunnel. Pings are sent by default at intervals of 10 seconds for up to 10 consecutive times. If no reply is received after 10 consecutive pings, the VPN is considered to be down and the IPsec security association (SA) is cleared.

VPN monitoring is enabled for a specified VPN by configuring the **vpn-monitor** option at the [**edit security ipsec vpn *vpn-name***] hierarchy level. The peer gateway's IP address is the default destination; however, you can specify a different destination IP address (such as a server) at the other end of the tunnel. The local tunnel endpoint is the default source interface, but you can specify a different interface name.

The VPN monitoring **optimized** option sends pings only when there is outgoing traffic and no incoming traffic through the VPN tunnel. If there is incoming traffic through the VPN tunnel, the security device considers the tunnel to be active and does not send pings to the peer. Configuring the **optimized** option can save resources on the security device because pings are only sent when peer liveliness needs to be determined. Sending pings can also activate costly backup links that would otherwise not be used.

You can configure the interval at which pings are sent and the number of consecutive pings that are sent without a reply before the VPN is considered to be down. These are configured with the **interval** and **threshold** options, respectively, at the [**edit security ipsec vpn-monitor-options**] hierarchy level.

NOTE: VPN monitoring may cause tunnel flapping in some VPN environments if ping packets are not accepted by the peer based on the packet's source or destination IP address.

Related Documentation

- Understanding VPN Monitoring and DPD on page 6782

Understanding Global SPI and VPN Monitoring Features

You can monitor and maintain the efficient operation of your VPN using the following global VPN features:

- SPI—Peers in a security association (SA) can become unsynchronized when one of the peers fails. For example, if one of the peers reboots, it might send an incorrect security parameter index (SPI). You can enable the device to detect such an event and resynchronize the peers by configuring the bad SPI response feature.

Copyright © 2016, Juniper Networks, Inc.

- VPN monitoring—You can use the global VPN monitoring feature to periodically send Internet Control Message Protocol (ICMP) requests to the peer to determine if the peer is reachable.

Related Documentation

- IPsec VPN Overview on page 6041
- Example: Configuring Global SPI and VPN Monitoring Features on page 6786

Example: Configuring Global SPI and VPN Monitoring Features

- Requirements on page 6786
- Overview on page 6786
- Configuration on page 6786

Requirements

Before you begin, understand global SPI and VPN monitoring features. See "Understanding Global SPI and VPN Monitoring Features" on page 6785.

Overview

In this example, you configure the device to detect and respond five times to a bad IPsec SPI before deleting the SA and initiating a new one. You also configure the device to monitor the VPN by sending ICMP requests to the peer every 15 seconds, and to declare the peer unreachable after 15 unsuccessful pings.

Configuration

Step-by-Step Procedure

To configure global VPN settings in the CLI editor:

1. Specify global VPN settings.

```
[edit]
user@host# set security ike respond-bad-spi 5
user@host# set security ipsec vpn-monitor-options interval 15 threshold 15
```

Related Documentation

- Example: Configuring a Policy-Based VPN on page 6222
- Example: Configuring a Route-Based VPN on page 6074

Understanding Tunnel Events

When there is a network problem related to a VPN, after the tunnel comes up only the tunnel status is tracked. Many issues can occur before the tunnel comes up. Hence, instead of tracking only the tunnel status, tunnel down issues, or negotiation failures, successful events such as successful IPsec SA negotiations, IPsec rekey, and IKE SA rekeys are now tracked. These events are called tunnel events.

For Phase 1 and Phase 2, negotiation events for a given tunnel are tracked along with the events that occur in external daemons like AUTHD or PKID. When a tunnel event occurs

Copyright © 2016, Juniper Networks, Inc.

multiple times, only one entry is maintained with the updated time and the number of times that event occurred.

Overall, 16 events are tracked: eight events for Phase 1 and eight events for Phase 2. Some events can reoccur and fill up the event memory, resulting in important events being removed. To avoid overwriting, an event is not stored unless a tunnel is down.

The following special events fall into this category:

- Lifetime in kilobytes expired for IPsec SA
- Hard lifetime of IPsec SA expired
- IPsec SA delete payload received from peer, corresponding IPsec SAs cleared
- Cleared unused redundant backup IPsec SA pairs
- IPsec SAs cleared as corresponding IKE SA deleted

AutoVPN tunnels are created and removed dynamically and consequently tunnel events corresponding to these tunnels are short lived. Sometimes these tunnel events cannot be associated with any tunnel so system logging is used for debugging instead.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 279

Improving IPsec VPN Traffic Performance

- Understanding VPN Session Affinity on page 6789
- Enabling VPN Session Affinity on page 6791
- Accelerating the IPsec VPN Traffic Performance on page 6792

Understanding VPN Session Affinity

VPN session affinity occurs when a clear-text session is located in a Services Processing Unit (SPU) that is different from the SPU where the IPsec tunnel session is located. The goal of VPN session affinity is to locate the clear-text and IPsec tunnel session in the same SPU.

Without VPN session affinity, a clear-text session created by a flow might be located in one SPU and the tunnel session created by IPsec might be located in another SPU. An SPU to SPU forward or hop is needed to route clear-text packets to the IPsec tunnel.

By default, VPN session affinity is disabled on SRX Series devices. When VPN session affinity is enabled, a new clear-text session is placed on the same SPU as the IPsec tunnel session. Existing clear-text sessions are not affected.

Junos OS Release 15.1X49-D10 introduces the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) for SRX5400, SRX5600, and SRX5800 devices.

The SRX5K-MPC (IOC2) and the IOC3 support VPN session affinity through improved flow module and session cache. With IOCs, the flow module creates sessions for IPsec tunnel-based traffic before encryption and after decryption on its tunnel-anchored SPU and installs the session cache for the sessions so that the IOC can redirect the packets to the same SPU to minimize packet forwarding overhead. Express Path (previously known as services offloading) traffic and NP cache traffic share the same session cache table on the IOCs.

Enabling VPN session affinity can improve VPN throughput under the following traffic conditions:

- A number of IPsec tunnels are needed and they are distributed evenly among SPUs. If IPsec tunnels are already concentrated on several SPUs, then enabling VPN session affinity allows all clear-text SPUs to also use those SPUs. This can cause those SPUs to be overutilized while other SPUs might be underutilized.

Copyright © 2016, Juniper Networks, Inc.

To display active tunnel sessions on SPUs, use the **show security ipsec security-association** command and specify the Flexible PIC Concentrator (FPC) and Physical Interface Card (PIC) slots that contain the SPU. For example:

```
user@host> show security ipsec security-association fpc 3 pic 0
    Total active tunnels: 1
    ID   Algorithm      SPI     Life:sec/kb  Mon vsys Port  Gateway
    <131073 ESP:aes-128/sha1 18c4fd00 491/  128000 - root 500   23.0.21.11
    >131073 ESP:aes-128/sha1 188c0750 491/  128000 - root 500   23.0.21.11
```

- Clear-text sessions passing through the tunnels should be at the highest volume for the longest periods of time as possible. Applying VPN session affinity to clear-text sessions of small volumes and short periods (for example, DNS sessions) will decrease the effect of session affinity and might even have a negative impact on VPN throughput under certain conditions.

NOTE: You need to evaluate the tunnel distribution and traffic patterns in your network to determine if VPN session affinity should be enabled.

The VPN session affinity limitations are as follows:

- Traffic across logical systems is not supported.
- If there is a route change, established clear-text sessions remain on an SPU and traffic is rerouted if possible. Sessions created after the route change may be set up on a different SPU.
- VPN session affinity only affects self traffic that terminates on the device (also known as host-inbound traffic); self traffic that originates from the device (also known as host-outbound traffic) is not affected.
- Multicast replication and forwarding performance is not affected.

Related Documentation

- Enabling VPN Session Affinity on page 6791
- SRX5000 Line Devices Processing Overview on page 1671
- Understanding Session Cache on page 1713
- Express Path Overview on page 1717
- Example: Enabling Express Path in Security Policies on page 1733
- Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path on page 1739

Copyright © 2016, Juniper Networks, Inc.

Enabling VPN Session Affinity

By default, VPN session affinity is disabled on SRX Series devices. Enabling VPN session affinity can improve VPN throughput under certain conditions. This section describes how to use the CLI to enable VPN session affinity.

Determine if clear-text sessions are being forwarded to IPsec tunnel sessions on a different SPU. Use the **show security flow session** command to display session information about clear-text sessions.

```
user@host> show security flow session
  Flow Sessions on FPC3 PIC0:

  Session ID: 60000001, Policy name: N/A, Timeout: N/A, Valid
   In: 23.0.21.11/6204 --> 23.0.21.6/41264;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 60000002, Policy name: N/A, Timeout: N/A, Valid
   In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 60000003, Policy name: self-traffic-policy/1, Timeout: 58, Valid
   In: 23.0.21.6/500 --> 23.0.21.11/500;udp, If: .local..0, Pkts: 105386, Bytes: 12026528
   Out: 23.0.21.11/500 --> 23.0.21.6/500;udp, If: ge-0/0/2.0, Pkts: 106462, Bytes: 12105912

  Session ID: 60017354, Policy name: N/A, Timeout: 1784, Valid
   In: 0.0.0.0/0 --> 0.0.0.0/0;0, If: N/A, Pkts: 0, Bytes: 0
   Out: 26.0.20.156/23 --> 22.0.20.155/53051;tcp, If: N/A, Pkts: 0, Bytes: 0
  Total sessions: 4

  Flow Sessions on FPC6 PIC0:

  Session ID: 120000001, Policy name: N/A, Timeout: N/A, Valid
   In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 120000002, Policy name: N/A, Timeout: N/A, Valid
   In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 120031730, Policy name: default-policy-00/2, Timeout: 1764, Valid
   In: 22.0.20.155/53051 --> 26.0.20.156/23;tcp, If: ge-0/0/1.0, Pkts: 44, Bytes: 2399
   Out: 26.0.20.156/23 --> 22.0.20.155/53051;tcp, If: st0.0, Pkts: 35, Bytes: 2449
  Total sessions: 3
```

In the example, there is a tunnel session on FPC 3, PIC 0 and a clear-text session on FPC 6, PIC 0. A forwarding session (session ID 60017354) is set up on FPC 3, PIC 0.

NOTE: Junos OS Release 15.1X49-D10 introduces session affinity support on the IOCs (SRX5K-MPC [IOC2], SRX5K-MPC3-100G10G [IOC3], and SRX5K-MPC3-40G10G [IOC3]). You can enable session affinity for the IPsec tunnel session on the IOC FPCs.

To enable VPN session affinity:

1. In configuration mode, use the **set** command to enable VPN session affinity.

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# set security flow load-distribution session-affinity ipsec
```

2. Check your changes to the configuration before committing.

```
[edit]
user@host# commit check
```

3. Commit the configuration.

```
[edit]
user@host# commit
```

After enabling VPN session affinity, use the **show security flow session** command to display session information about clear-text sessions.

```
user@host> show security flow session
  Flow Sessions on FPC3 PIC0:

  Session ID: 60000001, Policy name: N/A, Timeout: N/A, Valid
    In: 23.0.21.11/6352 --> 23.0.21.6/7927;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 60000002, Policy name: N/A, Timeout: N/A, Valid
    In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 60000003, Policy name: self-traffic-policy/1, Timeout: 56, Valid
    In: 23.0.21.6/500 --> 23.0.21.11/500;udp, If: .local..0, Pkts: 105425, Bytes: 12031144
    Out: 23.0.21.11/500 --> 23.0.21.6/500;udp, If: ge-0/0/2.0, Pkts: 106503, Bytes: 12110680

  Session ID: 60017387, Policy name: default-policy-00/2, Timeout: 1796, Valid
    In: 22.0.20.155/53053 --> 26.0.20.156/23;tcp, If: ge-0/0/1.0, Pkts: 10, Bytes: 610
    Out: 26.0.20.156/23 --> 22.0.20.155/53053;tcp, If: st0.0, Pkts: 9, Bytes: 602
  Total sessions: 4

  Flow Sessions on FPC6 PIC0:

  Session ID: 120000001, Policy name: N/A, Timeout: N/A, Valid
    In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0

  Session ID: 120000002, Policy name: N/A, Timeout: N/A, Valid
    In: 23.0.21.11/0 --> 23.0.21.6/0;esp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0
  Total sessions: 2
```

After VPN session affinity is enabled, the clear-text session is always located on FPC 3, PIC 0.

Related Documentation

- Understanding VPN Session Affinity on page 6789
- Understanding Session Cache on page 1713
- Express Path Overview on page 1717

Accelerating the IPsec VPN Traffic Performance

You can accelerate the IPsec VPN performance by configuring the performance acceleration parameter. By default, VPN performance acceleration is disabled on SRX

Copyright © 2016, Juniper Networks, Inc.

Series devices. Enabling the VPN performance acceleration can improve the VPN throughput with VPN session affinity enabled.

This topic describes how to use the CLI to enable VPN performance acceleration.

NOTE: To enable performance acceleration, you must ensure that cleartext sessions and IPsec tunnel sessions are established on the same Services Processing Unit (SPU). For more information on enabling the session affinity, see "Understanding VPN Session Affinity" on page 6789.

To enable IPsec VPN performance acceleration:

1. Enable VPN session affinity.

   ```
   [edit]
   user@host# set security flow load-distribution session-affinity ipsec
   ```

2. Enable IPsec performance acceleration.

   ```
   [edit]
   user@host# set security flow ipsec-performance-acceleration
   ```

3. Check your changes to the configuration before committing.

   ```
   [edit]
   user@host# commit check
   ```

4. Commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

After enabling VPN performance acceleration, use the **show security flow status** command to display flow status.

```
Flow forwarding mode:
   Inet forwarding mode: flow based
   Inet6 forwarding mode: drop
   MPLS forwarding mode: drop
   ISO forwarding mode: drop
  Flow trace status
    Flow tracing status: off
  Flow session distribution
    Distribution mode: Hash-based
     Flow packet ordering
    Ordering mode: Hardware
  Flow ipsec performance acceleration: on
```

Related Documentation

- Understanding VPN Session Affinity on page 6789
- Enabling VPN Session Affinity on page 6791
- ipsec-performance-acceleration (Security Flow) on page 6897
- show security flow status on page 2164

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 89

Troubleshooting

- Tunnel Events on page 6797

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 280

Tunnel Events

- Tunnel Events on page 6797

Tunnel Events

Tunnel events can include successful IPsec SA negotiations, IPsec and IKE SA rekeys, SA negotiation failures, and reasons for a tunnel going down. Tunnel events appear in the output for the **show security ipsec inactive-tunnel**, **show security ipsec inactive-tunnel detail**, and **show security ipsec security-association detail** commands. Table 571 lists the tunnel events in alphabetical order. Each event includes a description and the action you can take.

Table 571: IPsec VPN Tunnel Events

Tunnel Event	Description	Action
Bind-interface's address deleted. Existing IPSec SAs cleared	A configuration commit removed the IP address from the st0 interface, which resulted in the clearing of the IPsec SA for VPNs bound to the interface.	Review the VPN setup to determine the need for the IP address on the st0 tunnel interface. Review system logs for the commit change.
Bind-interface's address received. Information updated	A configuration commit changed or added an IP address to the st0 tunnel interface.	No action required.
Bind-interface's family deleted. Existing IPSec SAs cleared	A configuration commit removed the family inet or inet6 from the st0 interface, which resulted in the clearing of the IPsec SA for VPNs bound to the interface.	Verify in the configuration that st0.x has the family inet or inet6 associated with the interface. Review system logs for the commit changes.
Bind-interface's family received. Information updated	A configuration commit added the family inet or inet6 on the st0 interface.	No action required.
Bind-interface's zone received. Information updated	A configuration commit changed or added a security zone on the st0 tunnel interface.	No action required.
Bind-interface's zone status changed. Existing IPSec SAs cleared	The st0.x interface status changed from Up, which cleared the IPsec SA for VPNs bound to the st0 interface where the status changed.	Review system logs for the the interface status change reason.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
CA certificate for configured local certificate not found. Negotiation not initiated/successful	During VPN establishment using PKI certificates, the CA for the local certificate was not found on the device, which resulted in VPN establishment failure.	Verify the **ca-profile** configuration. Verify that the CA certificate is loaded on the device. Reload the CA certificate if necessary.
Certificate has expired. Refer to syslog for more information	An attempt to establish a VPN using PKI certificates failed because the CA or local certificate was expired.	Verify certificate validity dates. Verify the system date and time.
Cleared unused redundant backup IPSec SA pairs.	The IPsec SA count for a tunnel crossed two pairs.	No action required.
Configured local certificate has been revoked. Negotiation not initiated/successful	During a local certificate revocation check using the CRL, the local certificate was revoked or the CRL could not be downloaded to allow the revocation check, which resulted in VPN establishment failure or a failure to initiate the VPN tunnel.	Review system logs or PKI trace options for information about the CRL validation failure. Verify the downloaded CRL. Manually load an updated CRL. Consult the CA administrator about why the certificate is on the CRL. Disable the CRL revocation check.
CRL check failed as CA not reachable. Refer to syslog for more information	During a certificate revocation check using the CRL, the CA server could not be reached or did not respond, which resulted in VPN establishment failure.	Verify that the CA server and the CRL distribution point are reachable.
CRL check failed for a certificate. Refer to syslog for more information	During a certificate revocation check using the CRL, the received peer certificate was revoked or the CRL could not be downloaded to allow the revocation check, which resulted in VPN establishment failure.	Review system logs or PKI trace options for information about the CRL validation failure. Verify the downloaded CRL. Manually load an updated CRL. Consult the CA administrator about why the certificate is on the CRL. Disable the CRL revocation check.
Deactivated tunnel as interface information is not ready on new primary node	During a failover in a branch SRX Series chassis cluster, interface information was not available on the new primary node. This event is specific to branch SRX Series chassis clusters.	No action required.
DPD detected peer as down. Existing IKE/IPSec SAs cleared.	DPD is enabled and the peer was not reachable for the configured interval and threshold. When this happens, the corresponding IKE and IPsec SAs are cleared, causing the tunnel to flap.	Check peer connectivity. Verify peer gateway connectivity, increase DPD intervals or thresholds, or enable **probe-idle-tunnel**.
Duplicate IKE/IPSec session detected. Old session cleared	An established peer connected again with different information, such as IP address, username, or IKE ID. This event occurs for AutoVPN, dynamic endpoint, and dialup tunnels only.	No action required.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
External-interface's address deleted. Existing IPSec SAs cleared	A configuration commit removed or adjusted the IP address on the IKE external interface, which resulted in the clearing of the IPsec SA for IKE gateways bound to the interface.	Verify that an IP address is assigned to the IKE external interface. Review system logs for the commit change.
External interface's address received. Information updated	A configuration commit changed or added a security zone on the IKE external interface.	No action required.
External-interface's device status changed. Existing IPSec SAs cleared	The IKE gateway external interface status changed from Up, which resulted in clearing of the IPsec SA for all IKE gateways associated with the external interface where the status changed.	Review system logs for the interface status change reason.
External-interface's primary address change triggered clearing of IPSec SA	A configuration commit adjusted the IP address on the IKE external interface, which resulted in clearing of the IPsec SA for IKE gateways bound to the adjusted interface.	Verify the IP address assigned to the external interface. Verify use of the primary setting on the external interface. Review system logs for the commit change.
External-interface's sub-unit status changed. Existing IPSec SAs cleared	The IKE gateway external interface status changed from Up, which resulted in clearing of the IPsec SA for all IKE gateways associated with the external interface where the status changed.	Review system logs for the interface status change reason.
External interface's zone received. Information updated	A configuration commit changed or added a security zone on the IKE external interface.	No action required.
External interface's zone status changed. Existing IPSec SAs cleared	A configuration commit changed the security zone for the IKE external interface, which resulted in the clearing of the IPsec SA for all IKE gateways associated with the changed external interface.	Review system logs for commit changes.
Gateway configuration deletion triggered clearing of IPSec SA	A configuration commit deleted or deactivated the IKE gateway, which resulted in clearing of the IPsec SA.	Review system logs for commit changes.
Group VPN configuration change triggered clearing of IPSec SA	A configuration commit changed the group VPN configuration, which resulted in clearing of the IPsec SA.	Review system logs for commit changes.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
Hard lifetime of IPSec SA expired.	This event is tracked for a tunnel only if there are no more IPsec SAs. Otherwise, this event is tracked in statistics only to avoid multiple events being recorded during rekeys.	If the rekey fails or does not complete before the lifetime expires, this event is recorded and the statistics counter is incremented. If the hard lifetime expires before a rekey occurs, a higher lifetime value is recommended. If a rekey was triggered and failed, there might be some other issue noted in another tunnel event.
Idle timer triggered. Existing IPSec SAs cleared.	**idle-time** is configured at the [**edit security ipsec vpn** *vpn-name* **ike**] hierarchy level, and the tunnel was idle for the configured time.	Increase the idle tunnel interval.
IKE SA cleared as lifetime exipred	The IKE configured lifetime seconds expired. The default setting is 28,800 seconds. This event does not impact current IPsec SAs.	No action required. You can use DPD to maintain IKE establishment.
IKE SA cleared on backup HA node as requested from primary HA node	The primary chassis cluster node requested that the IKE SA be cleared on the backup node.	Review system logs on the primary node for the IKE SA clear reason.
IKE SA negotiation successfully completed	IKE Phase 1 negotiations were successfully completed.	No action required.
IKE SA rekey successfully completed	When using IKEv2, the IKE SA expired with an established IPsec SA. IKEv2 requires an established IKE SA while an IPsec SA is active.	No action required.
IKE SA UDP port change detected with peer. Existing IPSec SAs cleared.	There was a NAT-T port change, possibly caused by changed ports on the NAT device after the tunnel was established. An IPsec layer UDP packet was received from the peer with a different port for the established tunnel. This event resulted in the clearing of the IPsec SA.	Verify the NAT device behavior that led to the port change.
IKE version mismatch detected	The SRX Series device and the VPN peer attempted to use different IKE versions, which resulted in tunnel establishment failure. The SRX Series device is configured for IKEv1 usage by default.	Adjust the VPN peer to use the same IKE version as the SRX Series; or configure the SRX Series to use the same IKE version in as the peer with **set security ike gateway** *gateway-name* **version v1-only** or **set security ike gateway** *gateway-name* **version v2-only**.
Initial-Contact received from peer. Stale IKE/IPSec SAs cleared.	Initial contact was received from the peer.	No action required.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
IPSec SA delete payload received from peer, corresponding IPSec SAs cleared.	A peer or remote device sent a delete notification for a given IPsec SA, resulting in the deletion of that particular SA pair. If that SA is the last IPsec SA for that tunnel, the tunnel goes down. This event can occur for various reasons: for example, after a rekey the peer might send a delete for an old SA, or a configuration change triggered on a peer resulted in the clearing of the IPsec SA.	Review peer logs to locate the event that caused the SA deletion request to be sent.
IPSec SA negotiation successfully completed	IPsec Phase 2 negotiations were successfully completed.	No action required.
IPSec SA rekey successfully completed	The IPsec rekey was successfully completed.	No action required.
IPSec SA UDP port change detected with peer. Existing IPSec SAs cleared.	There was a NAT-T port change, possibly caused by changed ports on the NAT device after the tunnel was established. An IPsec layer UDP packet was received from the peer with a different port for the established tunnel. This event resulted in the clearing of the IPsec SA.	Verify the NAT device behavior that led to the port change.
IPSec SAs cleared as corresponding IKE SA deleted.	The IPsec SA was deleted.	No action required.
Key pair not found for configured local certificate. Negotiation failed	During VPN establishment using PKI certificates, a corrupt or missing key-pair file from the local device was detected, which resulted in VPN establishment failure.	Verify configuration of **local-certificate** in the IKE policy. Verify that the key-pair is located in **/var/db/certs/common/key-pair**. Generate a new key-pair and certificate-request, and load the new certificate.
Lifetime in kilobytes expired for IPSec SA.	The **lifetime-kilobytes** value has expired. Before this event, the soft lifetime triggered a rekey for the IPsec SA. The event is not captured and only a statistics counter is incremented.	If the rekey fails or does not complete before the lifetime expires, this event is recorded and the statistics counter is incremented. If the hard lifetime expires before a rekey occurs, a higher lifetime value is recommended. If a rekey was triggered and failed, there might be some other issue noted in another tunnel event.
Manual next-hop-tunnel configuration change triggered clearing of IPSec SA	A configuration commit changed the next-hop tunnel for the st0 interface, which resulted in the clearing of the IPsec SA for the VPN linked to the changed next-hop tunnel.	Review system logs for commit changes.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
Negotiation failed with error code INVALID_IKE_VERSION received from peer	The peer device rejected an incoming VPN tunnel setup request from the SRX Series device because of mismatched IKE versions, resulting in tunnel establishment failure.	Verify the VPN configuration and VPN peer configuration for IKE version usage. Configure the SRX Series device to use IKEv1 or IKEv2 based on the peer setup by entering **set security ike gateway *gateway-name* version v1-only** or **set security ike gateway *gateway-name* version v2-only**.
Negotiation failed with error code NO_PROPOSAL_CHOSEN received from peer	The VPN peer informed the SRX Series device of VPN failure based on a mismatch of proposals, IKE version, peer gateway match, proxy ID/traffic-selectors, DH groups, or PSK usage.	Review peer logs for the failure reason. Review configurations on the SRX Series device and the peer to ensure that expected VPN attributes match.
Negotiation failed with error code TS_UNACCEPTABLE received from peer	The VPN peer rejected the proxy ID/traffic selector requested by the SRX Series device, which resulted in tunnel establishment failure.	Review peer logs for the rejection reason. For route-based VPNs, verify the configured proxy ID/traffic selector. For policy-based VPNs, verify the source, policy, or application defined in the security policy bound to the VPN.
OCSP revocation check failed as server not reachable. Refer to syslog for more information	During a certificate revocation check using OCSP, the OCSP server could not be contacted, which resulted in VPN establishment failure.	Verify that the OCSP server is reachable. Verify the configured IP address of the OCSP server.
OCSP revocation check failed for a certificate. Refer to syslog for more information	During a certificate revocation check using OCSP, a revoke response was received, which resulted in VPN establishment failure.	Review OCSP server logs for the revocation reason.
Peer proposed phase1 negotiation mode (main/aggressive) does not match with configuration	The IKE negotiation mode configured on the SRX Series device for IKEv1 does not match the peer's proposed mode.	Revise the peer or SRX Series device configuration to match the other device.
Peer proposed phase1 proposal conflicts with local configuration. Negotiation failed	The Phase 1 proposal configured on the SRX Series device does not match the peer's proposal.	Revise the peer or SRX Series device configuration to match the other device.
Peer proposed traffic-selectors are not in configured range	The traffic selector configured on the SRX Series device does not match the peer's proposed traffic selectors.	Revise the peer or SRX Series device configuration to match the other device.
Peer proposed unsupported multiple traffic-selector attributes for a single IPSec SA. Negotiation failed.	During IKEv2 negotiations, the peer device sent a proposal containing multiple traffic selectors for a single VPN tunnel, which resulted in the failure of the VPN tunnel setup.	Review the peer configuration of ACLs or traffic selectors.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
Peer proposed unsupported port range in traffic-selector attribute. Phase 2 negotiation failed	During IPsec negotiation, the peer device sent a traffic selector that contained an unsupported port range, which resulted in the failure of the VPN tunnel setup.	Adjust the peer configuration for the port range setup for the ACLs or traffic selectors.
Peer proposed unsupported protocol in traffic-selector attribute. Phase 2 negotiation failed	During IPsec negotiation, the peer device sent a traffic selector that contained an unsupported protocol, which resulted in the failure of the VPN tunnel setup.	Adjust the peer configuration for the protocol setup for the ACLs or traffic selectors.
Peer's IKE-ID validation failed during negotiation	The received IKE ID did not match the expected IKE ID, which resulted in tunnel establishment failure. The default expected IKE ID is the IP address, peer, or **dynamic** setting configured for the IKE gateway.	Review the VPN peer configuration for the IKE ID the peer is sending. Configure the SRX Series device using **remote-identity** to adjust to the expected IKE ID of the peer.
Proposed peer's IKE-ID does not match with peer's certificate. Negotiation failed	When using PKI certificates, the peer IKE ID value was not in the SAN field of the received certificate, which resulted in VPN establishment failure.	Review the VPN peer and reissue a certificate with an updated SAN based on the IKE ID value. Adjust the VPN peer's IKE ID to match the SAN field of the certificate.
Received use IKEv1 message from peer	The peer device rejected an incoming VPN tunnel setup request from the SRX Series device to use IKEv2 when the peer is configured to use IKEv1, which resulted in tunnel establishment failure.	Adjust the VPN peer setup to use IKEv2, or adjust the SRX Series device's configuration to use IKEv1 by entering **set security ike gateway *gateway-name* version v1-only**.
Requested peer to use IKEv1 instead of IKEv2	The SRX Series device is configured to use IKEv1 by default, and the peer attempted to set up IKE using IKEv2, which resulted in tunnel establishment failure.	Adjust the VPN peer to use IKEv1, or configure the SRX Series device to use IKEv2 by entering **set security ike gateway *gateway-name* version v2-only**.
Security policy change triggered clearing of IPSec SA	A policy-based VPN configuration commit changed security policies bound to the IPsec VPN, which resulted in the clearing of the IPsec SA associated with the changed policies.	Review system logs for commit changes.
Shortcut Tunnel deleted because of inactivity	When using IKEv2 with ADVPN, the device received a shortcut suggestion. However, it did not receive a request from a partner to complete the setup of the shortcut tunnel.	Verify that shortcut tunnel peers can reach each other. Verify that the shortcut partners can exchange UDP500 IKEv2 traffic between them.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
Shortcut Tunnel deleted when idle-time is reached	When using IKEv2 with ADVPN, traffic flowing over the shortcut tunnel fell below the **idle-threshold** for longer than the **idle-time** (default is 5 packets per second for 900 seconds). Traffic continues to flow through the IPsec tunnel to the hub.	If traffic is sporadic, decrease **idle-threshold** and increase **idle-time**. The shortcut tunnel should remain established during times of low traffic throughput.
Tunnel configuration changed. Corresponding IKE/IPSec SAs are deleted	A configuration commit adjusted the IKE/IPsec configuration, which resulted in clearing of the IPsec SA.	Review system logs for commit changes.
Tunnel configuration is deleted. Corresponding IKE/IPSec SAs are deleted	A configuration commit deleted or deactivated the IKE/IPsec configuration, which resulted in clearing of the IPsec SA.	Review system logs for commit changes.
Tunnel deleted on backup HA node as requested from primary HA node	This event is generated on the backup chassis cluster node when the tunnel on the primary node is deleted.	No action required.
Tunnel ID reused for other tunnel on primary node. Cleared stale tunnel	On high-end SRX Series chassis clusters, if the tunnel ID becomes out of sync for a given tunnel, the old tunnel is removed on the backup chassis cluster node.	No action required.
Tunnel is ready. Waiting for trigger event or peer to trigger negotiation.	The required configuration is available for peer negotiation. The device is awaiting traffic for tunnel establishment or a tunnel setup request from the peer.	No action required.
Unsupported AH and ESP bundle negotiation request denied	The peer proposed AH and ESP protocols on the same IPsec tunnel, but the SRX Series device does not support this configuration.	Reconfigure the peer for either AH or ESP protocol on the tunnel.
User cleared IKE SA from CLI, corresponding IPSec SAs cleared	The IKE SA was manually cleared using the CLI, which cleared the IKE SA but which does not affect current established IPsec SAs.	No action required.
User cleared IPSec SA from CLI.	A user or an administrator has cleared the IPsec SA manually in the CLI.	No action required.
VPN monitoring detected tunnel as down. Existing IPSec SAs cleared.	VPN monitor is configured for the tunnel and the peer did not respond to VPN monitor keepalive messages, or the peer was not reachable. The corresponding IPsec SAs were cleared.	Check peer connectivity and the VPN monitor destination address.

Copyright © 2016, Juniper Networks, Inc.

Table 571: IPsec VPN Tunnel Events *(continued)*

Tunnel Event	Description	Action
Zone change for all interface detected. Existing IPSec SAs cleared	A configuration commit changed the security zone for all interfaces, which resulted in clearing of all device IPsec SAs.	Review system logs for commit changes.

Related Documentation

- show security ipsec inactive-tunnels on page 7058
- show security ipsec security-associations on page 3681

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

PART 90

Configuration Statements and Operational Commands

- Configuration Statements on page 6809
- Operational Commands on page 6973

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 281

Configuration Statements

- Security Configuration Statement Hierarchy on page 6814
- Access Configuration Statement Hierarchy on page 6815
- [edit security alarms] Hierarchy Level on page 6823
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security ike] Hierarchy Level on page 6827
- [edit security ipsec] Hierarchy Level on page 6829
- [edit security pki] Hierarchy Level on page 6831
- address (Security Group VPN Server IKE Gateway) on page 6832
- address (Security IKE Gateway) on page 6833
- administrator on page 6833
- advpn on page 6834
- algorithm (Security) on page 6835
- always-send on page 6836
- authentication (IPsec SA for OSPF) on page 6837
- authentication (Security IPsec) on page 6838
- authentication-algorithm (Security Group VPN IKE) on page 6839
- authentication-algorithm (Security Group VPN IPsec) on page 6839
- authentication-algorithm (Security IKE) on page 6840
- authentication-algorithm (Security IPsec) on page 6841
- authentication-method on page 6842
- authentication-method (Security Group VPN) on page 6843
- auto-re-enrollment (Security) on page 6844
- auxiliary-spi (IPsec SA for OSPF) on page 6845
- bind-interface on page 6845
- ca-identity (Security) on page 6846
- ca-profile (Security PKI) on page 6847
- ca-profile-name on page 6848
- certificate on page 6848

Copyright © 2016, Juniper Networks, Inc.

- certificate-id (Security) on page 6849
- challenge-password (Security) on page 6850
- connections-limit on page 6850
- container on page 6851
- crl (Security) on page 6852
- cryptographic-self-test on page 6853
- dead-peer-detection on page 6853
- dead-peer-detection (Security Group VPN Server) on page 6854
- decryption-failures on page 6854
- description (Security Policies) on page 6855
- destination-ip (Security IPsec) on page 6855
- df-bit on page 6856
- dh-group (Security IKE) on page 6857
- dh-group (Security Group VPN IKE) on page 6858
- disable (PKI) on page 6858
- distinguished-name (Security) on page 6859
- dynamic (Security) on page 6859
- dynamic (Security Group VPN) on page 6860
- encryption (IPsec SA for OSPF) on page 6861
- encryption (Security) on page 6862
- encryption-algorithm (Security Group VPN IKE) on page 6863
- encryption-algorithm (Security Group VPN IPsec) on page 6864
- encryption-algorithm (Security IKE) on page 6865
- encryption-algorithm (Security IPsec) on page 6866
- encryption-failures on page 6867
- enrollment (Security) on page 6868
- establish-tunnels on page 6869
- external-interface (Security IKE Gateway) on page 6869
- external-interface (Security Manual SA) on page 6870
- gateway (Security Group VPN Member IKE) on page 6870
- gateway (Security Group VPN Server IKE) on page 6871
- gateway (Security IKE) on page 6872
- gateway (Security IPsec VPN) on page 6873
- gateway (Security Manual SA) on page 6873
- general-ikeid on page 6874
- group (Security Group VPN) on page 6875
- group-vpn on page 6877

Copyright © 2016, Juniper Networks, Inc.

- hostname on page 6880
- idle-time on page 6880
- ike (Security) on page 6881
- ike (Security Group VPN Member) on page 6883
- ike (Security Group VPN Server) on page 6885
- ike (Security IPsec VPN) on page 6886
- ike-phase1-failures on page 6887
- ike-phase2-failures on page 6888
- ike-policy (Security Gateway) on page 6888
- ike-user-type on page 6889
- inet (Security Dynamic Peer) on page 6889
- inet6 (Security IKE Gateway) on page 6890
- install-interval on page 6890
- interval (Security IKE) on page 6891
- ipsec (Security) on page 6892
- ipsec (Security Group VPN Member) on page 6895
- ipsec (Security Group VPN Server) on page 6897
- ipsec-performance-acceleration (Security Flow) on page 6897
- ipsec-policy (Security) on page 6898
- ipsec-policy (Security Group VPN) on page 6898
- ipsec-sa (Security Group VPN) on page 6899
- ipsec-vpn (Security Flow) on page 6900
- key-generation-self-test on page 6900
- lifetime-kilobytes on page 6901
- lifetime-seconds (Security Group VPN) on page 6901
- lifetime-seconds (Security IKE) on page 6902
- lifetime-seconds (Security IPsec) on page 6902
- load-distribution on page 6903
- local (Security IPsec) on page 6903
- local-address on page 6904
- local-address (Security Group VPN Member) on page 6905
- local-address (Security Group VPN Server) on page 6905
- local-certificate (Security) on page 6906
- local-identity on page 6907
- local-identity (Security Group VPN) on page 6908
- manual (Security IPsec) on page 6909
- member (Security Group VPN) on page 6910

Copyright © 2016, Juniper Networks, Inc.

- member-threshold (Security Group VPN) on page 6912
- mode (Security Group VPN) on page 6913
- mode (Security IKE Policy) on page 6914
- nat-keepalive on page 6915
- no-anti-replay (Security) on page 6915
- no-nat-traversal on page 6916
- non-cryptographic-self-test on page 6916
- ocsp (Security PKI) on page 6917
- optimized on page 6918
- optimized (DPD) on page 6918
- peer-certificate-type on page 6919
- perfect-forward-secrecy (Security IPsec) on page 6920
- pki on page 6921
- pki-local-certificate on page 6922
- policy (Security Group VPN IKE) on page 6923
- policy (Security IKE) on page 6924
- policy (Security IPsec) on page 6925
- policy-oids on page 6926
- pre-shared-key (Security IKE Policy) on page 6927
- probe-idle-tunnel on page 6927
- profile (Access) on page 6928
- proposal (Security Group VPN Member IKE) on page 6931
- proposal (Security Group VPN Server IKE) on page 6932
- proposal (Security Group VPN Server IPsec) on page 6932
- proposal (Security IKE) on page 6933
- proposal (Security IPsec) on page 6934
- proposals (Security Group VPN) on page 6934
- proposals (Security IKE) on page 6935
- proposals (Security IPsec) on page 6935
- proposal-set (Security IKE) on page 6936
- proposal-set (Security IPsec) on page 6939
- protocol (IPsec SA for OSPF) on page 6941
- protocol (Security IPsec) on page 6941
- protocol (Security IPsec Manual SA) on page 6942
- proxy-identity on page 6942
- re-enroll-trigger-time-percentage (Security PKI) on page 6943
- re-generate-keypair on page 6943

Copyright © 2016, Juniper Networks, Inc.

- refresh-interval on page 6944
- remote (Security IPsec) on page 6944
- remote-identity on page 6945
- remote-identity (Security Group VPN) on page 6946
- replay-attacks on page 6947
- respond-bad-spi on page 6947
- revocation-check (Security PKI) on page 6948
- routing-instance (Security Group VPN) on page 6949
- routing-instance (Security PKI) on page 6949
- security-association on page 6950
- server (Security Group VPN) on page 6951
- server-address (Security Group VPN Member) on page 6953
- server-cluster (Security Group VPN Server) on page 6954
- server-member-communication (Security Group VPN Server) on page 6955
- service (Security IPsec) on page 6956
- session-affinity on page 6956
- source-interface (Security) on page 6957
- spi (IPsec SA for OSPF) on page 6957
- spi (Security IPsec) on page 6958
- threshold (Security IKE Gateway) on page 6958
- traceoptions (Security Group VPN) on page 6959
- traceoptions (Security IKE) on page 6962
- traceoptions (Security IPsec) on page 6964
- traceoptions (Security PKI) on page 6965
- traffic-selector on page 6966
- trusted-ca (Security IKE Policy) on page 6967
- use-ocsp (Security PKI) on page 6967
- user-at-hostname on page 6968
- version (Security IKE Gateway) on page 6968
- vpn (Security) on page 6969
- vpn-monitor on page 6970
- vpn-monitor-options on page 6971
- wildcard on page 6971
- xauth on page 6972

Copyright © 2016, Juniper Networks, Inc.

Security Configuration Statement Hierarchy

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- [edit security address-book] Hierarchy Level on page 643
- [edit security alarms] Hierarchy Level on page 6823
- [edit security alg] Hierarchy Level on page 312
- *[edit security analysis] Hierarchy Level*
- [edit security application-firewall] Hierarchy Level on page 644
- [edit security application-tracking] Hierarchy Level on page 644
- *[edit security certificates] Hierarchy Level*
- [edit security datapath-debug] Hierarchy Level on page 3526
- *[edit security dynamic-vpn] Hierarchy Level*
- [edit security firewall-authentication] Hierarchy Level on page 3527
- [edit security flow] Hierarchy Level on page 1837
- [edit security forwarding-options] Hierarchy Level on page 3006
- [edit security forwarding-process] Hierarchy Level on page 1839
- [edit security gprs] Hierarchy Level on page 2341
- [edit security group-vpn] Hierarchy Level on page 6824
- [edit security idp] Hierarchy Level on page 3528
- [edit security ike] Hierarchy Level on page 3537
- [edit security ipsec] Hierarchy Level on page 3539
- [edit security log] Hierarchy Level on page 645
- [edit security nat] Hierarchy Level on page 316
- [edit security pki] Hierarchy Level on page 646
- [edit security policies] Hierarchy Level on page 320
- *[edit security resource-manager] Hierarchy Level*

Copyright © 2016, Juniper Networks, Inc.

- [edit security screen] Hierarchy Level on page 979
- [edit security softwires] Hierarchy Level on page 3552
- *[edit security ssh-known-hosts] Hierarchy Level*
- [edit security traceoptions] Hierarchy Level on page 325
- [edit security user-identification] Hierarchy Level on page 1215
- [edit security utm] Hierarchy Level on page 5830
- [edit security zones] Hierarchy Level on page 324

Related Documentation

- *CLI User Guide*
- CLI Explorer

Access Configuration Statement Hierarchy

Use the statements in the **access** configuration hierarchy to configure access to the device and authentication methods, including address assignment and address pool, user and firewall authentication, a group profile, LDAP options and LDAP server configuration, an access profile, RADIUS options and RADIUS server configuration, and SecurID server configuration.

```
access {
  address-assignment {
    abated-utilization percentage;
    abated-utilization-v6 percentage;
    high-utilization percentage;
    high-utilization-v6 percentage;
    neighbor-discovery-router-advertisement ndra-name;
    pool pool-name {
      family {
        inet {
          dhcp-attributes {
            boot-file boot-file-name;
            boot-server boot-server-name;
            domain-name domain-name;
            grace-period seconds;
            maximum-lease-time (seconds | infinite);
            name-server ipv4-address;
            netbios-node-type (b-node | h-node | m-node | p-node);
            next-server next-server-name;
            option dhcp-option-identifier-code {
              array {
                byte [8-bit-value];
                flag [ false| off |on |true];
                integer [32-bit-numeric-values];
                ip-address [ip-address];
                short [signed-16-bit-numeric-value];
                string [character string value];
                unsigned-integer [unsigned-32-bit-numeric-value];
                unsigned-short [16-bit-numeric-value];
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        byte 8-bit-value;
        flag (false | off | on | true);
        integer 32-bit-numeric-values;
        ip-address ip-address;
        short signed-16-bit-numeric-value;
        string character string value;
        unsigned-integer unsigned-32-bit-numeric-value;
        unsigned-short 16-bit-numeric-value;
      }
      option-match {
        option-82 {
          circuit-id match-value {
            range range-name;
          }
          remote-id match-value;
            range range-name;
          }
        }
      }
      propagate-ppp-settings [interface-name];
      propagate-settings interface-name;
      router ipv4-address;
      server-identifier ip-address;
      sip-server {
        ip-address ipv4-address;
        name sip-server-name;
      }
      tftp-server server-name;
      wins-server ipv4-address;
    }
    host hostname {
      hardware-address mac-address;
      ip-address reserved-address;
    }
    network network address;
    range range-name {
      high upper-limit;
      low lower-limit;
    }
    xauth-attributes {
      primary-dns ip-address;
      primary-wins ip-address;
      secondary-dns ip-address;
      secondary-wins ip-address;
    }
  }
  inet6 {
    dhcp-attributes {
      dns-server ipv6-address;
      grace-period seconds;
      maximum-lease-time (seconds | infinite);
      option dhcp-option-identifier-code {
        array {
          byte [8-bit-value];
          flag [ false| off |on |true];
```


Copyright © 2016, Juniper Networks, Inc.

```
                integer [32-bit-numeric-values];
                ip-address [ip-address];
                short [signed-16-bit-numeric-value];
                string [character string value];
                unsigned-integer [unsigned-32-bit-numeric-value];
                unsigned-short [16-bit-numeric-value];
              }
              byte 8-bit-value;
              flag (false | off | on | true);
              integer 32-bit-numeric-values;
              ip-address ip-address;
              short signed-16-bit-numeric-value;
              string character string value;
              unsigned-integer unsigned-32-bit-numeric-value;
              unsigned-short 16-bit-numeric-value;
            }
            propagate-ppp-settings [interface-name];
            sip-server-address ipv6-address;
            sip-server-domain-name domain-name;
          }
          prefix ipv6-network-prefix;
          range range-name {
            high upper-limit;
            low lower-limit;
            prefix-length delegated-prefix-length;
          }
        }
      link pool-name;
    }
  }
  address-pool pool-name {
    (address address-or-address-prefix ) {
    address-range {
      high upper-limit;
      low lower-limit;
      mask network-mask;
    }
    primary-dns name;
    primary-wins name;
    secondary-dns name;
    secondary-wins name;
  }
  address-protection;
  domain {
    delimiter delimiter;
    map domain-map-name {
      aaa-logical-system logical-system-name;
      aaa-routing-instance routing-instance-name;
      access-profile access-profile-name;
      address-pool address-pool-name;
      dynamic-profile dynamic-profile-name;
      padn destination-address; {
        mask destination-mask;
        metric metric-value
      }
      strip-domain;
```

```
      target-logical-system logical-system-name;
      target-routing-instance target-routing-instance;
    }
    parse-direction (left-to-right | right-to-left);
  }
  firewall-authentication {
    pass-through {
      default-profile profile-name;
      ftp {
        banner {
          fail string;
          login string;
          success string;
        }
      }
      http {
        banner {
          fail string;
          login string;
          success string;
        }
      telnet {
        banner {
          fail string;
          login string;
          success string;
        }
    }
    traceoptions {
      file {
        filename;
        files number;
        flag flag;
        match regular-expression;
        no-remote-trace;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
    }
    web-authentication {
      banner {
        success string;
      }
      default-profile profile-name;
    }
  }
  group-profile profile-name {
    ppp {
      cell-overhead;
      encapsulated-overhead encapsulated-overhead-value;
      framed-pool address-pool-name;
      idle-timeout seconds;
      interface-id interface-identifier;
      keepalive seconds;
      ppp-options {
        chap;
```


Copyright © 2016, Juniper Networks, Inc.

```
        pap;
      }
      primary-dns name;
      primary-wins name;
      secondary-dns name;
      secondary-wins name;
    }
  }
  gx-plus {
    global {
      max-outstanding-requests max-outstanding-requests;
    }
    partition partition-name {
      destination-host gx-plus-destination-host;
      destination-realm gx-plus-destination-realm;
      diameter-instance gx-plus-diameter-instance;
    }
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter filter-name;
    }
  }
  ldap-server hostname-or-address; {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  ppp-options {
    compliance {
      rfc(2486 | [rfc-number]);
    }
  }
  profile profile-name  {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      coa-immediate-update;
      duplication;
      immediate-update;
      order [accounting-method];
      statistics (time | volume-time);
      update-interval minutes;
    }
    accounting-order [accounting-method];
```

```
address-assignment pool pool-name;
authentication-order [ldap | none | password | radius | securid];
authorization-order [jsrc];
client client-name  {
   chap-secret  chap-secret;
   client-group [ group-names ];
   firewall-user {
      password password;
   }
   no-rfc2486;
   pap-password pap-password;
   x-auth ip-address;
}
client-name-filter {
   count number;
   domain-name domain-name;
   separator special-character;
}
ldap-options {
   assemble {
      common-name common-name;
   }
   base-distinguished-name base-distinguished-name;
   revert-interval seconds;
   search {
      admin-search {
         distinguished-name distinguished-name;
         password password;
      }
      search-filter search-filter-name;
   }
}
ldap-server server-address {
   port port-number;
   retry attempts;
   routing-instance routing-instance-name;
   source-address source-address;
   timeout seconds;
}
provisioning-order (gx-plus | jsrc);
radius {
   accounting-server [server];
   attributes {
      exclude {
         acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
         acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
         acc-loop-cir-id [access-request | accounting-start | accounting-stop];
         accounting-authentic [accounting-off | accounting-on | accounting-start |
            accounting-stop];
         accounting-delay-time [accounting-off | accounting-on | accounting-start |
            accounting-stop];
         accounting-session-id [access-request];
         accounting-terminate-cause [accounting-off];
         act-data-rate-dn [access-request | accounting-start | accounting-stop];
         act-data-rate-up [access-request | accounting-start | accounting-stop];
         act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
```


Copyright © 2016, Juniper Networks, Inc.

act-interlv-delay-up [access-request | accounting-start | accounting-stop];
att-data-rate-dn [access-request | accounting-start | accounting-stop];
att-data-rate-up [access-request | accounting-start | accounting-stop];
called-station-id [access-request | accounting-start | accounting-stop];
calling-station-id [access-request | accounting-start | accounting-stop];
class [access-request | accounting-start | accounting-stop];
delegated-ipv6-prefix [accounting-start | accounting-stop];
dhcp-gi-address [access-request | accounting-start | accounting-stop];
dhcp-mac-address [access-request | accounting-start | accounting-stop];
dhcp-options [access-request | accounting-start | accounting-stop];
downstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];
dsl-forum-attributes [access-request | accounting-start | accounting-stop];
dsl-line-state [access-request | accounting-start | accounting-stop];
dsl-type [access-request | accounting-start | accounting-stop];
dynamic-iflset-name [accounting-start | accounting-stop];
event-time-stamp [accounting-off | accounting-on | accounting-start | accounting-stop];
framed-interface-id [access-request | accounting-start | accounting-stop];
framed-ip-address [access-request | accounting-start | accounting-stop];
framed-ip-netmask [access-request | accounting-start | accounting-stop];
framed-ip-route [access-request | accounting-start | accounting-stop];
framed-ipv6-pool [accounting-start | accounting-stop];
framed-ipv6-prefix [accounting-start | accounting-stop];
framed-ipv6-route [accounting-start | accounting-stop];
framed-pool [accounting-start | accounting-stop];
input-filter [accounting-start | accounting-stop];
input-gigapackets [accounting-stop];
input-gigawords [accounting-stop];
input-ipv6-gigawords [accounting-stop];
input-ipv6-octets [accounting-stop];
input-ipv6-packets [accounting-stop];
interface-description [access-request | accounting-start | accounting-stop];
l2c-downstream-data [access-request | accounting-start | accounting-stop];
l2c-upstream-data [access-request | accounting-start | accounting-stop];
max-data-rate-dn [access-request | accounting-start | accounting-stop];
max-data-rate-up [access-request | accounting-start | accounting-stop];
max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
max-interlv-delay-up [access-request | accounting-start | accounting-stop];
min-data-rate-dn [access-request | accounting-start | accounting-stop];
min-data-rate-up [access-request | accounting-start | accounting-stop];
min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
nas-identifier [access-request | accounting-start | accounting-stop];
nas-port [access-request | accounting-off | accounting-on | accounting-start | accounting-stop];
nas-port-id [access-request | accounting-start | accounting-stop];
nas-port-type [access-request | accounting-start | accounting-stop];
output-filter [accounting-start | accounting-stop];
output-gigapackets [accounting-stop];
output-gigawords [accounting-stop];
output-ipv6-gigawords [accounting-stop];
output-ipv6-octets [accounting-stop];
output-ipv6-packets [accounting-stop];
upstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];

Copyright © 2016, Juniper Networks, Inc.

```
        }
        ignore {
          dynamic-iflset-name;
          framed-ip-netmask;
          input-filter;
          logical-system-routing-instance;
          output-filter;
        }
      }
    authentication-server [server];
    radius-options {
      request-rate number;
      revert-interval seconds;
    }
    radius-server server-address {
      accounting-port port-number
      max-outstanding-requests number-of--outstanding-requests;
      port port-number;
      retry attempts;
      routing-instance routing-instance-name;
      secret password;
      source-address source-address;
      timeout seconds;
    }
    service {
      accounting-order {
        activation-protocol;
        radius;
      }
    }
    session-options {
      client-group [group-name];
      client-idle-timeout minutes;
      client-session-timeout minutes;
    }
}
radius-options {
  request-rate number;
  revert-interval seconds;
}
radius-server server-address {
  accounting-port port-number;
  max-outstanding-requests number-of-max-outstanding-requests;
  port port-number;
  retry attempts;
  routing-instance routing-instance-name;
  secret password;
  source-address source-address;
  timeout seconds;
}
securid-server server-name {
  configuration-file filepath;
}
terminate-code {
  aaa {
    deny {
```


Copyright © 2016, Juniper Networks, Inc.

```
            authentication-denied {
              radius acct-terminate-cause-value;
            }
            no-resources {
              radius acct-terminate-cause-value;
            }
            server-request-timeout {
              radius acct-terminate-cause-value;
            }
          }
          shutdown {
            administrative-reset {
              radius acct-terminate-cause-value;
            }
            remote-reset {
              radius acct-terminate-cause-value;
            }
          }
        }
        dhcp {
          client-request {
            radius acct-terminate-cause-value;
          }
          lost-carrier {
            radius acct-terminate-cause-value;
          }
          nak {
            radius acct-terminate-cause-value;
          }
          nas-logout {
            radius acct-terminate-cause-value;
          }
          no-offers {
            radius acct-terminate-cause-value;
          }
        }
      }
    }
```

Related Documentation

- *Understanding User Authentication Methods*
- Layer 2 Transparent and Switching Overview on page 2893
- Understanding User Authentication for Security Devices on page 5173

[edit security alarms] Hierarchy Level

```
security {
  alarms {
    audible {
      continuous;
    }
    potential-violation {
      authentication failures;
      cryptographic-self-test;
```

```
        decryption-failures {
          threshold value;
        }
        encryption-failures {
          threshold value;
        }
        idp;
        ike-phase1-failures {
          threshold value;
        }
        ike-phase2-failures {
          threshold value;
        }
        key-generation-self-test;
        non-cryptographic-self-test;
        policy {
          application {
            duration interval;
            size count;
            threshold value;
          }
          destination-ip {
            duration interval;
            size count;
            threshold value;
          }
          policy match {
            duration interval;
            size count;
            threshold value;
          }
          source-ip {
            duration interval;
            size count;
            threshold value;
          }
        }
        replay-attacks {
          threshold value;
        }
        security-log-percent-full percentage;
      }
    }
  }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603

[edit security group-vpn] Hierarchy Level

```
security {
  group-vpn {
    member {
      ike {
        gateway gateway-name {
```


Copyright © 2016, Juniper Networks, Inc.

```
        ike-policy policy-name;
        local address ip-address;
        local-identity {
          (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
            e-mail-address);
        }
        remote-identity {
          (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
        }
        routing-instance routing-instance;
        server-address [ip-address];
      }
      policy policy-name {
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [proposal-name];
      }
      proposal proposal-name {
        authentication-algorithm (sha-256 | sha-384);
        authentication-method pre-shared-keys;
        description description;
        dh-group (group14 | group24);
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        gateway-filter {
          local-address ip-address;
          remote-address ip-address;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
    }
    ipsec {
      vpn vpn-name {
        df-bit (clear | copy | set);
        exclude rule rule-name {
          source-address ip-address/mask;
          destination-address ip-address/mask;
          application application;
        }
        fail-open rule rule-name {
          source-address ip-address/mask;
          destination-address ip-address/mask;
          application application;
        }
```

```
        group id;
        group-vpn-external-interface interface;
        ike-gateway gateway-name;
        recovery-probe;
      }
    }
  }
  server {
    group name {
      anti-replay-time-window milliseconds;
      description description;
      group-id number;
      ike-gateway gateway-name;
      ipsec-sa name {
        match-policy policy-name {
          destination ip-address/netmask;
          destination-port number;
          protocol number;
          source ip-address/netmask;
          source-port number;
        }
        proposal proposal-name;
      }
      member-threshold number;
      server-cluster {
        ike-gateway gateway-name;
        retransmission-period seconds;
        server-role (root-server | sub-server);
      }
      server-member-communication {
        certificate certificate-id;
        communication-type unicast;
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
        number-of-retransmission number;
        retransmission-period seconds;
        sig-hash-algorithm (sha-256 | sha-384);
      }
    }
    ike {
      gateway gateway-name {
        address ip-address ;
        dead-peer-detection {
          always-send;
          interval seconds;
          threshold number;
        }
        dynamic {
          (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        ike-policy policy-name;
        local-address ip-address;
        local-identity {
          (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
        }
        remote-identity {
```


Copyright © 2016, Juniper Networks, Inc.

```
                (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
              }
              routing-instance routing-instance;
            }
            policy policy-name {
              description text;
              mode (aggressive | main);
              pre-shared-key (ascii-text key | hexadecimal key);
              proposals [proposal-name];
            }
            proposal proposal-name {
              authentication-algorithm (sha-256 | sha-384);
              authentication-method pre-shared-keys;
              description description;
              dh-group (group14 | group24);
              encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
            }
          }
          ipsec {
            proposal proposal-name {
              authentication-algorithm hmac-sha-256-128;
              description description;
              encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
              lifetime-seconds seconds;
            }
          }
          traceoptions {
            file {
              filename;
              files number;
              match regular-expression;
              size maximum-file-size;
              (world-readable | no-world-readable);
            }
            flag flag;
            gateway-filter {
              local-address ip-address;
              remote-address ip-address;
            }
            level (all | error | info | notice | verbose | warning);
            no-remote-trace;
          }
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Group VPNv2 Overview on page 6649

[edit security ike] Hierarchy Level

```
security {
  ike {
    gateway gateway-name {
```

```
    address [ip-address-or-hostname];
    dead-peer-detection {
      (always-send | optimized | probe-idle-tunnel);
      interval seconds;
      threshold number;
    }
    dynamic {
      connections-limit number;
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
      ike-user-type (group-ike-id | shared-ike-id);
    }
    external-interface external-interface-name;
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
      (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address
        | user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    remote-identity {
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    version (v1-only | v2-only);
    xauth {
      access-profile profile-name;
    }
  }
  policy policy-name {
    certificate {
      local-certificate certificate-id;
      peer-certificate-type (pkcs7 | x509-signature);
    }
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1);
    authentication-method (dsa-signatures | ecdsa-signatures-256 |
      ecdsa-signatures-384 | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
  }
  respond-bad-spi <max-responses>;
  traceoptions {
    file {
```


Copyright © 2016, Juniper Networks, Inc.

```
            filename;
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
          rate-limit messages-per-second;
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041

[edit security ipsec] Hierarchy Level

```
security {
  ipsec {
    internal {
      security-association {
        manual encryption {
          iked_encryption enabled;
          algorithm 3des-cbc;
          key ascii-text key;
          }
      }
    }
    policy policy-name {
      description description;
      perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24
        | group5);
      proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
      proposals [proposal-name];
    }
    proposal proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
      description description;
      encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc |
        aes-192-gcm | aes-256-cbc | aes-256-gcm | des-cbc);
      lifetime-kilobytes kilobytes;
      lifetime-seconds seconds;
      protocol (ah | esp);
    }
    security-association sa-name {
      manual {
        direction bidirectional {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key {
              ascii-text key;
              hexadecimal key;
              }
```

Copyright © 2016, Juniper Networks, Inc.

```
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
          algorithm (3des-cbc | des-cbc | null);
          key {
            ascii-text key;
            hexadecimal key;
          }
        }
        protocol (ah | esp);
        spi spi-value;
      }
    }
    mode transport;
  }
  traceoptions {
    flag flag;
  }
  vpn vpn-name {
    bind-interface interface-name;
    copy-outer-dscp;
    df-bit (clear | copy | set);
    establish-tunnels (immediately | on-traffic);
    ike {
      gateway gateway-name;
      idle-time seconds;
      install-interval seconds;
      ipsec-policy ipsec-policy-name;
      no-anti-replay;
      proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
      }
    }
    manual {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      encryption {
        algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
        key (ascii-text key | hexadecimal key);
      }
      external-interface external-interface-name;
      gateway ip-address;
      protocol (ah | esp);
      spi spi-value;
    }
    traffic-selector traffic-selector-name {
      local-ip ip-address/netmask;
      remote-ip ip-address/netmask;
    }
    vpn-monitor {
      destination-ip ip-address;
      optimized;
```


Copyright © 2016, Juniper Networks, Inc.

```
            source-interface interface-name;
        }
    }
    vpn-monitor-options {
        interval seconds;
        threshold number;
    }
  }
}
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- IPsec VPN Overview on page 6041
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

[edit security pki] Hierarchy Level

```
security {
  pki {
    auto-re-enrollment {
      cmpv2 {
        certificate-id certificate-id-name {
          ca-profile-name ca-profile-name ;
          re-enroll-trigger-time-percentage percentage ;
          re-generate-keypair;
        }
      }
      scep {
        certificate-id certificate-id-name {
          ca-profile-name ca-profile-name ;
          challenge-password password ;
          re-enroll-trigger-time-percentage percentage ;
          re-generate-keypair;
        }
      }
    }
    ca-profile ca-profile-name {
      administrator {
        e-mail-address e-mail-address;
      }
      ca-identity ca-identity ;
      enrollment {
          retry number;
          retry-interval seconds;
          url url-name;
      }
      revocation-check {
        crl {
          disable {
            on-download-failure;
          }
          refresh-interval hours;
          url url-name;
        }
```

Copyright © 2016, Juniper Networks, Inc.

```
            disable;
            ocsp {
              connection-failure (disable | fallback-crl);
              disable-responder-revocation-check;
              nonce-payload (enable | disable);
              url ocsp-url;
            }
            use-ocsp;
          }
          routing-instance routing-instance-name ;
        }
        traceoptions {
          file filename {
            files number;
            match regular-expression;
            (no-world-readable | world-readable);
            size maximum-file-size;
          }
          flag flag;
          no-remote-trace;
        }
      }
    }
```

Related Documentation

- Security Configuration Statement Hierarchy on page 603
- Understanding Security Building Blocks for Security Devices on page 1045

address (Security Group VPN Server IKE Gateway)

Syntax address *ip-address-or-hostname* ;

Hierarchy Level [edit security group-vpn server ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.2 for **group-vpn** hierarchy.

Description Specify the IPv4 address of the primary Internet Key Exchange (IKE) gateway.

Options ***ip-address-or-hostname***—IPv4 address of an IKE gateway.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

address (Security IKE Gateway)

Syntax address [*ip-address-or-hostname*];

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description Specify the IPv4 or IPv6 address or the hostname of the primary Internet Key Exchange (IKE) gateway and up to four backup gateways.

Options ***ip-address-or-hostname***—IPv4 or IPv6 address or hostname of an IKE gateway.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

administrator

Syntax
```
administrator {
  e-mail-address e-mail-address ;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify an administrator e-mail address to which the certificate request is sent.

Options **e-mail-address *e-mail-address*** —E-mail address where the certificate request is sent. By default, there is no preset e-mail address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

advpn

Syntax

```
advpn {
  suggester {
    disable;
  }
  partner {
    connection-limit number;
    idle-threshold packets/sec;
    idle-time seconds;
    disable;
  }
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 12.3X48-D10. The range for the **idle-threshold** option and the range and default value for the **idle-time** option revised in Junos OS Release 12.3X48-D20.

Description Enable Auto Discovery VPN (ADVPN) protocol on the specified gateway.

Options

suggester—VPN peer that can initiate a shortcut exchange to allow shortcut partners to establish dynamic security associations (SAs) with each other. Specify **disable** to disable this role on the gateway.

NOTE: Both suggester and partner roles are enabled if **advpn** is configured without explicitly configuring **suggester** or **partner** keywords. We do not support suggester and partner roles on the same gateway. You must explicitly configure **disable** with the **suggester** or **partner** keyword to disable that particular role. You cannot disable both suggester and partner roles on the same gateway.

partner—VPN peer that can receive a shortcut exchange suggesting that it should establish dynamic SAs with another peer. Specify **disable** to disable this role on the gateway. The following options can be configured for the partner role:

connection-limit—Maximum number of shortcut tunnels that can be created with different shortcut partners using a particular gateway. The maximum number, which is also the default, is platform-dependent.

NOTE: Reducing the configured **connection-limit** value causes all active shortcut tunnels to be brought down. For example, if **connection-limit** is configured as 100 and you later reconfigure the number to 80, all active shortcut tunnels are brought down. Increasing the configured **connection-limit** value does not cause shortcut tunnels to go down.

Copyright © 2016, Juniper Networks, Inc.

idle-threshold—Rate, in packets per second, below which the shortcut is brought down.

Range: 3 through 5,000 packets per second.

Default: 5 packets per second.

idle-time—Duration, in seconds, after which the shortcut is deleted if the traffic remains below the **idle-threshold** value.

Range: 60 seconds through 86,400 seconds.

Default: 300 seconds.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Auto Discovery VPN on page 6559

algorithm (Security)

Syntax
algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);

Hierarchy Level
[edit security ipsec vpn *vpn-name* manual encryption]

Release Information
Statement modified in Junos OS Release 8.5.

Description
Select the encryption algorithm for the internal Routing-Engine-to-Routing-Engine IPsec security association (SA) configuration.

Options
- **3des-cbc**—3DES-CBC encryption algorithm.
- **aes-128-cbc**—AES-CBC 128-bit encryption algorithm.
- **aes-192-cbc**—AES-CBC 192-bit encryption algorithm.
- **aes-256-cbc**—AES-CBC 256-bit encryption algorithm.
- **des-cbc**—DES-CBC encryption algorithm.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

always-send

Syntax always-send;

Hierarchy Level [edit security ike gateway *gateway-name* dead-peer-detection]

Release Information Statement introduced in Junos OS Release 8.5.

Description Instructs the device to send dead peer detection (DPD) requests regardless of whether there is outgoing IPsec traffic to the peer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

authentication (IPsec SA for OSPF)

Syntax

```
authentication {
  algorithm (hmac-md5-96 | hmac-sha1-96);
  key {
    ascii-text key;
    hexadecimal key;
  }
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure authentication parameters for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **algorithm**—Hash algorithm that authenticates packet data. It can be one of the following:

- **hmac-md5-96**—Produces a 128-bit digest. This is the default.
- **hmac-sha1-96**—Produces a 160-bit digest.

key—Type of authentication key. It can be one of the following:

- **ascii-text *key***—ASCII text key. For **hmac-md5-96**, the key is 16 ASCII characters; for **hmac-sha1-96**, the key is 20 ASCII characters.
- **hexadecimal *key***—Hexadecimal key. For **hmac-md5-96**, the key is 32 hexadecimal characters; for **hmac-sha1-96**, the key is 40 hexadecimal characters.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

authentication (Security IPsec)

Syntax

```
authentication {
  algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
  key (ascii-text key | hexadecimal key );
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement modified in Junos OS Release 8.5. Support for **hmac-sha-256-128** added to high-end SRX Series devices in Junos OS Release 12.1X46-D20.

Description Configure IPsec authentication parameters for a manual security association.

Options

- **algorithm**—Hash algorithm that authenticates packet data. It can be one of the following:
  - **hmac-md5-96**—Produces a 128-bit digest.
  - **hmac-sha-256-128**—Produces a 256-bit digest, truncated to 128 bits.
  - **hmac-sha1-96**—Produces a 160-bit digest.
- **key**—Type of authentication key. It can be one of the following:
  - **ascii-text *key***—ASCII text key. For **hmac-md5-96**, the key is 16 ASCII characters; for **hmac-sha1-96**, the key is 20 ASCII characters.
  - **hexadecimal *key***—Hexadecimal key. For **hmac-md5-96**, the key is 32 hexadecimal characters; for **hmac-sha1-96**, the key is 40 hexadecimal characters.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

authentication-algorithm (Security Group VPN IKE)

Syntax authentication-algorithm (sha-256 | sha-384);

Hierarchy Level [edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the Internet Key Exchange (IKE) authentication algorithm.

NOTE: The device does not delete existing IPsec SAs when you update the **authentication-algorithm** configuration in the IKE proposal. The device deletes existing IPsec SAs when you update the **authentication-algorithm** configuration in the IPsec proposal.

Options **authentication-algorithm**—Hash algorithm that authenticates packet data. It can be one of the following algorithms:

- **sha-256**—Produces a 256-bit digest. This is the default value.
- **sha-384**—Produces a 384-bit digest.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

authentication-algorithm (Security Group VPN IPsec)

Syntax authentication-algorithm hmac-sha-256-128;

Hierarchy Level [edit security group-vpn server ipsec proposal *proposal-name*]

Release Information Statement added in Junos OS Release 10.2.

Description Configure the IPsec authentication algorithm.

Options **hmac-sha-256-128**—Produces a 256-bit digest, truncated to 128 bits. This is the default value.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

authentication-algorithm (Security IKE)

Syntax authentication-algorithm (md5 | sha-256 | sha-384 | sha1);

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **sha-384** option added in Junos OS Release 12.1X45-D10.

Description Configure the Internet Key Exchange (IKE) authentication algorithm.

NOTE: The device does not delete existing IPsec SAs when you update the **authentication-algorithm** configuration in the IKE proposal. The device deletes existing IPsec SAs when you update the **authentication-algorithm** configuration in the IPsec proposal.

Options **authentication-algorithm**—Hash algorithm that authenticates packet data. It can be one of the following algorithms:

- **md5**—Produces a 128-bit digest.
- **sha-256**—Produces a 256-bit digest.
- **sha-384**—Produces a 384-bit digest.
- **sha1**—Produces a 160-bit digest.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

 Copyright © 2016, Juniper Networks, Inc.

authentication-algorithm (Security IPsec)

Syntax authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);

Hierarchy Level [edit security ipsec proposal *proposal-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **hmac-sha-256-128** added to high-end SRX Series devices in Junos OS Release 12.1X46-D20.

Description Configure the IPsec authentication algorithm.

Options The hash algorithm to authenticate data can be one of the following:

- **hmac-md5-96**—Produces a 128-bit digest.
- **hmac-sha-256-128**—Produces a 256-bit digest, truncated to 128 bits.
- **hmac-sha1-96**—Produces a 160-bit digest.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

authentication-method

Syntax authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384 | pre-shared-keys | rsa-signatures);

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **ecdsa-signatures-256** and **ecdsa-signatures-384** options added in Junos OS Release 12.1X45-D10.

Description Specify the method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The **pre-shared-keys** option refers to a preshared key, which is a key for encryption and decryption that both participants must have before beginning tunnel negotiations. The other options refer to types of digital signatures, which are certificates that confirm the identity of the certificate holder.

NOTE: The device does not delete existing IPsec SAs when you update the authentication-method configuration in the IKE proposal.

Options

- **dsa-signatures**—Specify that the Digital Signature Algorithm (DSA) is used.
- **ecdsa-signatures-256**—Specify that the Elliptic Curve DSA (ECDSA) using the 256-bit elliptic curve secp256r1, as specified in the *Federal Information Processing Standard (FIPS) Digital Signature Standard (DSS) 186-3*, is used.
- **ecdsa-signatures-384**—Specify that the ECDSA using the 384-bit elliptic curve secp384r1, as specified in the *FIPS DSS 186-3*, is used.
- **pre-shared-keys**—Specify that a preshared key, which is a secret key shared between the two peers, is used during authentication to identify the peers to each other. The same key must be configured for each peer. This is the default method.
- **rsa-signatures**—Specify that a public key algorithm, which supports encryption and digital signatures, is used.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

 Copyright © 2016, Juniper Networks, Inc.

authentication-method (Security Group VPN)

Syntax authentication-method pre-shared-keys;

Hierarchy Level [edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the method the device uses to authenticate the source of Internet Key Exchange (IKE) messages. The **pre-shared-keys** option refers to a preshared key, which is a key for encryption and decryption that both participants must have before beginning tunnel negotiations.

NOTE: The device does not delete existing IPsec SAs when you update the **authentication-method** configuration in the IKE proposal.

Options

- **pre-shared-keys**—Specify that a preshared key, which is a secret key shared between the two peers, is used during authentication to identify the peers to each other. The same key must be configured for each peer. This is the default method.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

auto-re-enrollment (Security)

Syntax

```
auto-re-enrollment {
  cmpv2 {
    certificate-id certificate-id-name {
      ca-profile-name ca-profile-name ;
      re-enroll-trigger-time-percentage percentage ;
      re-generate-keypair;
    }
  }
  scep {
    certificate-id certificate-id-name {
      ca-profile-name ca-profile-name ;
      challenge-password password ;
      re-enroll-trigger-time-percentage percentage ;
      re-generate-keypair;
    }
  }
}
```

Hierarchy Level [edit security pki]

Release Information Statement modified in Junos OS Release 9.0. **cmpv2** and **scep** keywords and options added in Junos OS Release 15.1X49-D40.

Description Configure the automatic reenrollment of a local end-entity (EE) certificate.

Options **cmpv2**—Configure automatic reenrollment of a local certificate using CMPv2.

scep—Configure automatic reenrollment of a local certificate using Simple Certificate Enrollment Protocol (SCEP).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

auxiliary-spi (IPsec SA for OSPF)

Syntax auxiliary-spi *auxiliary-spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure an auxiliary security parameter index (SPI) for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **auxiliary-spi**—Auxiliary SPI for the manual IPsec SA. The SPI uniquely identifies the SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16639

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

bind-interface

Syntax bind-interface ***interface-name***;

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure the tunnel interface to which the route-based virtual private network (VPN) is bound.

Options ***interface-name*** —Tunnel interface.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ca-identity (Security)

Syntax ca-identity *ca-identity*;

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement modified in Junos OS Release 11.1.

Description Specify the certificate authority (CA) identity to use in requesting digital certificates.

Options

- ***ca-identity*** —Name of CA identity. This name is typically the domain name of the CA.
- ***routing-instance-name*** —Name of routing instance. The routing instance name is chosen from the list of configured routing instances.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

ca-profile (Security PKI)

Syntax

```
ca-profile ca-profile-name {
  administrator {
    e-mail-address e-mail-address;
  }
  ca-identity ca-identity ;
  enrollment {
      retry number;
      retry-interval seconds;
      url url-name;
  }
  revocation-check {
    disable;
    crl {
      disable {
        on-download-failure;
      }
      refresh-interval hours;
      url url-name;
    }
    disable;
    ocsp {
      connection-failure (disable | fallback-crl);
      disable-responder-revocation-check;
      nonce-payload (enable | disable);
      url ocsp-url;
    }
    use-ocsp;
  }
  routing-instance routing-instance-name ;
}
```

Hierarchy Level [edit security pki]

Release Information Statement modified in Junos OS Release 8.5. Support for **ocsp** and **use-ocsp** options added in Junos OS Release 12.1X46-D20.

Description Configure certificate authority (CA) profile.

Options *ca-profile-name* —Name of a trusted CA.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

ca-profile-name

Syntax ca-profile-name *ca-profile-name*;

Hierarchy Level [edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]
[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]

Release Information Statement modified in Junos OS Release 9.0. Support for **[edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]** and **[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]** hierarchies added in Junos OS Release 15.1X49-D40.

Description Specify the name of the certificate authority (CA) profile to be used for automatic reenrollment. The CA certificate must be present to initiate reenrollment.

Options ***ca-profile-name*** —Name of the CA profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

certificate

Syntax
```
certificate {
  local-certificate certificate-id;
  peer-certificate-type (pkcs7 | x509-signature);
  policy-oids [ oid ];
}
```

Hierarchy Level [edit security ike policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. **policy-oids** option added in Junos OS Release 12.3X48-D10.

Description Specify usage of a digital certificate to authenticate the virtual private network (VPN) initiator and recipient.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

certificate-id (Security)

Syntax

```
certificate-id certificate-id-name {
  ca-profile-name ca-profile-name;
  challenge-password password;
  re-enroll-trigger-time-percentage percentage;
  re-generate-keypair;
}
```

Hierarchy Level

[edit security pki auto-re-enrollment cmpv2]
[edit security pki auto-re-enrollment scep]

Release Information

Statement modified in Junos OS Release 9.0. Support for **[edit security pki auto-re-enrollment cmpv2]** and **[edit security pki auto-re-enrollment scep]** hierarchies added in Junos OS Release 15.1X49-D40.

Description

Specify the certificate authority (CA) certificate to use for automatic reenrollment.

NOTE: The **challenge-password** option is only applicable for SCEP reenrollment.

Options

certificate-id-name —Identifier of the end-entity (EE) certificate to be automatically reenrolled. The certificate must be already enrolled for reenrollment to be initiated.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

challenge-password (Security)

Syntax challenge-password *password*;

Hierarchy Level [edit security pki auto-re-enrollment certificate-id *certificate-id-name*]
[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]

Release Information Statement modified in Junos OS Release 9.0. Support for **[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]** hierarchy added in Junos OS Release 15.1X49-D40.

Description Specify the password used by the certificate authority (CA) for enrollment and revocation. If the CA does not provide the challenge password, choose your own password.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

connections-limit

Syntax connections-limit *number*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the number of concurrent connections that the group profile supports. When the maximum number of connections is reached, no more dynamic virtual private network (VPN) endpoints dialup users attempting to access an IPsec VPN are allowed to begin Internet Key Exchange (IKE) negotiations.

Options *number*—Maximum number of concurrent connections allowed.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041
- *Monitoring VPNs*

Copyright © 2016, Juniper Networks, Inc.

container

Syntax container *container-string*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic distinguished-name]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that the value in the identity fields of a dynamic virtual private network (VPN) endpoint user's distinguished name exactly match the values in the group IKE user's distinguished name. The order of the identity fields in the fields of the distinguished name strings must be identical when matching.

Options ***container-string***—Distinguished name identity value to be matched. For example, **cn**=admin, **ou**=eng, **o**=example, **dc**=net.

NOTE: Add a space between each container string. For example, **edit security ike gateway jsr_gateway dynamic distinguished-name container o=example, ou=eng;**

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- *Monitoring VPNs*

Copyright © 2016, Juniper Networks, Inc.

crl (Security)

Syntax

```
crl {
  disable {
    on-download-failure;
  }
  refresh-interval hours;
  url url-name;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name* revocation-check]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

Options

- **disable on-download-failure**—(Optional) Override the default behavior and permit certificate verification even if the CRL fails to download.
- **refresh-interval** *hours*—Time interval, in hours, between CRL updates.

  Range — 0 through 8784 hours.
- **url** ***url-name*** —Name of the location from which to retrieve the CRL through HTTP or Lightweight Directory Access Protocol (LADP). You can specify one URL for each configured CA profile. By default, no location is specified. Use a fully qualified domain name (FQDN) or an IP address and, optionally, a port number. If no port number is specified, port 80 is used for HTTP and port 443 is used for LDAP.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

cryptographic-self-test

Syntax cryptographic-self-test;

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm when the device or switch detects a cryptographic self-test failure. Cryptographic self-tests are a set of preoperational tests that are performed after the device or switch is powered on. The self-test run without operator intervention.

Default No alarm is raised upon failure of a cryptographic self-test.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

dead-peer-detection

Syntax
```
dead-peer-detection {
  (always-send | optimized | probe-idle-tunnel);
  interval seconds;
  threshold number;
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **optimized** and **probe-idle-tunnel** options added in Junos OS Release 12.1X46-D10.

Description Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peers. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding AutoVPN on page 6451
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

dead-peer-detection (Security Group VPN Server)

Syntax

```
dead-peer-detection {
  always-send;
  interval seconds;
  threshold number;
}
```

Hierarchy Level [edit security group-vpn server ike gateway *gateway-name*]

Release Information Support for the Group VPN server added in Junos OS Release 15.1X49-D30 for vSRX.

Description Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peers. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

decryption-failures

Syntax

```
decryption-failures {
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm after exceeding a specified number of decryption failures.

Default Multiple decryption failures do not cause an alarm to be raised.

Options *failures*—Number of decryption failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

Range: 0 through 1 through 1,000,000,000.

Default: 1000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

description (Security Policies)

Syntax description *description*;

Hierarchy Level [edit security group-vpn member ike policy *policy-name*]
[edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike policy *policy-name*]
[edit security group-vpn server ipsec proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]
[edit security ike policy *policy-name*],
[edit security ike proposal *proposal-name*],
[edit security ipsec policy *policy-name*],
[edit security ipsec proposal *proposal-name*]
[edit security polices from-zone *zone-name* to-zone *zone-name* policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2. Support for the **security policies** hierarchy added in Junos OS Release 12.1.

Description Specify descriptive text for an IKE policy, an IPsec policy, an IKE proposal, an IPsec proposal, or a security policy.

Options ***description*** —Descriptive text about an IKE policy, an IPsec policy, an IKE proposal, an IPsec proposal, or a security policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- [edit security policies] Hierarchy Level on page 320
- IPsec VPN Overview on page 6041

destination-ip (Security IPsec)

Syntax destination-ip ***ip-address***;

Hierarchy Level [edit security ipsec vpn *vpn-name* vpn-monitor]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the destination of the Internet Control Message Protocol (ICMP) pings. If this statement is used, the device uses the peer's gateway address by default.

Options ***ip-address***—Destination IP address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

df-bit

Syntax df-bit (clear | copy | set);

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify how the device handles the Don't Fragment (DF) bit in the outer header.

NOTE: On high-end SRX Series devices, the DF-bit configuration for VPN only works if the original packet size is smaller than the st0 interface MTU, and larger than the external interface-ipsec overhead.

Options

- **clear**—Clear (disable) the DF bit from the outer header. This is the default.
- **copy**—Copy the DF bit to the outer header.
- **set**—Set (enable) the DF bit in the outer header.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

dh-group (Security IKE)

Syntax dh-group (group1 | group14 | group19 | group2 | group20 | group24| group5);

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for the **group14** option added in Junos OS Release 11.1. Support for **group19**, **group20**, and **group24** options added in Junos OS Release 12.1X45-D10.

Description Specify the IKE Diffie-Hellman group.

NOTE: The device does not delete existing IPsec SAs when you update the **dh-group** configuration in the IKE proposal.

Options **dh-group**—Diffie-Hellman group for key establishment.

- **group1**—768-bit Modular Exponential (MODP) algorithm.
- **group14**—2048-bit MODP group.
- **group19**—256-bit random Elliptic Curve Groups modulo a Prime (ECP groups) algorithm.
- **group2**—1024-bit MODP algorithm.
- **group20**—384-bit random ECP groups algorithm.
- **group24**—2048-bit MODP Group with 256-bit prime order subgroup.
- **group5**—1536-bit MODP algorithm.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

dh-group (Security Group VPN IKE)

Syntax dh-group (group14 | group24);

Hierarchy Level [edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 10.2. Support for the **group14** option added in Junos OS Release 11.1. Support for the **group24** option added in Junos OS Release 15.1X49-D30 for vSRX.

Description Specify the IKE Diffie-Hellman group.

NOTE: The device does not delete existing IPsec SAs when you update the **dh-group** configuration in the IKE proposal.

Options **dh-group**—Diffie-Hellman group for key establishment.

- **group14**—2048-bit group. This is the default value.
- **group24**—2048-bit, 256 bit subgroup.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

disable (PKI)

Syntax disable;

Hierarchy Level [edit security pki ca-profile *profile-name* revocation-check]

Release Information Statement modified in Junos OS Release 9.0.

Description Disable revocation checks.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

distinguished-name (Security)

Syntax distinguished-name <container *container-string*> <wildcard *wildcard-string*>

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a distinguished name as the identifier for the remote gateway with a dynamic IP address.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

dynamic (Security)

Syntax

```
dynamic {
  connections-limit number;
  (distinguished-name <container container-string> <wildcard wildcard-string> | hostname
    domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-address);
  ike-user-type (group-ike-id | shared-ike-id);
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for the **inet6** option added in Junos OS Release 11.1.

Description Specify the identifier for the remote gateway with a dynamic IPv4 or IPv6 address. Use this statement to set up a VPN with a gateway that has an unspecified IPv4 or IPv6 address.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

dynamic (Security Group VPN)

Syntax

```
dynamic {
   (hostname domain-name | inet ip-address | user-at-hostname e-mail-address);
}
```

Hierarchy Level [edit security group-vpn server ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the identifier for the remote gateway with a dynamic IPv4 address. Use this statement to set up a VPN with a gateway that has an unspecified IPv4 address.

NOTE: Configuring **mode main** for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is **pre-shared-keys**.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

encryption (IPsec SA for OSPF)

Syntax

```
encryption {
  algorithm (3des-cbc | des-cbc | null);
  key {
    ascii-text key;
    hexadecimal key;
  }
}
```

Hierarchy Level [edit security ipsec security-association *sa-name* manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure encryption parameters for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **algorithm**—Type of encryption algorithm. It can be one of the following:

- **3des-cbc**—Has block size of 8 bytes (64 bits); its key size is 192 bits long.
- **des-cbc**—Has a block size of 8 bytes (64 bits); its key size is 48 bits long.
- **null**—With null encryption, you are choosing not to provide encryption on OSPFv3 headers.

key—Type of encryption key. It can be one of the following:

- **ascii-text *key***—ASCII text key. For the **des-cbc** option, the key contains 8 ASCII characters; for **3des-cbc**, the key contains 24 ASCII characters.
- **hexadecimal *key***—Hexadecimal key. For the **des-cbc** option, the key contains 16 hexadecimal characters; for the **3des-cbc** option, the key contains 48 hexadecimal characters.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

encryption (Security)

Syntax

```
encryption {
  algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
  key (ascii-text key | hexadecimal key) ;
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure an encryption algorithm and key for a manual Security Association (SA).

Options

- **algorithm**—Type of encryption algorithm. It can be one of the following:
  - **des-cbc**—Has a block size of 8 bytes (64 bits); its key size is 48 bits long.
  - **3des-cbc**—Has block size of 8 bytes (64 bits); its key size is 192 bits long

    NOTE: For 3des-cbc, we recommend that the first 8 bytes be different from the second 8 bytes, and the second 8 bytes be the same as the third 8 bytes.
  - **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
  - **aes-192-cbc**—Advanced Encryption Standard (AES) 192-bit encryption algorithm.
  - **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption algorithm.
- **key**—Type of encryption key. It can be one of the following:
  - **ascii-text key**—ASCII text key. For the **des-cbc** option, the key contains 8 ASCII characters; for **3des-cbc**, the key contains 24 ASCII characters.
  - **hexadecimal key**—Hexadecimal key. For the **des-cbc** option, the key contains 16 hexadecimal characters; for the **3des-cbc** option, the key contains 48 hexadecimal characters.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

encryption-algorithm (Security Group VPN IKE)

Syntax encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);

Hierarchy Level [edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2.

Description Configure an encryption algorithm for an IKE proposal.

NOTE: The device does not delete existing IPsec SAs when you update the **encryption-algorithm** configuration in the IKE proposal.

Options

- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—AES 192-bit encryption algorithm.
- **aes-256-cbc**—AES 256-bit encryption algorithm.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

encryption-algorithm (Security Group VPN IPsec)

Syntax encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);

Hierarchy Level [edit security group-vpn server ipsec proposal *proposal-name*]
[edit security group-vpn member ipsec proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure an encryption algorithm.

NOTE: The device deletes existing IPsec SAs when you update the encryption-algorithm configuration in the IPsec proposal.

Options

- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-192-cbc**—AES 192-bit encryption algorithm.
- **aes-256-cbc**—AES 256-bit encryption algorithm. This is the default value.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

encryption-algorithm (Security IKE)

Syntax encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-256-cbc | aes-256-gcm | des-cbc);

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **aes-128-gcm** and **aes-256-gcm** options added in Junos OS Release 15.1X49-D40.

Description Configure an encryption algorithm for an IKE proposal.

NOTE: The device does not delete existing IPsec SAs when you update the **encryption-algorithm** configuration in the IKE proposal.

Options

3des-cbc—Has a block size of 24 bytes; the key size is 192 bits long.

aes-128-cbc—Advanced Encryption Standard (AES) 128-bit encryption algorithm.

aes-128-gcm—AES 128-bit authenticated encryption algorithm supported with IKEv2 only. When this option is used, **aes-128-gcm** should be configured at the [**edit security ipsec proposal** *proposal-name*] hierarchy level, and the **authentication-algorithm** option should not be configured at the [**edit security ike proposal** *proposal-name*] hierarchy level.

aes-192-cbc—AES 192-bit encryption algorithm.

aes-256-cbc—AES 256-bit encryption algorithm.

aes-256-gcm—AES 256-bit authenticated encryption algorithm supported with IKEv2 only. When this option is used, **aes-256-gcm** should be configured at the [**edit security ipsec proposal** *proposal-name*] hierarchy level, and the **authentication-algorithm** option should not be configured at the [**edit security ike proposal** *proposal-name*] hierarchy level.

des-cbc—Has a block size of 8 bytes; the key size is 48 bits long.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

encryption-algorithm (Security IPsec)

Syntax encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm | aes-256-cbc | aes-256-gcm | des-cbc);

Hierarchy Level [edit security ipsec proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **aes-128-gcm**, **aes-192-gcm**, and **aes-256-gcm** options added in Junos OS Release 12.1X45-D10.

Description Configure an encryption algorithm.

NOTE: The device deletes existing IPsec SAs when you update the encryption-algorithm configuration in the IPsec proposal.

Options

- **3des-cbc**—Has a block size of 24 bytes; the key size is 192 bits long.
- **aes-128-cbc**—Advanced Encryption Standard (AES) 128-bit encryption algorithm.
- **aes-128-gcm**—AES Galois/Counter Mode (GCM) 128-bit encryption algorithm. This option is for IPsec proposals only.
- **aes-192-cbc**—AES 192-bit encryption algorithm.
- **aes-192-gcm**—AES GCM 192-bit encryption algorithm. This option is for IPsec proposals only.
- **aes-256-cbc**—AES 256-bit encryption algorithm.
- **aes-256-gcm**—AES GCM 256-bit encryption algorithm. This option is for IPsec proposals only.
- **des-cbc**—Has a block size of 8 bytes; the key size is 48 bits long.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- *Monitoring VPNs*

Copyright © 2016, Juniper Networks, Inc.

encryption-failures

Syntax

```
encryption-failures {
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm after exceeding a specified number of encryption failures.

Default Multiple encryption failures do not cause an alarm to be raised.

Options *failures*—Number of encryption failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

Range: 1 through 1,000,000,000.

Default: 1000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

enrollment (Security)

Syntax

```
enrollment {
  retry number;
  retry-interval seconds;
  url url-name;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the enrollment parameters for a certificate authority (CA).

Options

- **retry** ***number*** —Number of automated attempts for online enrollment to be retried in case enrollment response is pending.

  Range: 0 through 1080
  Default: 10

- **retry-interval** ***seconds*** —Time interval, in seconds, between the enrollment retries.

  Range: 0 through 3600
  Default: 900 seconds

- **url** ***url-name*** —Enrollment URL where the Simple Certificate Enrollment Protocol (SCEP) or CMPv2 request is sent to the certification authority (CA) as configured in this profile.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

establish-tunnels

Syntax establish-tunnels (immediately | on-traffic);

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify when IKE is activated: immediately after VPN information is configured and configuration changes are committed, or only when data traffic flows. In the second case, IKE needs to be negotiated with the peer gateway.

Options
- **immediately**—IKE is activated immediately after VPN configuration and configuration changes are committed.
- **on-traffic**—IKE is activated only when data traffic flows and must to be negotiated.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

external-interface (Security IKE Gateway)

Syntax external-interface *external-interface-name*;

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the outgoing interface for IKE SAs. This interface is associated with a zone that acts as its carrier, providing firewall security for it.

Options ***external-interface-name*** —Name of the interface to be used to send traffic to the IPsec VPN.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

external-interface (Security Manual SA)

Syntax external-interface *external-interface-name*;

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the outgoing interface for the manual SA.

Options ***external-interface-name*** —Name of the outgoing interface.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

gateway (Security Group VPN Member IKE)

Syntax

```
gateway gateway-name {
   ike-policy policy-name;
   local address ip-address;
   local-identity {
      (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
         e-mail-address);
   }
   remote-identity {
      (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
   }
   routing-instance routing-instance;
   server-address [ip-address];
}
```

Hierarchy Level [edit security group-vpn member ike]

Release Information Statement introduced in Junos OS Release 10.2. Support for the **routing-instance** option added in Junos OS Release 15.1X49-D30 for vSRX.

Description Configure IKE gateway for group VPN member.

Options ***gateway-name*** —Name of the gateway.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

gateway (Security Group VPN Server IKE)

Syntax

```
gateway gateway-name {
  address ip-address ;
  dead-peer-detection {
    always-send;
    interval seconds;
    threshold number;
  }
  dynamic {
    (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
  }
  ike-policy policy-name;
  local-address ip-address;
  local-identity {
    (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
  }
  remote-identity {
    (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
  }
  routing-instance routing-instance;
}
```

Hierarchy Level [edit security group-vpn server ike]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure IKE gateway for group VPN server.

Options ***gateway-name*** —Name of the gateway.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

gateway (Security IKE)

Syntax

```
gateway gateway-name {
  address [ip-address-or-hostname];
  advpn {
    suggester {
      disable;
    }
    partner {
      connection-limit <number>;
      idle-threshold <packets/sec>;
      idle-time <seconds>;
      disable;
    }
  }
  dead-peer-detection {
    (always-send | optimized | probe-idle-tunnel);
    interval seconds;
    threshold number;
  }
  dynamic {
    connections-limit number;
    (distinguished-name <container container-string> <wildcard wildcard-string> | hostname
      domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-address);
    ike-user-type (group-ike-id | shared-ike-id);
  }
  external-interface external-interface-name;
  general-ikeid;
  ike-policy policy-name;
  local-address (ipv4-address | ipv6-address);
  local-identity {
    (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address |
      user-at-hostname e-mail-address);
  }
  nat-keepalive seconds;
  no-nat-traversal;
  remote-identity {
    (distinguished-name <container container-string> <wildcard wildcard-string> | hostname
      hostname | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-address);
  }
  version (v1-only | v2-only);
  xauth {
    access-profile profile-name;
  }
}
```

Hierarchy Level [edit security ike]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The **inet6** option added in Junos OS Release 11.1. Support for the **advpn** option added in Junos OS Release 12.3X48-D10.

Description Configure an IKE gateway.

Copyright © 2016, Juniper Networks, Inc.

Options ***gateway-name*** —Name of the gateway.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

gateway (Security IPsec VPN)

Syntax gateway *ip-address*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the IP address of the peer.

Options *ip-address*—IP address of the peer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

gateway (Security Manual SA)

Syntax gateway ***ip-address***;

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description For a manual security association, specify the IPv4 or IPv6 address of the peer.

Options ***ip-address*** —IPv4 or IPv6 address of the peer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

general-ikeid

Syntax general-ikeid;

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Accept general peer IKE ID.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

group (Security Group VPN)

Syntax

```
group name {
  anti-replay-time-window milliseconds;
  description description;
  group-id number;
  ike-gateway gateway-name;
  ipsec-sa name {
    match-policy policy-name {
      destination ip-address/netmask;
      destination-port number;
      protocol number;
      source ip-address/netmask;
      source-port number;
    }
    proposal proposal-name;
  }
  member-threshold number;
  server-cluster {
    ike-gateway gateway-name;
    retransmission-period seconds;
    server-role (root-server | sub-server);
  }
  server-member-communication {
    certificate certificate-id;
    communication-type unicast;
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
    number-of-retransmission number;
    retransmission-period seconds;
    sig-hash-algorithm (sha-256 | sha-384);
  }
}
```

Hierarchy Level [edit security group-vpn server]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure group VPN on the group server.

Options *name*—Name of the group.

- **anti-replay-time-window** *milliseconds*—Configure antireplay time in milliseconds. Specify a value from 1 to 60,000.

NOTE: We recommend that NTP be configured on Group VPNv2 devices to ensure proper antireplay operation.

Copyright © 2016, Juniper Networks, Inc.

NOTE: Group members that are running on vSRX instances on a host machine where the hypervisor is running under a heavy load may experience issues that can be corrected by reconfiguring the **anti-replay-time-window** value. If data that matches the IPsec policy on the group member is not being transferred, check the **show security group-vpn member ipsec statistics** output for D3P errors. Make sure that NTP is operating correctly. If there are errors, adjust the **anti-replay-time-window** value.

- **description** ***description***—Description of the group.
- **group-id** ***number***—Identifier for this group VPN. Specify a value from 1 to 4,294,967,295.
- **ike-gateway** ***gateway-name***—Define the group member for Phase 1 negotiation. There can be multiple instances of this option configured. When a group member sends its registration request to the server, the server checks to see that the member is configured for the group.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

group-vpn

Syntax

```
group-vpn {
  member {
    ike {
      gateway gateway-name {
        ike-policy policy-name;
        local address ip-address;
        local-identity {
          (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
            e-mail-address);
        }
        remote-identity {
          (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
        }
        routing-instance routing-instance;
        server-address [ip-address];
      }
      policy policy-name {
        description description;
        mode (aggressive | main);
        pre-shared-key (ascii-text key | hexadecimal key);
        proposals [proposal-name];
      }
      proposal proposal-name {
        authentication-algorithm (sha-256 | sha-384);
        authentication-method pre-shared-keys;
        description description;
        dh-group (group14 | group24);
        encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        lifetime-seconds seconds;
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        gateway-filter {
          local-address ip-address;
          remote-address ip-address;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
    }
    ipsec {
      vpn vpn-name {
        df-bit (clear | copy | set);
        exclude rule rule-name {
          source-address ip-address/mask;
```

Copyright © 2016, Juniper Networks, Inc.

```
            destination-address ip-address/mask;
            application application;
          }
          fail-open rule rule-name {
            source-address ip-address/mask;
            destination-address ip-address/mask;
            application application;
          }
          group id;
          group-vpn-external-interface interface;
          ike-gateway gateway-name;
          recovery-probe;
        }
      }
    }
    server {
      group name {
        anti-replay-time-window milliseconds;
        description description;
        group-id number;
        ike-gateway gateway-name;
        ipsec-sa name {
          match-policy policy-name {
            destination ip-address/netmask;
            destination-port number;
            protocol number;
            source ip-address/netmask;
            source-port number;
          }
          proposal proposal-name;
        }
        member-threshold number;
        server-cluster {
          ike-gateway gateway-name;
          retransmission-period seconds;
          server-role (root-server | sub-server);
        }
        server-member-communication {
          certificate certificate-id;
          communication-type unicast;
          encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
          lifetime-seconds seconds;
          number-of-retransmission number;
          retransmission-period seconds;
          sig-hash-algorithm (sha-256 | sha-384);
        }
      }
      ike {
        gateway gateway-name {
          address ip-address ;
          dead-peer-detection {
            always-send;
            interval seconds;
            threshold number;
          }
          dynamic {
```


Copyright © 2016, Juniper Networks, Inc.

```
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
          }
          ike-policy policy-name;
          local-address ip-address;
          local-identity {
            (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
          }
          remote-identity {
            (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
          }
          routing-instance routing-instance;
        }
        policy policy-name {
          description text;
          mode (aggressive | main);
          pre-shared-key (ascii-text key | hexadecimal key);
          proposals [proposal-name];
        }
        proposal proposal-name {
          authentication-algorithm (sha-256 | sha-384);
          authentication-method pre-shared-keys;
          description description;
          dh-group (group14 | group24);
          encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
        }
      }
      ipsec {
        proposal proposal-name {
          authentication-algorithm hmac-sha-256-128;
          description description;
          encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
          lifetime-seconds seconds;
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        gateway-filter {
          local-address ip-address;
          remote-address ip-address;
        }
        level (all | error | info | notice | verbose | warning);
        no-remote-trace;
      }
    }
  }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure group VPNs.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

hostname

Syntax hostname *domain-name*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Unique name by which a network-attached device is known on a network.

Options ***domain-name***—A fully qualified domain name (FQDN).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

idle-time

Syntax idle-time *seconds*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum amount of idle time to delete a security association (SA).

Options ***seconds*** —Maximum amount of idle time.
Range: 60 through 999,999 seconds
Default: To be disabled

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ike (Security)

Syntax

```
ike {
  gateway gateway-name {
    address [ip-address-or-hostname];
    advpn {
      suggester {
        disable;
      }
      partner {
        connection-limit <number>;
        idle-threshold <packets/sec>;
        idle-time <seconds>;
        disable;
      }
    }
    dead-peer-detection {
      (always-send | optimized | probe-idle-tunnel);
      interval seconds;
      threshold number;
    }
    dynamic {
      connections-limit number;
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
      ike-user-type (group-ike-id | shared-ike-id);
    }
    external-interface external-interface-name;
    general-ikeid;
    ike-policy policy-name;
    local-address (ipv4-address | ipv6-address);
    local-identity {
      (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address |
        user-at-hostname e-mail-address);
    }
    nat-keepalive seconds;
    no-nat-traversal;
    remote-identity {
      (distinguished-name <container container-string> <wildcard wildcard-string> |
        hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    version (v1-only | v2-only);
    xauth {
      access-profile profile-name;
    }
  }
  policy policy-name {
    certificate {
      local-certificate certificate-id;
      peer-certificate-type (pkcs7 | x509-signature);
      policy-oids [ oid ];
    }
```

Copyright © 2016, Juniper Networks, Inc.

```
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (md5 | sha-256 | sha-384| sha1);
    authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384
      | pre-shared-keys | rsa-signatures);
    description description;
    dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    lifetime-seconds seconds;
  }
  respond-bad-spi <max-responses>;
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    rate-limit messages-per-second;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The **inet6** option added in Junos OS Release 11.1.

Description Define Internet Key Exchange (IKE) configuration.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- ALG Overview on page 3
- Understanding Logical Systems for SRX Series Services Gateways on page 3205

Copyright © 2016, Juniper Networks, Inc.

ike (Security Group VPN Member)

Syntax

```
ike {
  gateway gateway-name {
    ike-policy policy-name;
    local address ip-address;
    local-identity {
      (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
        e-mail-address);
    }
    remote-identity {
      (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
    }
    routing-instance routing-instance;
    server-address [ip-address];
  }
  policy policy-name {
    description description;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (sha-256 | sha-384);
    authentication-method pre-shared-keys;
    description description;
    dh-group (group14 | group24);
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    gateway-filter {
      local-address ip-address;
      remote-address ip-address;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security group-vpn member]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure IPsec group VPN on the group member.

Copyright © 2016, Juniper Networks, Inc.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

ike (Security Group VPN Server)

Syntax

```
ike {
  gateway gateway-name {
    address ip-address ;
    dead-peer-detection {
      always-send;
      interval seconds;
      threshold number;
    }
    dynamic {
      (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    ike-policy policy-name;
    local-address ip-address;
    local-identity {
      (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
    }
    remote-identity {
      (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
    }
    routing-instance routing-instance;
  }
  policy policy-name {
    description text;
    mode (aggressive | main);
    pre-shared-key (ascii-text key | hexadecimal key);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (sha-256 | sha-384);
    authentication-method pre-shared-keys;
    description description;
    dh-group (group14 | group24);
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
  }
}
```

Hierarchy Level [edit security group-vpn server]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure Phase 1 security association (SA) with a member on the group server. The gateway is the group member.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

ike (Security IPsec VPN)

Syntax

```
ike {
    gateway gateway-name;
    idle-time seconds;
    install-interval seconds;
    ipsec-policy ipsec-policy-name;
    no-anti-replay;
    proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
    }
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description Define an IKE-keyed IPsec VPN.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ike-phase1-failures

Syntax

```
ike-phase1-failures {
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) Phase 1 failures.

Default Multiple IKE phase 1 failures do not cause an alarm to be raised.

Options ***failures***—Number of IKE phase 1 failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

Range: 1 through 1,000,000,000.

Default: 20

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ike-phase2-failures

Syntax

```
ike-phase2-failures {
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm after exceeding a specified number of Internet Key Exchange (IKE) phase 2 failures.

Default Multiple IKE phase 2 failures do not cause an alarm to be raised.

Options *failures*—Number of IKE phase 2 failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

Range: 1 through 1,000,000,000.

Default: 20

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

ike-policy (Security Gateway)

Syntax ike-policy *policy-name*;

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the IKE policy to be used for the gateway.

Options ***policy-name*** —IKE policy name.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ike-user-type

Syntax ike-user-type (group-ike-id | shared-ike-id);

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the type of IKE user for a remote access connection.

Options

- **group-ike-id**—E-mail address or fully qualified domain name (FQDN) shared by a group of remote access users so that each user does not need to configure a separate IKE profile. When group IKE IDs are configured, the IKE ID of each user is a concatenation of a user-specific part and a part that is common to all group IKE ID users. For example, the user Bob might use "Bob.example.net" as his full IKE ID, where ".example.net" is common to all users. The full IKE ID is used to uniquely identify each user connection. Group IKE IDs require the generation of a unique preshared key based on the username supplied during VPN connection, which can be viewed with the **show security ike pre-shared-key** command.
- **shared-ike-id**—E-mail address shared by a large number of remote access users so that each user does not need to configure a separate IKE profile. When a shared IKE ID is configured, all users share a single IKE ID and a single IKE preshared key. Each user is authenticated through the mandatory XAuth phase, where the credentials of individual users are verified either with an external RADIUS server or with a local access database. XAuth is required for shared IKE IDs.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

inet (Security Dynamic Peer)

Syntax inet *ip-address*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify IP address to identify the dynamic peer.

Options *ip-address*—IP address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

inet6 (Security IKE Gateway)

Syntax inet6 *ipv6-address*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 11.1.

Description Specify an IPv6 address to identify the dynamic peer.

Options ***ipv6-address***—IPv6 address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

install-interval

Syntax install-interval *seconds*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum number of seconds to allow for the installation of a rekeyed outbound security association (SA) on the device.

Options ***seconds*** —Maximum amount of idle time.
Range: 0 through 10 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

interval (Security IKE)

Syntax interval *seconds*;

Hierarchy Level [edit security ike gateway *gateway-name* dead-peer-detection]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the amount of time that the peer waits for traffic from its destination peer before sending a dead-peer-detection (DPD) request packet.

Options ***seconds*** —Number of seconds that the peer waits before sending a DPD request packet.

Range: 10 through 60 seconds

Default: 10 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ipsec (Security)

Syntax

```
ipsec {
  policy policy-name {
    description description;
    perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 |
      group5);
    proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
    proposals [proposal-name];
  }
  proposal proposal-name {
    authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
    description description;
    encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm
      | aes-256-cbc | aes-256-gcm | des-cbc);
    lifetime-kilobytes kilobytes;
    lifetime-seconds seconds;
    protocol (ah | esp);
  }
  security-association sa-name {
    manual {
      direction bidirectional {
        authentication {
          algorithm (hmac-md5-96 | hmac-sha1-96);
          key {
            ascii-text key;
            hexadecimal key;
          }
        }
        auxiliary-spi auxiliary-spi-value;
        encryption {
          algorithm (3des-cbc | des-cbc | null);
          key {
            ascii-text key;
            hexadecimal key;
          }
        }
        protocol (ah | esp);
        spi spi-value;
      }
    }
    mode transport;
  }
  traceoptions {
    flag flag;
  }
  vpn vpn-name {
    bind-interface interface-name;
    copy-outer-dscp;
    ike {
    gateway gateway-name;
    ipsec-policy ipsec-phase2-policy;
    }
    establish-tunnels (immediately | on-traffic);
```


Copyright © 2016, Juniper Networks, Inc.

```
    }
    ike {
      gateway gateway-name;
      idle-time seconds;
      install-interval seconds;
      ipsec-policy ipsec-policy-name;
      no-anti-replay;
      proxy-identity {
        local ip-prefix;
        remote ip-prefix;
        service (any | service-name);
      }
    }
    manual {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        key (ascii-text key | hexadecimal key);
      }
      encryption {
        algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
        key (ascii-text key | hexadecimal key);
      }
      external-interface external-interface-name;
      gateway ip-address;
      protocol (ah | esp);
      spi spi-value;
    }
    traffic-selector traffic-selector-name {
      local-ip ip-address/netmask;
      remote-ip ip-address/netmask;
    }
    vpn-monitor {
      destination-ip ip-address;
      optimized;
      source-interface interface-name;
    }
  }
  vpn-monitor-options {
    interval seconds;
    threshold number;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5.

Description Define IPsec configuration.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ipsec (Security Group VPN Member)

Syntax

```
ipsec {
  vpn vpn-name {
    df-bit (clear | copy | set);
    exclude rule rule-name {
      source-address ip-address/mask;
      destination-address ip-address/mask;
      application application;
    }
    fail-open rule rule-name {
      source-address ip-address/mask;
      destination-address ip-address/mask;
      application application;
    }
    group id;
    group-vpn-external-interface interface;
    ike-gateway gateway-name;
    recovery-probe;
  }
}
```

Hierarchy Level [edit security group-vpn member]

Release Information Statement introduced in Junos OS Release 10.2. **df-bit**, **exclude rule**, **fail-open rule**, and **recovery-probe** options added in Junos OS Release 15.1X49-D30 for vSRX.

Description Configure IPsec for Phase 2 exchange on the group member.

Options **vpn *vpn-name***—Name of the VPN.

df-bit—Specifies pre-fragmentation and post-fragmentation of IPsec traffic on the group member. One of the following options can be configured:

- clear—Sets the outer IP do not fragment (DF) bit to 0. When the packet size is larger than the path maximum transmission unit (path MTU), pre-fragmentation is done if the DF bit is not set in the inner packet and post-fragmentation is done if the DF bit is set in the inner packet. This is the default.
- copy—Copies the DF bit from the inner header to the outer header. When the packet size is larger than the path PMTU, pre-fragmentation is done if the DF bit is not set in the inner packet. If the DF bit is set in the inner packet, the packet is dropped and an ICMP message is sent back.
- set—Sets the outer IP DF bit to 1. When the packet size is larger than the path MTU, pre-fragmentation is done if the DF bit is not set in the inner packet. If the DF bit is set in the inner packet, the packet is dropped and an ICMP message is sent back

exclude rule—Specifies traffic to be excluded from Group VPN encryption. A maximum of 10 exclude rules can be configured. Source and destination addresses must be specified in *ip-address/mask* format; address books and address sets are not

Copyright © 2016, Juniper Networks, Inc.

supported. Predefined and user-defined applications are supported, but application sets are not supported.

fail-open rule—Specifies the traffic to be sent in cleartext mode if there is no valid SA key available to protect the traffic. Traffic that is not specified by the fail-open rule is blocked if there is no valid SA key available to protect the traffic. A maximum of 10 fail-open rules can be configured. Source and destination addresses must be specified in *ip-address/mask* format; address books and address sets are not supported. Predefined and user-defined applications are supported, but application sets are not supported.

group *id*—Identifier configured for the Group VPN.

group-vpn-external-interface *interface*—Interface used by the group member to connect to the Group VPN peers. The interface must belong to the same zone as the **to-zone** configured at the [**edit security ipsec-policy**] hierarchy level for Group VPN traffic.

ike-gateway *gateway-name*—Name of the IKE gateway for the Group VPN.

recovery-probe—Enables initiation of **groupkey-pull** exchanges at specific intervals to update the member's SA from the group server if the group member is determined to be out of synchronization with the group server and other group members. This option is disabled by default.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

ipsec (Security Group VPN Server)

Syntax

```
ipsec {
  proposal proposal-name {
    authentication-algorithm hmac-sha-256-128;
    description description;
    encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    lifetime-seconds seconds;
  }
}
```

Hierarchy Level [edit security group-vpn server]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure IPsec proposal for Phase 2 exchange on the group server.

Options **proposal *proposal-name***—Name of the proposal. The proposal name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

ipsec-performance-acceleration (Security Flow)

Syntax ipsec-performance-acceleration;

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Enables IPsec VPN performance acceleration.

Options None.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041
- **show security flow status on page 2164**

Copyright © 2016, Juniper Networks, Inc.

ipsec-policy (Security)

Syntax ipsec-policy *ipsec-policy-name*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the IPsec policy name.

Options ***ipsec-policy-name*** —Name of the IPsec policy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

ipsec-policy (Security Group VPN)

Syntax ipsec-policy from-zone *zone-name* to-zone *zone-name* ipsec-group-vpn *vpn-name*;

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Specifies that matching traffic is checked against rules associated with the specified Group VPN. Exclude and fail-open rules are configured at the [**edit security group-vpn member ipsec vpn *vpn-name***] hierarchy level.

Options **from-zone *zone-name***—Specify the incoming zone for Group VPN traffic.

to-zone *zone-name*—Specify the outgoing zone for Group VPN traffic.

NOTE: The **to-zone** zone must include the interface configured with the **group-vpn-external-interface** option at the [**edit security group-vpn member ipsec vpn *vpn-name***] hierarchy level.

ipsec-group-vpn *vpn-name*—Specify the Group VPN to which the traffic applies. Only one Group VPN can be referenced by a specific from-zone/to-zone pair.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Traffic Steering
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

ipsec-sa (Security Group VPN)

Syntax

```
ipsec-sa name {
  match-policy policy-name {
    destination ip-address/netmask;
    destination-port number;
    protocol number;
    source ip-address/netmask;
    source-port number;
  }
  proposal proposal-name;
}
```

Hierarchy Level [edit security group-vpn server group *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the group SAs to be downloaded to members. There can be multiple group SAs downloaded to group members.

Options **ipsec-sa *name***—Define the group SAs to be downloaded to members.

- **match-policy *policy-name***—Configure the group policy with source address, source port, destination address, destination port, and protocol.

  Use **0.0.0.0** to specify any source or destination. Use **0** to specify any source port, destination port, or protocol.

- **proposal *proposal-name***—Specify the name of the IPsec proposal configured with the **proposal** configuration statement at the [**edit security group-vpn server ipsec**] hierarchy.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

ipsec-vpn (Security Flow)

Syntax

```
ipsec-vpn {
  mss value;
}
```

Hierarchy Level [edit security flow tcp-mss]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the TCP maximum segment size (TCP MSS) for the TCP packets that are about to go into an IPsec VPN tunnel. This value overrides the value specified in the **all-tcp-mss** statement.

Options

mss *value*—TCP MSS value for TCP packets entering an IPsec VPN tunnel. Value is optional.

Range: 64 through 65,535 bytes

Default: 1320 bytes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

key-generation-self-test

Syntax key-generation-self-test;

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm when the device or switch detects a key generation self-test failure. Key generation is the process of generating keys for cryptography. A key is used to encrypt and decrypt data. The self-tests run without operator intervention.

Default No alarm is raised upon failure of a key generation self-test.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

lifetime-kilobytes

Syntax lifetime-kilobytes *kilobytes*;

Hierarchy Level [edit security ipsec proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the lifetime (in kilobytes) of an IPsec security association (SA).

Options ***kilobytes*** —Lifetime of the IPsec security association (SA). If this statement is not configured, the number of kilobytes used for the SA lifetime is unlimited.

Range: 64 through 1,048,576 kilobytes

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

lifetime-seconds (Security Group VPN)

Syntax lifetime-seconds *seconds*;

Hierarchy Level [edit security group-vpn member ike proposal *proposal-name*]
[edit security group-vpn server ipsec proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify the lifetime (in seconds) of an IKE or IPsec security association (SA) for group VPN. When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

Options ***seconds***—Lifetime of the IKE or IPsec SA.

Range: 180 through 86,400 seconds

Default: 3600 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

lifetime-seconds (Security IKE)

Syntax lifetime-seconds *seconds*;

Hierarchy Level [edit security ike proposal *proposal-name*]

Release Information Statement introduced in Junos OS Release 8.5. Default value modified in Junos OS Release 10.2.

Description Specify the lifetime (in seconds) of an IKE security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

Options ***seconds***—Lifetime of the IKE SA.
Range: 180 through 86,400 seconds
Default: 28,800 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041
- *Understanding User Authentication Methods*

lifetime-seconds (Security IPsec)

Syntax lifetime-seconds ***seconds*** ;

Hierarchy Level [edit security ipsec proposal ***proposal-name***]

Release Information Statement introduced in Junos OS Release 8.5. Default value modified in Junos OS Release 10.2.

Description Specify the lifetime (in seconds) of an IPsec security association (SA). When the SA expires, it is replaced by a new SA and security parameter index (SPI) or terminated.

Options ***seconds***—Lifetime of the IPsec SA.
Range: 180 through 86,400 seconds
Default: 3600 seconds

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

load-distribution

Syntax
```
load distribution {
   session-affinity ipsec;
}
```

Hierarchy Level [edit security flow]

Release Information Statement introduced in Junos OS Release 11.4R5.

Description Enable load distribution for a data flow.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

local (Security IPsec)

Syntax local *ip-prefix*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike proxy-identity]

Release Information Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description Specify the local IPv4 or IPv6 address and subnet mask for the proxy identity.

Options ***ip-prefix*** —IPv4 or IPv6 address and subnet mask.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

local-address

Syntax local-address (*ipv4-address* | *ipv6-address*);

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Specify the local gateway address. Multiple addresses in the same address family can be configured on an external physical interface to a VPN peer. If this is the case, we recommend that **local-address** be configured. If there is only one IPv4 and one IPv6 address configured on an external physical interface, **local-address** configuration is not necessary.

NOTE: **local-address** must be an IP address that is configured on an interface on the SRX Series device. We recommend that **local-address** belong to the external interface of the IKE gateway. If **local-address** does not belong to the external interface of the IKE gateway, the interface must be in the same zone as the external interface of the IKE gateway and an intra-zone security policy must be configured to permit traffic.

local-address and the remote IKE gateway address must be in the same address family, either IPv4 or IPv6.

Options ***ipv4-address***—IPv4 address for the local gateway.

ipv6-address—IPv6 address for the local gateway.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

local-address (Security Group VPN Member)

Syntax local–address *ip-address*;

Hierarchy Level [edit security group-vpn member ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the IP address the member uses when accessing the group server.

Options *ip-address*—IPv4 address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

local-address (Security Group VPN Server)

Syntax local-address *ip-address*;

Hierarchy Level [edit security group-vpn server ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Configure the source IP address the group VPN server uses when communicating with a group member or a root-server. This statement is normally used when there are multiple IP addresses bound to an interface.

Options **local-address *ip-address***—Specify an IPv4 address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

local-certificate (Security)

Syntax local-certificate *certificate-id*;

Hierarchy Level [edit security ike policy *policy-name* certificate]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify a particular certificate when the local device has multiple loaded certificates.

NOTE: The device deletes existing IKE and IPsec SAs when you update the local-certificate configuration in the IKE policy.

Options ***certificate-id*** —Name of the specific certificate to be used.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

local-identity

Syntax

```
local-identity {
  (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
    e-mail-address);
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5. The **inet6** option added in Junos OS Release 11.1.

Description Specify the local IKE identity to send in the exchange with the destination peer to establish communication. If you do not configure a local-identity, the device uses the IPv4 or IPv6 address corresponding to the local endpoint by default.

NOTE: For Network Address Translation Traversal (NAT-T), both local identity and remote identity must be configured.

Options

- **hostname *hostname***—Specify identity as a fully qualified domain name (FQDN).
- **inet *ip-address***—Specify identity as an IPv4 address.
- **user-at-hostname *e-mail-address***—Specify identity as an e-mail address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

local-identity (Security Group VPN)

Syntax

```
local-identity {
    (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
      e-mail-address);
}
```

Hierarchy Level

[edit security group-vpn member ike gateway *gateway-name*]
[edit security group-vpn server ike gateway *gateway-name*]

Release Information

Support for **group-vpn** hierarchies added in Junos OS Release 10.2.

Description

Specify the local IKE identity to send in the exchange with the destination peer to establish communication. If you do not configure a local-identity, the device uses the IPv4 corresponding to the local endpoint by default.

Options

- **hostname *hostname***—Specify identity as a fully qualified domain name (FQDN).
- **inet *ip-address***—Specify identity as an IPv4 address.
- **user-at-hostname *e-mail-address***—Specify identity as an e-mail address.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

manual (Security IPsec)

Syntax

```
manual {
  authentication {
    algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
    key (ascii-text key | hexadecimal key );
  }
  encryption {
    algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
    key (ascii-text key | hexadecimal key );
  }
  external-interface external-interface-name;
  gateway ip-address;
  protocol (ah | esp);
  spi spi-value ;
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description Define a manual IPsec security association (SA).

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

member (Security Group VPN)

Syntax

```
member {
  ike {
    gateway gateway-name {
      ike-policy policy-name;
      local address ip-address;
      local-identity {
        (hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
          e-mail-address);
      }
      remote-identity {
        (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
      }
      routing-instance routing-instance;
      server-address [ip-address];
    }
    policy policy-name {
      description description;
      mode (aggressive | main);
      pre-shared-key (ascii-text key | hexadecimal key);
      proposals [proposal-name];
    }
    proposal proposal-name {
      authentication-algorithm (sha-256 | sha-384);
      authentication-method pre-shared-keys;
      description description;
      dh-group (group14 | group24);
      encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
      lifetime-seconds seconds;
    }
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      gateway-filter {
        local-address ip-address;
        remote-address ip-address;
      }
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
  ipsec {
    vpn vpn-name {
      df-bit (clear | copy | set);
      exclude rule rule-name {
        source-address ip-address/mask;
        destination-address ip-address/mask;
```


Copyright © 2016, Juniper Networks, Inc.

```
        application application;
      }
      fail-open rule rule-name {
        source-address ip-address/mask;
        destination-address ip-address/mask;
        application application;
      }
      group id;
      group-vpn-external-interface interface;
      ike-gateway gateway-name;
      recovery-probe;
    }
  }
}
```

Hierarchy Level [edit security group-vpn]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure group VPN member.

Options Configure group VPN member. You configure the following on the group member:

- Phase 1 IKE SA with the group server. The IKE gateway is the group server.

  NOTE: We recommend that you do not change the default value for **lifetime-seconds** for the IKE proposal on the member. Increasing the value might cause the member device to continue to use the existing Phase 1 IKE SA key even in the event of a crash; this can delay the recovery process.

- IPsec group VPN.

  NOTE: A scope policy must also be configured on the group member. To configure a scope policy, use the **policies** configuration statement at the [**edit security**] hierarchy and specify the IPsec group VPN for the **ipsec-group-vpn** option.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

member-threshold (Security Group VPN)

Syntax member-threshold *number*;

Hierarchy Level [edit security group-vpn server group *group-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Specify the maximum number of group VPN members that can be accepted in the group. The same **member-threshold** value must be configured on the root-server and all sub-servers in a group server cluster.

Options **member-threshold *number***—Specify the maximum number of group VPN members that can be accepted in the group. There is no default number.

Range: 1 to 2,000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699

Copyright © 2016, Juniper Networks, Inc.

mode (Security Group VPN)

Syntax mode (aggressive | main);

Hierarchy Level [edit security group-vpn member ike policy *policy-name*]
[edit security group-vpn server ike policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2.

Description Define the mode used for Internet Key Exchange (IKE) Phase 1 negotiations. Use aggressive mode only when you need to initiate an IKE key exchange without ID protection, as when a peer unit has a dynamically assigned IP address. (The **main** option is not supported on dynamic VPN implementations.)

NOTE:

- IKEv2 protocol does not negotiate using mode configuration.
- The device deletes existing IKE and IPsec SAs when you update the mode configuration in the IKE policy.

Options

- **aggressive**—Aggressive mode.
- **main**—Main mode. Main mode is the recommended key-exchange method because it conceals the identities of the parties during the key exchange.

NOTE: Configuring mode main for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is pre-shared-keys.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

mode (Security IKE Policy)

Syntax mode (aggressive | main);

Hierarchy Level [edit security ike policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the mode used for Internet Key Exchange (IKE) Phase 1 negotiations. Use aggressive mode only when you need to initiate an IKE key exchange without ID protection, as when a peer unit has a dynamically assigned IP address.

NOTE:

- IKEv2 protocol does not negotiate using mode configuration.
- The device deletes existing IKE and IPsec SAs when you update the mode configuration in the IKE policy.

Options

- **aggressive**—Aggressive mode.
- **main**—Main mode. Main mode is the recommended key-exchange method because it conceals the identities of the parties during the key exchange.

NOTE: Configuring mode main for group VPN servers or members is not supported when the remote gateway has a dynamic address and the authentication method is pre-shared-keys.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

nat-keepalive

Syntax nat-keepalive *seconds*;

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5. Default value changed from 5 seconds to 20 seconds in Junos OS Release 12.1X46-D10.

Description Specify the interval at which NAT keepalive packets can be sent so that NAT translation continues.

Options ***seconds*** —Maximum interval in seconds at which NAT keepalive packets can be sent.
Range: 1 through 300 seconds.
Default: 20 seconds.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

no-anti-replay (Security)

Syntax no-anti-replay;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Disable the antireplay checking feature of IPsec. By default, antireplay checking is enabled.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

no-nat-traversal

Syntax no-nat-traversal;

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Disables UDP encapsulation of IPsec Encapsulating Security Payload (ESP) packets, otherwise known as Network Address Translation Traversal (NAT-T). NAT-T is enabled by default.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

non-cryptographic-self-test

Syntax non-cryptographic-self-test;

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm when the device or switch detects a noncryptographic self-test failure. The self-tests run without operator intervention.

Default No alarm is raised upon failure of a noncryptographic self-test.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

ocsp (Security PKI)

Syntax

```
ocsp {
  connection-failure (disable | fallback-crl);
  disable-responder-revocation-check;
  nonce-payload (enable | disable);
  url ocsp-url;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name* revocation-check]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure Online Certificate Status Protocol (OCSP) to check the revocation status of a certificate.

Options

connection-failure—(Optional) Specify action to take if there is a connection failure to the OCSP responder. If this option is not configured and there is no response from the OCSP responder, certificate validation will fail.

disable—Skip the revocation check if the OCSP responder is not reachable.

fallback-crl—Use CRL to check the revocation status of the certificate.

disable-responder-revocation-check —(Optional) Disable revocation check for the CA certificate received in an OCSP response. The certificates received in an OCSP response generally have shorter lifetimes and revocation check is not required.

nonce-payload—(Optional) Send a nonce payload to prevent replay attack. A nonce payload is sent by default unless it is explicitly disabled. If enabled, the SRX Series device expects OCSP responses to contain a nonce payload, otherwise the revocation check will fail. If OCSP responders are not capable of responding with a nonce payload, disable this option.

disable—Explicitly disable the sending of a nonce payload.

enable—Enable the sending of a nonce payload. This is the default.

url *ocsp-url*—Specify HTTP addresses for OCSP responders. A maximum of two HTTP URL addresses can be configured. If the configured URLs are not reachable, or URLs are not configured, the URL from the certificate being verified is checked.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

optimized

Syntax optimized;

Hierarchy Level [edit security ipsec vpn *vpn-name* vpn-monitor]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that VPN monitoring optimization is enabled for the VPN object. When VPN monitoring optimization is enabled, the SRX Series device only sends ICMP echo requests (pings) when there is outgoing traffic and no incoming traffic from the configured peer through the VPN tunnel. If there is incoming traffic through the VPN tunnel, the SRX Series device considers the tunnel to be active and does not send pings to the peer.

Because ICMP echo requests are only sent when needed to determine peer liveliness, VPN monitoring optimization can save resources on the SRX Series device. Also, ICMP echo requests can activate costly backup links that would otherwise not be used.

This option is disabled by default.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

optimized (DPD)

Syntax optimized;

Hierarchy Level [edit security ike gateway *gateway-name* dead-peer-detection]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Send dead peer detection (DPD) messages if there is no incoming IKE or IPsec traffic within the configured interval after outgoing packets are sent to the peer. This is the default DPD mode.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

peer-certificate-type

Syntax peer-certificate-type (pkcs7 | x509-signature);

Hierarchy Level [edit security ike policy *policy-name* certificate]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify a preferred type of certificate (PKCS7 or X509).

Options
- **pkcs7**—Public-Key Cryptography Standard #7.
- **x509-signature**—X509 is an ITU-T standard for public key infrastructure. This is the default value.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

perfect-forward-secrecy (Security IPsec)

Syntax perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 | group5);

Hierarchy Level [edit security ipsec policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **group14** options added in Junos OS Release 11.1. Support for **group19**, **group20**, and **group24** options added in Junos OS Release 12.1X45-D10.

Description Specify Perfect Forward Secrecy (PFS) as the method that the device uses to generate the encryption key. PFS generates each new encryption key independently from the previous key.

NOTE: The device deletes existing IPsec SAs when you update the **perfect-forward-secrecy** configuration in the IPsec policy.

Options

- **group1**—Diffie-Hellman Group 1.
- **group14**—Diffie-Hellman Group 14.
- **group19**—Diffie-Hellman Group 19.
- **group2**—Diffie-Hellman Group 2.
- **group20**—Diffie-Hellman Group 20.
- **group24**—Diffie-Hellman Group 24.
- **group5**—Diffie-Hellman Group 5.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

pki

Syntax

```
pki {
  auto-re-enrollment {
    certificate-id certificate-id-name {
      ca-profile-name ca-profile-name ;
      challenge-password password ;
      re-enroll-trigger-time-percentage percentage ;
      re-generate-keypair;
    }
  }
  ca-profile ca-profile-name {
    administrator {
      e-mail-address e-mail-address;
    }
    ca-identity ca-identity;
    enrollment {
        retry number;
        retry-interval seconds;
        url url-name;
    }
    revocation-check {
      crl {
        disable {
          on-download-failure;
        }
        refresh-interval hours;
        url url-name;
      }
      disable;
      ocsp {
        connection-failure (disable | fallback-crl);
        disable-responder-revocation-check;
        nonce-payload (enable | disable);
        url ocsp-url;
      }
      use-ocsp;
    }
    routing-instance routing-instance-name;
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security]

Copyright © 2016, Juniper Networks, Inc.

Release Information Statement modified in Junos OS Release 8.5.

Description Configure an IPsec profile to request digital certificates.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

pki-local-certificate

Syntax pki-local-certificate *name*;

Hierarchy Level [edit system services web-management https]

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify the name of the certificate that is generated by public key infrastructure (PKI) and authenticated by certificate authority (CA).

Options *name* —Name of certificate.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

policy (Security Group VPN IKE)

Syntax

```
policy policy-name {
  description description;
  mode (aggressive | main);
  pre-shared-key (ascii-text key | hexadecimal key);
  proposals [proposal-name];
}
```

Hierarchy Level

[edit security group-vpn member ike]
[edit security group-vpn server ike]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Configure an IKE policy.

Options

policy-name—Name of the IKE policy. The policy name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

policy (Security IKE)

Syntax

```
policy policy-name {
  certificate {
    local-certificate certificate-id;
    peer-certificate-type (pkcs7 | x509-signature);
    policy-oids [ oid ];
  }
  description description;
  mode (aggressive | main);
  pre-shared-key (ascii-text key | hexadecimal key);
  proposal-set (basic | compatible | standard } suiteb-gcm-128 | suiteb-gcm-256);
  proposals [proposal-name];
}
```

Hierarchy Level [edit security ike]

Release Information Statement modified in Junos OS Release 8.5. Support for **suiteb-gcm-128** and **suiteb-gcm-256** options added in Junos OS Release 12.1X45-D10. Support for **policy-oids** option added in Junos OS Release 12.3X48-D10.

Description Configure an IKE policy.

Options *policy-name*—Name of the IKE policy. The policy name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

policy (Security IPsec)

Syntax

```
policy policy-name {
  description description;
  perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 |
    group5);
  proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
  proposals [proposal-name];
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement modified in Junos OS Release 8.5. Support for group 14 is added in Junos OS Release 11.1.

Description Define an IPsec policy.

Options ***policy-name*** —Name of the IPsec policy.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

policy-oids

Syntax policy-oids [*oid*];

Hierarchy Level [edit security ike policy *policy-name* certificate]

Release Information Statement introduced in Junos OS Release 12.3X48-D10.

Description Configure policy object identifiers (OIDs). This configuration is optional.

Options ***oid***—Policy OID contained in a peer's certificate or certificate chain. Up to five policy OIDs can be configured. Each OID can be up to 63 bytes long.

NOTE: You must ensure that at least one of the configured policy OIDs is included in a peer's certificate or certificate chain. Note that the **policy-oids** field in a peer's certificate is optional. If you configure policy OIDs in an IKE policy and the peer's certificate chain does not contain any policy OIDs, certificate validation for the peer fails.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Understanding Digital Certificate Validation on page 6357

 Copyright © 2016, Juniper Networks, Inc.

pre-shared-key (Security IKE Policy)

Syntax pre-shared-key (ascii-text *key* | hexadecimal *key*);

Hierarchy Level [edit security ike policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Define a preshared key for an IKE policy.

NOTE: The device deletes existing IKE and IPsec SAs when you update the pre-shared-key configuration in the IKE policy.

Options

ascii-text *key*—Specify a string of 1 to 255 ASCII text characters for the key. Characters **@ + -** or **=** are not allowed. To include the special characters **( ) [ ] { } , ;** enclose either the entire key string or the special character in quotation marks; for example **"str)ng"** or **str")"ng**. Other use of quotation marks within the string is not allowed. With **des-cbc** encryption, the key contains 8 ASCII characters. With **3des-cbc** encryption, the key contains 24 ASCII characters.

hexadecimal *key*—Specify a string of 1 to 255 hexadecimal characters for the key. Characters must be hexadecimal digits **0** through **9**, or letters **a** through **f** or **A** through **F**. With **des-cbc** encryption, the key contains 16 hexadecimal characters. With **3des-cbc** encryption, the key contains 48 hexadecimal characters.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

probe-idle-tunnel

Syntax probe-idle-tunnel;

Hierarchy Level [edit security ike gateway *gateway-name* dead-peer-detection]

Release Information Statement introduced in Junos OS Release 12.1X46-D10.

Description Send dead peer detection (DPD) messages during idle traffic time between peers.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order [accounting-method];
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order [accounting-method];
  address-assignment pool pool-name;
  authentication-order [ldap | none | password | radius | securid];
  authorization-order [jsrc];
  client client-name {
    chap-secret chap-secret;
    client-group [ group-names ];
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
  provisioning-order (gx-plus | jsrc);
```


Copyright © 2016, Juniper Networks, Inc.

```
radius {
  accounting-server [server];
  attributes {
    exclude {
      acc-aggr-cir-id-asc [access-request | accounting-start | accounting-stop];
      acc-aggr-cir-id-bin [access-request | accounting-start | accounting-stop];
      acc-loop-cir-id [access-request | accounting-start | accounting-stop];
      accounting-authentic [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-delay-time [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      accounting-session-id [access-request];
      accounting-terminate-cause [accounting-off];
      act-data-rate-dn [access-request | accounting-start | accounting-stop];
      act-data-rate-up [access-request | accounting-start | accounting-stop];
      act-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      act-interlv-delay-up [access-request | accounting-start | accounting-stop];
      att-data-rate-dn [access-request | accounting-start | accounting-stop];
      att-data-rate-up [access-request | accounting-start | accounting-stop];
      called-station-id [access-request | accounting-start | accounting-stop];
      calling-station-id [access-request | accounting-start | accounting-stop];
      class [access-request | accounting-start | accounting-stop];
      delegated-ipv6-prefix [accounting-start | accounting-stop];
      dhcp-gi-address [access-request | accounting-start | accounting-stop];
      dhcp-mac-address [access-request | accounting-start | accounting-stop];
      dhcp-options [access-request | accounting-start | accounting-stop];
      downstream-calculated-qos-rate [access-request | accounting-start |
        accounting-stop];
      dsl-forum-attributes [access-request | accounting-start | accounting-stop];
      dsl-line-state [access-request | accounting-start | accounting-stop];
      dsl-type [access-request | accounting-start | accounting-stop];
      dynamic-iflset-name [accounting-start | accounting-stop];
      event-time-stamp [accounting-off | accounting-on | accounting-start |
        accounting-stop];
      framed-interface-id [access-request | accounting-start | accounting-stop];
      framed-ip-address [access-request | accounting-start | accounting-stop];
      framed-ip-netmask [access-request | accounting-start | accounting-stop];
      framed-ip-route [access-request | accounting-start | accounting-stop];
      framed-ipv6-pool [accounting-start | accounting-stop];
      framed-ipv6-prefix [accounting-start | accounting-stop];
      framed-ipv6-route [accounting-start | accounting-stop];
      framed-pool [accounting-start | accounting-stop];
      input-filter [accounting-start | accounting-stop];
      input-gigapackets [accounting-stop];
      input-gigawords [accounting-stop];
      input-ipv6-gigawords [accounting-stop];
      input-ipv6-octets [accounting-stop];
      input-ipv6-packets [accounting-stop];
      interface-description [access-request | accounting-start | accounting-stop];
      l2c-downstream-data [access-request | accounting-start | accounting-stop];
      l2c-upstream-data [access-request | accounting-start | accounting-stop];
      max-data-rate-dn [access-request | accounting-start | accounting-stop];
      max-data-rate-up [access-request | accounting-start | accounting-stop];
      max-interlv-delay-dn [access-request | accounting-start | accounting-stop];
      max-interlv-delay-up [access-request | accounting-start | accounting-stop];
      min-data-rate-dn [access-request | accounting-start | accounting-stop];
```

Copyright © 2016, Juniper Networks, Inc.

```
            min-data-rate-up [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-dn [access-request | accounting-start | accounting-stop];
            min-lp-data-rate-up [access-request | accounting-start | accounting-stop];
            nas-identifier [access-request | accounting-start | accounting-stop];
            nas-port [access-request | accounting-off | accounting-on | accounting-start |
              accounting-stop];
            nas-port-id [access-request | accounting-start | accounting-stop];
            nas-port-type [access-request | accounting-start | accounting-stop];
            output-filter [accounting-start | accounting-stop];
            output-gigapackets [accounting-stop];
            output-gigawords [accounting-stop];
            output-ipv6-gigawords [accounting-stop];
            output-ipv6-octets [accounting-stop];
            output-ipv6-packets [accounting-stop];
            upstream-calculated-qos-rate [access-request | accounting-start | accounting-stop];
          }
          ignore {
            dynamic-iflset-name;
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
          }
        }
      authentication-server [server];
      radius-options {
        request-rate number;
        revert-interval seconds;
      }
      radius-server server-address {
        accounting-port port-number
        max-outstanding-requests number-of--outstanding-requests;
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        secret password;
        source-address source-address;
        timeout seconds;
      }
      service {
        accounting-order {
          activation-protocol;
          radius;
        }
      }
      session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
      }
    }
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Copyright © 2016, Juniper Networks, Inc.

Description Create a profile containing a set of attributes that define device management access.

Required Privilege Level
access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation
- Understanding Interfaces on page 2441
- Understanding User Authentication for Security Devices on page 5173
- Layer 2 Transparent and Switching Overview on page 2893

proposal (Security Group VPN Member IKE)

Syntax

```
proposal proposal-name {
  authentication-algorithm (sha-256 | sha-384);
  authentication-method pre-shared-keys;
  description description;
  dh-group (group14 | group24);
  encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
  lifetime-seconds seconds;
}
```

Hierarchy Level [edit security group-vpn member ike]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define an IKE proposal.

Options ***proposal-name***—Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

proposal (Security Group VPN Server IKE)

Syntax

```
proposal proposal-name {
  authentication-algorithm (sha-256 | sha-384);
  authentication-method pre-shared-keys;
  description description;
  dh-group (group14 | group24);
  encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
}
```

Hierarchy Level [edit security group-vpn server ike]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define an IKE proposal for group VPN server.

Options ***proposal-name***—Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

proposal (Security Group VPN Server IPsec)

Syntax

```
proposal proposal-name {
  authentication-algorithm hmac-sha-256-128;
  description description;
  encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
  lifetime-seconds seconds;
}
```

Hierarchy Level [edit security group-vpn server ipsec]

Release Information Statement introduced in Junos OS Release 10.2.

Description Define an IPsec proposal.

Options *proposal-name*—Name of the IPsec proposal.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

 Copyright © 2016, Juniper Networks, Inc.

proposal (Security IKE)

Syntax

```
proposal proposal-name {
  authentication-algorithm (md5 | sha-256 | sha-384| sha1);
  authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384
    | pre-shared-keys | rsa-signatures);
  description description;
  dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
  encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
  lifetime-seconds seconds;
}
```

Hierarchy Level [edit security ike]

Release Information Statement modified in Junos OS Release 8.5. Support for **dh-group group 14** and **dsa-signatures** added in Junos OS Release 11.1. Support for **sha-384**, e**cdsa-signatures-256**, **ecdsa-signatures-384**, **group19**, **group20**, and **group24** options added in Junos OS Release 12.1X45-D10.

Description Define an IKE proposal.

Options ***proposal-name***—Name of the IKE proposal. The proposal name can be up to 32 alphanumeric characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

proposal (Security IPsec)

Syntax

```
proposal proposal-name {
  authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
  description description;
  encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm
    | aes-256-cbc | aes-256-gcm | des-cbc);
  lifetime-kilobytes kilobytes;
  lifetime-seconds seconds;
  protocol (ah | esp);
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement modified in Junos OS Release 8.5.

Description Define an IPsec proposal.

Options *proposal-name*—Name of the IPsec proposal.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

proposals (Security Group VPN)

Syntax proposals [*proposal-name*];

Hierarchy Level [edit security group-vpn member ike policy *policy-name*]
[edit security group-vpn server ike policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2.

Description Specify up to four Phase 1 proposals for an IKE policy. If you include multiple proposals, use the same Diffie-Hellman group in all of the proposals.

Options ***proposal-name***—Names of up to four configured Phase 1 proposals.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

proposals (Security IKE)

Syntax proposals [*proposal-name*];

Hierarchy Level [edit security ike policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **group-vpn** hierarchies added in Junos OS Release 10.2.

Description Specify up to four Phase 1 proposals for an IKE policy. If you include multiple proposals, use the same Diffie-Hellman group in all of the proposals.

Options ***proposal-name***—Names of up to four configured Phase 1 proposals.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

proposals (Security IPsec)

Syntax proposals [*proposal-name*];

Hierarchy Level [edit security ipsec policy *policy-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Specify one or more proposals for an IPsec policy.

Options *proposal-name*—Name of a configured proposal.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

proposal-set (Security IKE)

Syntax proposal-set (basic | compatible | prime-128 | prime-256 | standard | suiteb-gcm-128 | suiteb-gcm-256);

Hierarchy Level [edit security ike policy *policy-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **suiteb-gcm-128** and **suiteb-gcm-256** options added in Junos OS Release 12.1X45-D10. Support for **prime-128** and **prime-256** options added in Junos OS Release 15.1X49-D40.

Description Specify a set of default Internet Key Exchange (IKE) proposals.

NOTE: The **prime-128** and **prime-256** proposal sets require IKEv2 and certificate-based authentication.

Options

- **basic**—Includes a basic set of two IKE proposals:
  - Proposal 1—Preshared key, Data Encryption Standard (DES) encryption, and Diffie-Hellman (DH) group 1 and Secure Hash Algorithm 1 (SHA-1) authentication.
  - Proposal 2—Preshared key, DES encryption, and DH group 1 and Message Digest 5 (MD5) authentication.
- **compatible**—Includes a set of four commonly used IKE proposals:
  - Proposal 1—Preshared key, triple DES (3DES) encryption, and Gnutella2 (G2) and SHA-1 authentication.
  - Proposal 2—Preshared key, 3DES encryption, and DH group 2 and MD5 authentication.
  - Proposal 3—Preshared key, DES encryption, and DH group 2 and SHA-1 authentication.
  - Proposal 4—Preshared key, DES encryption, and DH group 2 and MD5 authentication.
- **prime-128**—Provides the following proposal set (this option is not supported on Group VPNv2):
  - Authentication method—Elliptic Curve Digital Signature Algorithm (ECDSA) 256-bit signatures.
  - Diffie-Hellman Group—19.
  - Encryption algorithm—Advanced Encryption Standard (AES) 128-bit Galois/Counter Mode (GCM).
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication).

  When this option is used, **prime-128** should also be configured at the [**edit security ipsec policy *policy-name* proposal-set**] hierarchy level.

Copyright © 2016, Juniper Networks, Inc.

- **prime-256**—Provides the following proposal set (this option is not supported on Group VPNv2):
  - Authentication method—ECDSA 384-bit signatures.
  - Diffie-Hellman Group—20.
  - Encryption algorithm—AES 256-bit GCM.
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication).

  When this option is used, **prime-256** should also be configured at the [**edit security ipsec policy *policy-name* proposal-set**] hierarchy level.
- **standard**—Includes a standard set of two IKE proposals:
  - Proposal 1— Preshared key, 3DES encryption, and DH group 2 and SHA-1 authentication.
  - Proposal 2—Preshared key, AES 128-bit encryption, and DH group 2 and SHA-1 authentication.
- **suiteb-gcm-128**—Provides the following Suite B proposal set (this option is not supported on Group VPNv2):
  - Authentication method—ECDSA 256-bit signatures
  - Diffie-Hellman Group—19
  - Encryption algorithm—Advanced Encryption Standard (AES) 128-bit cipher block chaining (CBC)

    NOTE: CBC mode is used instead of GCM.
  - Authentication algorithm—SHA-256
- **suiteb-gcm-256**—Provides the following Suite B proposal set (this option is not supported on Group VPNv2):
  - Authentication method—ECDSA 384-bit signatures
  - Diffie-Hellman Group—20
  - Encryption algorithm—AES 256-bit CBC

    NOTE: CBC mode is used instead of GCM.
  - Authentication algorithm—SHA-384

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

proposal-set (Security IPsec)

Syntax proposal-set (basic | compatible | prime-128 | prime-256 | standard | suiteb-gcm-128 | suiteb-gcm-256);

Hierarchy Level [edit security ipsec policy *policy-name*]

Release Information Statement introduced in Junos OS Release 10.4. Support for **suiteb-gcm-128** and **suiteb-gcm-256** options added in Junos OS Release 12.1X45-D10. Support for **prime-128** and **prime-256** options added in Junos OS Release 15.1X49-D40.

Description Define a set of default IPsec proposals.

Options

- **basic** —nopfs-esp-des-sha and nopfs-esp-des-md5
- **compatible**—nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5
- **prime-128**—Provides the following proposal set:
  - Encapsulating Security Payload (ESP) protocol
  - Encryption algorithm—Advanced Encryption Standard Galois/Counter mode (AES-GCM)128-bit
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

  NOTE: This option is not supported on Group VPNv2.

- **prime-256**—Provides the following proposal set:
  - ESP protocol
  - Encryption algorithm—AES-GCM 256-bit
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

  NOTE: This option is not supported on Group VPNv2.

- **standard**—g2-esp-3des-sha and g2-esp-aes128-sha
- **suiteb-gcm-128**—Provides the following proposal set:
  - ESP protocol
  - Encryption algorithm—AES-GCM 128-bit

Copyright © 2016, Juniper Networks, Inc.

- Authentication algorithm—None (AES-GCM provides both encryption and authentication)

NOTE: This option is not supported on Group VPNv2.

- **suiteb-gcm-256**—Provides the following proposal set:
  - ESP protocol
  - Encryption algorithm—AES-GCM 256-bit
  - Authentication algorithm—None (AES-GCM provides both encryption and authentication)

NOTE: This option is not supported on Group VPNv2.

NOTE: The Perfect Forward Secrecy setting in IPsec policy overrides the settings in proposal sets.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

protocol (IPsec SA for OSPF)

Syntax protocol (ah | esp);

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure the IPsec protocol for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **protocol**—Define the IPsec protocol for the manual SA. The protocol can be one of the following:

- **ah**—Authentication Header (AH) protocol.
- **esp**—Encapsulating Security Payload (ESP) protocol. This is the default.

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

protocol (Security IPsec)

Syntax protocol (ah | esp);

Hierarchy Level [edit security ipsec proposal *proposal-name*]

Release Information Statement modified in Junos OS Release 8.5.

Description Define the IPsec protocol for a manual or dynamic security association (SA).

NOTE: The device deletes existing IPsec SAs when you update the encryption-algorithm configuration in the IPsec proposal.

Options

- **ah**—Authentication Header protocol.
- **esp**—Encapsulating Security Payload (ESP) protocol.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

protocol (Security IPsec Manual SA)

Syntax protocol (ah | esp)

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement modified in Junos OS Release 8.5.

Description Define the IPsec protocol for the manual security association.

Options

- **ah**—Authentication Header protocol.
- **esp**—ESP protocol (To use the ESP protocol, you must also use the **tunnel** statement at the **[edit security ipsec security-association *sa-name* mode]** hierarchy level.)

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

proxy-identity

Syntax

```
proxy-identity {
  local ip-prefix;
  remote ip-prefix;
  service (all | service-name);
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name* ike]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 added in Junos OS Release 12.1X46-D10.

Description Optionally specify the IPsec proxy ID to use in negotiations. The default is the identity based on the IKE gateway. If the IKE gateway is an IPv6 site-to-site gateway, the default proxy ID is ::/0. If the IKE gateway is an IPv4 gateway or a dynamic endpoint or dialup gateway, the default proxy ID is 0.0.0.0/0.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

re-enroll-trigger-time-percentage (Security PKI)

Syntax re-enroll-trigger-time-percentage *percentage*;

Hierarchy Level [edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]
[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]

Release Information Statement modified in Junos OS Release 9.0. Support for **[edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]** and **[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]** hierarchies added in Junos OS Release 15.1X49-D40.

Description Specify the certificate reenrollment trigger as a percentage of the end-entity (EE) certificate's lifetime that remains before certificate reenrollment is initiated. For example, if the renewal request is to be sent when the certificate's remaining lifetime is 10 percent, then configure 10 for **re-enroll-trigger-time-percentage** value. The time at which the certificate reenrollment is initiated is based on the certificate expiry date.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

re-generate-keypair

Syntax re-generate-keypair;

Hierarchy Level [edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]
[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]

Release Information Statement introduced in Junos OS Release 8.5. Support for **[edit security pki auto-re-enrollment cmpv2 certificate-id *certificate-id-name*]** and **[edit security pki auto-re-enrollment scep certificate-id *certificate-id-name*]** hierarchies added in Junos OS Release 15.1X49-D40.

Description Specifies new key pair generation for automatic certificate reenrollment. If this statement is not configured, the current key pair is used. If the key pair does not change, the CA may not issue new certificates. We recommend that a new key pair be generated during reenrollment as it provides better security.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

refresh-interval

Syntax refresh-interval *hours*;

Hierarchy Level [edit security pki ca-profile *ca-profile-name* revocation-check crl]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the amount of time between certificate revocation list (CRL) updates.

Options ***number-of-hours***—Time interval, in hours, between CRL updates.
Range: 0 through 8784
Default: 6

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Security Configuration Statement Hierarchy on page 603
- [edit security pki] Hierarchy Level on page 646
- **crl (Security) on page 6852**

remote (Security IPsec)

Syntax remote *ip-prefix*;

Hierarchy Level [edit security ipsec vpn *vpn-name* ike proxy-identity]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1.

Description Specify the remote IPv4 or IPv6 address and subnet mask for the proxy identity.

Options ***ip-prefix*** —IPv4 or IPv6 address and subnet mask.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

remote-identity

Syntax

```
remote-identity {
    (distinguished-name <container container-string> <wildcard wildcard-string> | hostname
      hostname | inet ip-address | inet6 ipv6-address | user-at-hostname e-mail-address);
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the remote IKE identity to exchange with the destination peer to establish communication. If you do not configure a remote-identity, the device uses the IPv4 or IPv6 address corresponding to the remote endpoint by default.

NOTE: For Network Address Translation Traversal (NAT-T), both remote identity and local identity must be configured.

Options

- **distinguished-name**—Specify identity as the distinguished name (DN) from the certificate. If there is more than one certificate on the device, use the **security ike gateway** *gateway-name* **policy** *policy-name* **certificate local-certificate** *certificate-id*.

  Optional container and wildcard strings may be specified:

  - **container** ***container-string***—Specify a string for the container.
  - **wildcard** ***wildcard-string***—Specify a string for the wildcard.

- **hostname** ***hostname***—Specify identity as a fully qualified domain name (FQDN).
- **inet** ***ip-address***—Specify identity as an IPv4 address.
- **inet6** ***ipv6-address***—Specify identity as an IPv6 address.
- **user-at-hostname** ***e-mail-address***—Specify identity as an e-mail address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

remote-identity (Security Group VPN)

Syntax

```
remote-identity {
    (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
}
```

Hierarchy Level

[edit security group-vpn member ike gateway *gateway-name*]
[edit security group-vpn server ike gateway *gateway-name*]

Release Information

Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description

Specify the remote IKE identity of the destination peer. If you do not configure a remote identity, the device uses, by default, the IPv4 address that corresponds to the destination peer.

Options

hostname *hostname*—Specify a fully qualified domain name (FQDN).

inet *ip-address*—Specify an IPv4 address.

user-at-hostname *username_FQDN*—Specify a fully qualified username.

Required Privilege Level

security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

replay-attacks

Syntax
```
replay-attacks {
  threshold value;
}
```

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm when the device detects a replay attack. A replay attack is a form of network attack in which a valid data transmission is maliciously or fraudulently repeated or delayed.

Default Replay attacks do not raise security alarms.

Options
- **threshold *value***—Number of reply attacks up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

  Range: Range: 0 through 100,00,00,000.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

respond-bad-spi

Syntax respond-bad-spi *max-responses*;

Hierarchy Level [edit security ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Enable response to invalid IPsec Security Parameter Index (SPI) values. If the security associations (SAs) between two peers of an IPsec VPN become unsynchronized, the device resets the state of a peer so that the two peers are synchronized.

Options ***max-responses***—Number of times to respond to invalid SPI values per gateway.

Range: 1 through 30

Default: 5

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

revocation-check (Security PKI)

Syntax

```
revocation-check {
  crl {
    disable {
      on-download-failure;
    }
    refresh-interval hours;
    url url-name;
  }
  disable;
  ocsp {
    connection-failure (disable | fallback-crl);
    disable-responder-revocation-check;
    nonce-payload (enable | disable);
    url ocsp-url;
  }
  use-ocsp;
}
```

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement modified in Junos OS Release 8.5. Support for **ocsp** and **use-ocsp** options added in Junos OS Release 12.1X46-D20.

Description Specify the method the device uses to verify the revocation status of digital certificates.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

routing-instance (Security Group VPN)

Syntax routing-instance *routing-instance*;

Hierarchy Level [edit security group-vpn member ike gateway *gateway-name*]
[edit security group-vpn server ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Configure the routing instance that the group VPN server or member uses when communicating with a group member or server. This statement is used when the IKE gateway is not configured in the default routing instance.

Options **routing-instance *routing-instance***—Specify the name of a routing instance. If this is not specified, the default inet.0 routing instance is used.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

routing-instance (Security PKI)

Syntax routing-instance *routing-instance-name*

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement modified in Junos OS Release 9.0.

Description Specify the routing-instance to be used.

Options
- ***routing-instance-name***—Name of the routing instance.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

security-association

Syntax

```
security-association sa-name {
  manual {
    direction bidirectional {
      authentication {
        algorithm (hmac-md5-96 | hmac-sha1-96);
        key {
          ascii-text key;
          hexadecimal key;
        }
      }
      auxiliary-spi auxiliary-spi-value;
      encryption {
        algorithm (3des-cbc | des-cbc | null);
        key {
          ascii-text key;
          hexadecimal key;
        }
      }
      protocol (ah | esp);
      spi spi-value;
    }
  }
  mode transport;
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link. IPsec can provide authentication and confidentiality to OSPF or OSPFv3 routing packets.

Options

sa-name—Name of the SA.

mode—SA mode. For this feature, the mode must be **transport**.

direction—Direction of the manual SA. For this feature, the direction must be **bidirectional**.

The remaining statements are explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

server (Security Group VPN)

Syntax

```
server {
  group name {
    anti-replay-time-window milliseconds;
    description description;
    group-id number;
    ike-gateway [gateway-name];
    ipsec-sa name {
      match-policy policy-name {
        destination ip-address/netmask;
        destination-port number;
        protocol number;
        source ip-address/netmask;
        source-port number;
      }
      proposal proposal-name;
    }
    member-threshold number;
    server-cluster {
      ike-gateway gateway-name;
      retransmission-period seconds;
      server-role (root-server | sub-server);
    }
    server-member-communication {
      certificate certificate-id;
      communication-type unicast;
      encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
      lifetime-seconds seconds;
      number-of-retransmission number;
      retransmission-period seconds;
      sig-hash-algorithm (sha-256 | sha-384);
    }
  }
  ike {
    gateway gateway-name {
      address ip-address ;
      dead-peer-detection {
        always-send;
        interval seconds;
        threshold number;
      }
      dynamic {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
      }
      ike-policy policy-name;
      local-address ip-address;
      local-identity {
        (hostname hostname | inet ip-address | user-at-hostname e-mail-address);
      }
      remote-identity {
        (hostname [hostname] | inet ip-address | user-at-hostname e-mail-address);
      }
      routing-instance routing-instance;
```

```
    }
    policy policy-name {
      description text;
      mode (aggressive | main);
      pre-shared-key (ascii-text key | hexadecimal key);
      proposals [proposal-name];
    }
    proposal proposal-name {
      authentication-algorithm (sha-256 | sha-384);
      authentication-method pre-shared-keys;
      description description;
      dh-group (group14 | group24);
      encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
    }
  }
  ipsec {
    proposal proposal-name {
      authentication-algorithm hmac-sha-256-128;
      description description;
      encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
      lifetime-seconds seconds;
    }
  }
  traceoptions {
    file {
      filename;
      files number;
      match regular-expression;
      size maximum-file-size;
      (world-readable | no-world-readable);
    }
    flag flag;
    gateway-filter {
      local-address ip-address;
      remote-address ip-address;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level [edit security group-vpn]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure group VPN server. You configure the following on the group server:

- Phase 1 IKE SA for group members
- Phase 2 IPsec proposal
- Group identifier, group members, server-member communications, and group policies to be downloaded to members
- Group VPN trace options

Copyright © 2016, Juniper Networks, Inc.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

server-address (Security Group VPN Member)

Syntax server-address [*ip-address*];

Hierarchy Level [edit security group-vpn member ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Specify the group server that this member registers through a **groupkey-pull** exchange. Up to four server IP addresses can be configured. The group member attempts to register with the first configured server. If registration with a configured server is not successful, the group member tries to register with the next configured server.

NOTE: We recommend that group members only register with sub-servers in a server cluster and not the root-server.

Options ***ip-address***—IPv4 address of an IKE gateway.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

server-cluster (Security Group VPN Server)

Syntax

```
server-cluster {
  ike-gateway gateway-name;
  retransmission-period seconds;
  server-role (root-server | sub-server);
}
```

Hierarchy Level [edit security group-vpn server group *name*]

Release Information Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Configure the Group Domain of Interpretation (GDOI) group controller/key server (GCKS) cluster for the specified group. All servers in a group VPN server cluster must be SRX Series devices.

Options

ike-gateway *gateway-name*—(Required) Specify the name of the IKE gateway for the local device in the group server cluster. IKE gateways are configured at the [**edit security group-vpn server ike**] hierarchy level.

If the local device is a root-server, the IKE gateway name must be a sub-server in the cluster; up to four sub-server IKE gateways can be specified.

If the local device is a sub-server, the IKE gateway name must be the root-server.

retransmission-period *seconds*—(Optional) Specify the time after which the root-server retransmits a **cluster-update** message if it has not received an acknowledgement from a sub-server.

Range: 2 to 60 seconds.

Default: 10 seconds.

server-role—(Required) Assign the role of the local device in the group server cluster, either **root-server** or **sub-server**. Only one device in the cluster can be configured as the root-server. You can configure up to four other devices as a sub-server in a group server cluster.

NOTE: You must ensure that there is only one root-server at any time for a group VPN server cluster.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699

Copyright © 2016, Juniper Networks, Inc.

server-member-communication (Security Group VPN Server)

Syntax

```
server-member-communication {
  certificate certificate-id;
  communication-type unicast;
  encryption-algorithm (aes-128-cbc | aes-192-cbc | aes-256-cbc);
  lifetime-seconds seconds;
  number-of-retransmission number;
  retransmission-period seconds;
  sig-hash-algorithm (sha-256 | sha-384);
}
```

Hierarchy Level [edit security group-vpn server group *name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable and configure server to member communication. When these options are configured, group members receive new keys before current keys expire.

Options

- **certificate** ***certificate-id***—Specify the certificate identification. Only RSA keys are supported.
- **communication-type**—Configure **unicast** (the default).
- **encryption-algorithm**—Encryption used for communications between the group server and group member. Specify **aes-128-cbc**, **aes-192-cbc**, or **aes-256-cbc**.
- **lifetime-seconds** ***seconds***—Lifetime, in seconds, of the key encryption key (KEK). Specify a value from 180 to 86,400. The default is 3600 seconds.
- **number-of-retransmission** ***number***—For unicast communications, the number of times the group server retransmits messages to a group member when there is no reply. Specify a value from 0 to 60. The default is 2.
- **retransmission-period** ***seconds***—The time period between a transmission and the first retransmission when there is no reply from the group member. Specify a value from 2 to 60. The default is 10 seconds.
- **sig-hash-algorithm**—Authentication algorithm used to authenticate the group member to the group server. Specify **sha-256** or **sha-384**.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

service (Security IPsec)

Syntax service (all | *service-name*);

Hierarchy Level [edit security ipsec vpn *vpn-name* ike proxy-identity]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the service (port and protocol combination) to protect.

Options ***service-name*** —Name of the service, as defined with **system-services** (**Interface Host-Inbound Traffic**) and **system-services** (**Zone Host-Inbound Traffic**).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

session-affinity

Syntax session-affinity ipsec

Hierarchy Level [edit security flow load-distribution]

Release Information Statement introduced in Junos OS Release 11.4R5.
Starting with Junos OS Release 15.1X49-D10, IPsec session affinity is supported for IPsec tunnel-based traffic by the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) for SRX5400, SRX5600, and SRX5800 devices through improved flow module and session cache.

Description Enable VPN session affinity.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

source-interface (Security)

Syntax source-interface *interface-name* ;

Hierarchy Level [edit security ipsec vpn *vpn-name* vpn-monitor]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the source interface for ICMP requests (VPN monitoring “hellos”). If no source interface is specified, the device automatically uses the local tunnel endpoint interface.

Options *interface-name* —Name of the interface for the ICMP requests.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

spi (IPsec SA for OSPF)

Syntax spi *spi-value*;

Hierarchy Level [edit security ipsec security-association *sa-name* mode transport manual direction bidirectional]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Configure a security parameter index (SPI) for a manual IPsec security association (SA) to be applied to an OSPF or OSPFv3 interface or virtual link.

Options **spi**—SPI for the manual SA. The SPI uniquely identifies the SA to use at the receiving host (the destination address in the packet).

Range: 256 through 16,639

Required Privilege Level view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

Copyright © 2016, Juniper Networks, Inc.

spi (Security IPsec)

Syntax spi *spi-value*;

Hierarchy Level [edit security ipsec vpn *vpn-name* manual]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure a security parameter index (SPI) for a security association (SA).

Options ***spi-value*** —An arbitrary value that uniquely identifies which security association (SA) to use at the receiving host (the destination address in the packet).

Range: 256 through 16,639

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

threshold (Security IKE Gateway)

Syntax threshold *number*;

Hierarchy Level [edit security ike gateway *gateway-name* dead-peer-detection]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable.

Options ***number*** —Maximum number of unsuccessful DPD requests to be sent.

Range: 1 through 5

Output: 5

NOTE: **The threshold number for the IKEv2 protocol is predefined as 5.**

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security Group VPN)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  gateway-filter {
    local-address ip-address;
    remote-address ip-address;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level

[edit security group-vpn member ike]
[edit security group-vpn server]

Release Information

Statement introduced in Junos OS Release 10.2. Support for **gateway-filter** options and for the [**edit security group-vpn member ike**] hierarchy level added in Junos OS Release 15.1X49-D30 for vSRX.

Description

Configure group VPN trace options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Copyright © 2016, Juniper Networks, Inc.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: ***x*** **k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace all activity.
  - **certificates**—Trace certificate-related activity.
  - **config**—Trace configuration activity.
  - **database**—Trace SA-related database activity.
  - **general**—Trace general activity.
  - **high-availability**—Trace high-availability operations.
  - **ike**—Trace IKE protocol activity.
  - **next-hop-tunnels**—Trace next-hop tunnel operations.
  - **parse**—Trace configuration processing.
  - **policy-manager**—Trace IKE callback activity.
  - **routing-socket**—Trace routing socket activity.
  - **thread**—Trace thread processing.
  - **timer**—Trace timer activity.
- **gateway-filter**—Configure debugging for the tunnel between the group VPN server and a group member. This option is configured on a group VPN server or member.
  - **local-address**—When configured on a server, the IP address of the group VPN server. When configured on a member, the IP address of the group VPN member.
  - **remote-address**—When configured on a server, the IP address of the group VPN member. When configured on a member, the IP address of the group VPN server.

Copyright © 2016, Juniper Networks, Inc.

- **level**—Set the level of debugging.
  - **all**—Match all levels.
  - **error**—Match error conditions.
  - **info**—Match informational messages.
  - **notice**—Match conditions that should be handled specifically.
  - **verbose**—Match verbose messages.
  - **warning**—Match warning messages.
- **no-remote-trace**—Disable remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Group VPNv2 Overview on page 6649

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security IKE)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
  rate-limit messages-per-second;
}
```

Hierarchy Level [edit security ike]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure IKE tracing options.

Options

- file—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.
  - **files** ***number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0,** then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match** ***regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size** ***maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

    Syntax: ***x*** **k** to specify KB, ***x*** **m** to specify MB, or ***x*** **g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace all iked process modules activity
  - **certificates**—Trace certificate-related activity
  - **config**—Trace configuration download processing
  - **database**—Trace VPN-related database activity
  - **general**—Trace general activity
  - **high-availability**—Trace high-availability operations
  - **ike**—Trace IKE protocol activity
  - **next-hop-tunnels**—Trace next-hop tunnels operations
  - **parse**—Trace VPN parsing activity
  - **policy-manager**—Trace iked callback activity
  - **routing-socket**—Trace routing socket activity
  - **thread**—Trace thread processing
  - **timer**—Trace timer activity
- **no-remote-trace**—Set remote tracing as disabled.
- **rate-limit *messages-per-second***—Configure the incoming rate of trace messages.

  Range: 0 through 4,294,967,295

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

traceoptions (Security IPsec)

Syntax

```
traceoptions {
  flag flag;
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure IPsec tracing options.

NOTE: Configure IPsec tracing options only when instructed to do so by your Juniper support representative.

Trace operations are written to the trace file **/var/log/kmd**.

Options

- **flag**—To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
  - **next-hop-tunnel-binding**—Trace next-hop tunnel binding events
  - **packet-drops**—Trace packet drop activity
  - **packet-processing**—Trace data packet processing events
  - **security-associations**—Trace security association (SA) management events

Required Privilege Level trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Security PKI)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit security pki]

Release Information Statement modified in Junos OS Release 8.5.

Description Configure public key infrastructure (PKI) tracing options.

Options

- **file**—Configure the trace file options.
  - ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.
  - **files *number***—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

    If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

    Range: 2 through 1000 files

    Default: 10 files
  - **match *regular-expression***—Refine the output to include lines that contain the regular expression.
  - **size *maximum-file-size***—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the **trace-file** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

    If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

    Syntax: ***x* K** to specify KB, ***x* m** to specify MB, or ***x* g** to specify GB

    Range: 10 KB through 1 GB

Copyright © 2016, Juniper Networks, Inc.

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.
  - **all**—Trace with all flags enabled
  - **certificate-verification**—Trace PKI certificate verification events
  - **online-crl-check**—Trace PKI online certificate revocation list (CRL) events
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- Understanding Certificates and PKI on page 6347

traffic-selector

Syntax

```
traffic-selector traffic-selector-name {
  local-ip ip-address/netmask;
  remote-ip ip-address/netmask;
}
```

Hierarchy Level
[edit security ipsec vpn *vpn-name*]

Release Information
Statement introduced in Junos OS Release 12.1X46-D10.

Description
Configure local and remote IP addresses for a traffic selector.

Options
local-ip *ip-address/netmask*—A local IP address or a local subnetwork protected by the local VPN device.

remote-ip *ip-address/netmask*—A remote IP address or a remote subnetwork protected by the peer VPN device.

Required Privilege Level
security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

trusted-ca (Security IKE Policy)

Syntax trusted-ca (*ca-index* | use-all);

Hierarchy Level [edit security ike policy *policy-name* certificate]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify the preferred certificate authority (CA) to use when requesting a certificate from the peer. If no value is specified, then no certificate request is sent (although incoming certificates are still accepted).

Options

- ***ca-index***—Preferred certificate authority ID for the device to use.
- **use-all**—Device uses all configured certificate authorities.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

use-ocsp (Security PKI)

Syntax use-ocsp;

Hierarchy Level [edit security pki ca-profile *ca-profile-name*]

Release Information Statement introduced in Junos OS Release 12.1X46-D20.

Description Specify the Online Certificate Status Protocol (OCSP) as the method to check the revocation status of a certificate. CRL is the default method.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- Understanding Certificates and PKI on page 6347

Copyright © 2016, Juniper Networks, Inc.

user-at-hostname

Syntax user-at-hostname *e-mail-address*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure an e-mail address.

Options ***e-mail-address*** —Valid e-mail address.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

version (Security IKE Gateway)

Syntax version (v1-only | v2-only);

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 11.3.

Description Specify the IKE version to use to initiate the connection.

Options **v1-only**—The connection must be initiated using IKE version 1. This is the default.

v2-only—The connection must be initiated using IKE version 2.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

vpn (Security)

Syntax

```
vpn vpn-name {
  bind-interface interface-name;
  copy-outer-dscp;
  ike {
  gateway gateway-name;
  ipsec-policy ipsec-phase2-policy;
  }
  establish-tunnels (immediately | on-traffic);
  }
  ike {
    gateway gateway-name;
    idle-time seconds;
    install-interval seconds;
    ipsec-policy ipsec-policy-name;
    no-anti-replay;
    proxy-identity {
      local ip-prefix;
      remote ip-prefix;
      service (any | service-name);
    }
  }
  manual {
    authentication {
      algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
      key (ascii-text key | hexadecimal key);
    }
    encryption {
      algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
      key (ascii-text key | hexadecimal key);
    }
    external-interface external-interface-name;
    gateway ip-address;
    protocol (ah | esp);
    spi spi-value;
  }
  traffic-selector traffic-selector-name {
    local-ip ip-address/netmask;
    remote-ip ip-address/netmask;
  }
  vpn-monitor {
    destination-ip ip-address;
    optimized;
    source-interface interface-name;
  }
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. Support for **copy-outer-dscp** added in Junos OS Release 15.1X49-D30.

Copyright © 2016, Juniper Networks, Inc.

Description Configure an IPsec VPN.

Options *vpn-name* —Name of the VPN.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

vpn-monitor

Syntax

```
vpn-monitor {
  destination-ip ip-address;
  optimized;
  source-interface interface-name;
}
```

Hierarchy Level [edit security ipsec vpn *vpn-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure settings for VPN monitoring. This feature cannot be configured simultaneously with the **dead-peer-detection** statement.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation
- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

vpn-monitor-options

Syntax

```
vpn-monitor-options {
    interval seconds;
    threshold number;
}
```

Hierarchy Level [edit security ipsec]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure VPN monitoring options.

Options

- **interval** ***seconds*** —Interval at which to send ICMP requests to the peer.

  Range: 2 through 3600 seconds
  Default: 10 seconds

- **threshold** ***number*** —Number of consecutive unsuccessful pings before the peer is declared unreachable.

  Range: 1 through 65,536 pings
  Default: 10 pings

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

wildcard

Syntax wildcard *string*;

Hierarchy Level [edit security ike gateway *gateway-name* dynamic distinguished-name]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that the values of a dynamic virtual private network (VPN) endpoint user's distinguished name's identity fields match the values in the group IKE user's distinguished name's fields. The order of the identity fields in the distinguished name strings does not matter during a match.

Options ***string***—Distinguished name identity values to be matched.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

xauth

Syntax

```
xauth {
  access-profile profile-name;
}
```

Hierarchy Level [edit security ike gateway *gateway-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify that Extended authentication (XAuth) is performed in addition to IKE authentication for remote users trying to access a VPN tunnel. Include a previously created access profile, created with the **edit access profile** statement, to specify the access profile to be used for authentication information.

Options

access-profile ***profile-name***—Name of previously created access profile to reference for authentication information.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- IPsec VPN Overview on page 6041

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 282

Operational Commands

- clear security group-vpn member group
- clear security group-vpn member ike security-associations
- clear security group-vpn member ipsec security-associations
- clear security group-vpn member ipsec security-associations statistics
- clear security group-vpn member ipsec statistics
- clear security group-vpn server
- clear security group-vpn server server-cluster statistics
- clear security group-vpn server statistics
- clear security ike respond-bad-spi-count
- clear security ike security-associations
- clear security ipsec security-associations
- clear security ipsec statistics
- clear security ipsec tunnel-events-statistics
- clear security pki key-pair (Local Certificate)
- clear security pki local-certificate (Device)
- request security pki ca-certificate ca-profile-group load
- request security pki ca-certificate enroll (Security)
- request security pki ca-certificate load (Security)
- request security pki ca-certificate verify (Security)
- request security pki crl load (Security)
- request security pki generate-certificate-request (Security)
- request security pki generate-key-pair (Security)
- request security pki local-certificate enroll cmpv2
- request security pki local-certificate enroll scep
- request security pki local-certificate export
- request security pki local-certificate generate-self-signed (Security)
- request security pki local-certificate load
- request security pki local-certificate verify (Security)

Copyright © 2016, Juniper Networks, Inc.

- request security pki verify-integrity-status
- show security group-vpn member ike security-associations
- show security group-vpn member ipsec inactive-tunnels
- show security group-vpn member ipsec security-associations
- show security group-vpn member ipsec statistics
- show security group-vpn member kek security-associations
- show security group-vpn member policy
- show security group-vpn server ike security-associations
- show security group-vpn server ipsec security-associations
- show security group-vpn server kek security-associations
- show security group-vpn server registered-members
- show security group-vpn server server-cluster
- show security group-vpn server statistics
- show security ike active-peer
- show security ike pre-shared-key
- show security ike security-associations
- show security ike tunnel-map
- show security ipsec control-plane-security-associations
- show security ipsec inactive-tunnels
- show security ipsec next-hop-tunnels
- show security ipsec security-associations
- show security ipsec statistics
- show security ipsec traffic-selector
- show security ipsec tunnel-events-statistics
- show security pki ca-certificate (View)
- show security pki certificate-request (View)
- show security pki crl (View)
- show security pki local-certificate (View)

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn member group

Syntax clear security group-vpn member group <vpn *vpn-name*> <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Clear all current information for IKE, TEK, and KEK SAs.

Options **none**—Clear SA information for all groups.

vpn *vpn-name*—(Optional) Clear SA information for the specified VPN name.

group-id *group-id*—(Optional) Clear SA information for the specified group identifier.

Required Privilege Level clear

Related Documentation

- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn member ike security-associations

Syntax clear security group-vpn member ike security-associations [index *SA-index*] [*peer-ipaddress*]

Release Information Command introduced in Junos OS Release 10.2.

Description Clear IKE security association (SA) for a group member.

Options
- none—Clear all IKE SAs for the group member.
- **index**—(Optional) Clear the IKE SA with this index number.
- *peer-ipaddress*—(Optional) Clear the IKE SA with this peer.

Required Privilege Level clear

Related Documentation
- show security group-vpn member ike security-associations on page 7010
- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn member ipsec security-associations

Syntax clear security group-vpn member ipsec security-associations [index *SA-index*]

Release Information Command introduced in Junos OS Release 10.2.

Description Clear group VPN SA for a group member.

Options
- none—Clear all group VPN SAs for the group member.
- **index**—(Optional) Clear the group VPN SA with this index number.

Required Privilege Level clear

Related Documentation
- show security group-vpn member ipsec security-associations on page 7016
- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn member ipsec security-associations statistics

Syntax clear security group-vpn member ipsec security-associations statistics <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Clear IPsec SA statistics.

Options **none**—Clear IPsec SA statistics for all groups.

group-id *group-id*—(Optional) Clear IPsec SA statistics for the specified group identifier.

Required Privilege Level clear

Related Documentation
- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn member ipsec statistics

Syntax clear security group-vpn member ipsec statistics <index *index*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Clear IPsec statistics.

Options **none**—Clear IPsec statistics for all groups.

index *index*—(Optional) Clear the IPsec statistics for the SA with this index number.

Required Privilege Level clear

Related Documentation

- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn server

Syntax clear security group-vpn server [group *group-name* | group-id *group-id*] [now]

Description Clear active members for a specified group. If no options are specified, members are cleared from all groups. After this command is issued, members will need to reregister.

NOTE: An IKE SA can be used by a group member to register to multiple groups. When you clear members for a specified group, all existing IKE SAs that could be used to register to the group are also cleared.

Options

- none—All members are cleared from all groups.
- **group**—(Optional) Clear members and SAs for the specified group name.
- **group-id**—(Optional) Clear members and SAs for the specified group identifier.
- **now**—(Optional) Immediately clear all group-related information.

Required Privilege Level clear

Related Documentation

- show security group-vpn server registered-members on page 7037
- Group VPNv2 Overview on page 6649

Output Fields If there is a problem with the command, one of the following messages appears:

- Group does not exist
- Group is in the process of deletion
- Error in clear members
- Warning Message; Fail to push delete to members as server-member-communication is not configured.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn server server-cluster statistics

Syntax clear security group-vpn server server-cluster statistics <group *group-name*> <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Clear Group VPNv2 server cluster statistics.

Options **none**—Clear Group VPNv2 server cluster statistics for all groups.

group ***group-name***—(Optional) Clear Group VPNv2 server cluster statistics for the specified group name.

group-id ***group-id***—(Optional) Clear Group VPNv2 server cluster statistics for the specified group identifier.

Required Privilege Level clear

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security group-vpn server statistics

Syntax clear security group-vpn server statistics <group *group-name*> <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Clear group statistics.

Options

none—Clear statistics for all groups.

group *group-name*—(Optional) Clear statistics for the specified group name.

group-id *group-id*—(Optional) Clear statistics for the specified group identifier.

Required Privilege Level clear

Related Documentation

- show security group-vpn server statistics on page 7042
- Group VPNv2 Overview on page 6649

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security ike respond-bad-spi-count

Syntax
clear security ike respond-bad-spi-count
< *gateway-name* >

Release Information
Command introduced in Junos OS Release 8.5.

Description
Clear information about invalid Internet Key Exchange (IKE) security parameter index (SPI) counters.

Options
- none—Clear all invalid SPI counters.
- ***gateway-name*** —(Optional) Clear the invalid SPI counters for the given gateway.

Required Privilege Level
clear

Related Documentation
- respond-bad-spi on page 6947

Output Fields
This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security ike security-associations

Syntax

clear security ike security-associations
< *peer-address* >
< **port** >
<fpc *slot-number*>
<index *SA-index-number*>
<kmd-instance (all | *kmd-instance-name*)>
<pic *slot-number*>
port
<family (inet | inet6)>

Release Information

Command introduced in Junos OS Release 8.5. The **fpc, pic,** and **kmd-instance** options added in Junos OS Release 9.3. The **port** option added in Junos OS Release 10.0. The **family** option added in Junos OS Release 11.1.

Description

Clear information about the current Internet Key Exchange security associations (IKE SAs). For IKEv2, the device clears the information about the IKE SAs and the associated IPSec SA.

Options

- none—Clear all IKE SAs.
- ***peer-address*** —(Optional) Clear IKE SAs for the destination peer at this IP address.
- **fpc** ***slot-number*** —Specific to SRX Series devices. Clear information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot.
- **index** ***SA-index-number*** —(Optional) Clear the IKE SA with this index number.
- **port**—(Optional) Port number of SA (1 through 65,535).
- **kmd-instance**—Specific to SRX Series devices. Clear information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC ***slot-number*** and PIC ***slot-number***.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number*** —Specific to SRX Series devices. Clear information about existing IKE SAs in this PIC slot.
- **family**—(Optional) Clear IKE SAs by family.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.

Required Privilege Level

clear

Related Documentation

- show security ike security-associations on page 3673

Copyright © 2016, Juniper Networks, Inc.

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security ipsec security-associations

Syntax

clear security ipsec security-associations
fpc *slot-number*
<index *SA-index-number*>
kmd-instance (all | *kmd-instance-name*)
pic *slot-number*
<family (inet | inet6)>

Release Information

Command introduced in Junos OS Release 8.5. The **fpc**, **pic,** and **kmd-instance** options added in Junos OS Release 9.3. The **family** option added in Junos OS Release 11.1.

Description

Clear information about IPsec security associations (SAs).

Options

- none—Clear all IPsec SAs.
- **fpc** *slot-number* —Specific to SRX Series devices. Clear information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot.
- **index** *SA-index-number* —(Optional) Clear the IPsec SA with this index number.
- **kmd-instance**—Specific to SRX Series devices. Clear information about existing IPsec SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *slot-number* and PIC *slot-number* .
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - *kmd-instance-name*—Name of the KMD instance running on the SPU.

pic *slot-number* —Specific to SRX Series devices. Clear information about existing IPsec SAs in this PIC slot.

family—(Optional) Clear SAs by family.

- **inet**—IPv4 address family.
- **inet6**—IPv6 address family.

Required Privilege Level

clear

Related Documentation

- show security ipsec security-associations on page 3681

Output Fields

This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security ipsec statistics

Syntax

clear security ike statistics
<fpc *slot-number*>
<index *SA-index-number*>
<kmd-instance (all | *kmd-instance-name*)>
<pic *slot-number*>

Release Information

Command introduced in Junos OS Release 8.5. **fpc and pic** options added in Junos OS Release 9.3. **kmd-instance** option added in Junos OS Release 10.4.

Description

Clear IPsec statistics on the device.

Options

- none—Clear all IPsec statistics.
- **fpc** ***slot-number*** —Specific to SRX Series devices. Clear statistics about existing IPsec security associations (SAs) in this Flexible PIC Concentrator (FPC) slot.
- **index** ***SA-index-number*** —(Optional) Clear the IPsec statistics for the SA with this index number.
- **kmd-instance**—Specific to SRX Series devices. Clear information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC ***slot-number*** and PIC ***slot-number*** .
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number*** —Specific to SRX Series devices. Clear statistics about existing IPsec SAs in this PIC slot.

Required Privilege Level

clear

Related Documentation

- show security ipsec statistics on page 4011

Output Fields

This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security ipsec tunnel-events-statistics

Syntax clear security ipsec tunnel-events-statistics

Release Information Command introduced in Junos OS Release 12.3X48-D10.

Description Clear IPsec tunnel event statistics.

Required Privilege Level clear

Related Documentation
- *show security ipsec tunnel-events-statistics*

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security pki key-pair (Local Certificate)

Syntax clear security pki key-pair (all | certificate-id *certificate-id*)

Release Information Command introduced in Junos OS Release 8.5.

Description Clear public key infrastructure (PKI) key pair information for local digital certificates on the device.

Options

- **all**—Clear key pair information for all local certificates.
- **certificate-id** ***certificate-id*** —Clear key pair information for the local certificate with this certificate ID.

Required Privilege Level clear and security

Related Documentation

- show security pki certificate-request (View) on page 7084

Output Fields This command produces no output.

Copyright © 2016, Juniper Networks, Inc.

clear security pki local-certificate (Device)

Syntax clear security pki local-certificate (all | certificate-id *certificate-id* | system-generated)

Release Information Command modified in Junos OS Release 9.1.

Description Clear public key infrastructure (PKI) information for local digital certificates on the device.

Options

- **all**—Clear information for all the local digital certificates on the device.

NOTE: You cannot clear the automatically generated self-signed certificate using clear security pki local-certificate all command. To clear the self-signed certificate you need to use system-generated as an option.

- **certificate-id** ***certificate-id*** —Clear the specified local digital certificate with this certificate ID.
- **system-generated**—Clear the existing automatically generated self-signed certificate and generate a new self-signed certificate.

Required Privilege Level clear and security

Related Documentation

- show security pki local-certificate (View) on page 797
- request security pki local-certificate generate-self-signed (Security) on page 7005

List of Sample Output

clear security pki local-certificate all on page 6990
clear security pki local-certificate system-generated on page 6990

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear security pki local-certificate all

```
user@host> clear security pki local-certificate all
```

Sample Output

clear security pki local-certificate system-generated

```
user@host> clear security pki local-certificate system-generated
```


Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate ca-profile-group load

Syntax request security pki ca-certificate ca-profile-group load ca-group-name *ca-group-name* filename [*path/filename* | default]

Release Information Command introduced in Junos OS Release 12.1; **default** option added in Junos OS Release 12.1X47-D10.

Description For SSL forward proxy, you need to load trusted CA certificates on your system. By default, Junos OS provides a list of trusted CA certificates that include default certificates used by common browsers. Alternatively, you can define your own list of trusted CA certificates and import them on to your system.

Use this command to load the default certificates or to specify a path and filename of trusted CA certificates that you define.

Options **ca-group-name *ca-group-name***—Load the specified CA group profile.

filename *path/filename*—Directory location and filename of the trusted CA certificates defined by you.

filename default—Load the trusted CA certificates available by default.

Required Privilege Level maintenance

Related Documentation
- show security pki ca-certificate on page 793
- Understanding Certificates and PKI on page 6347

List of Sample Output
request security pki ca-certificate ca-profile-group load (default) on page 6991
request security pki ca-certificate ca-profile-group load (path/filename) on page 6992

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate ca-profile-group load (default)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-default
filename default

Do you want to load this CA certificate ? [yes,no] (no) yes
Loading 157 certificates for group 'ca-default'.
ca-default_1: Loading done.
ca-default_2: Loading done.
ca-default_3: Loading done.
.....
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

request security pki ca-certificate ca-profile-group load (path/filename)

```
user@host> request security pki ca-certificate ca-profile-group load ca-group-name ca-manual
filename /var/tmp/firefox-all.pem

Do you want to load this CA certificate ? [yes,no] (no) yes

Loading 196 certificates for group 'ca-manual'.
ca-manual_1_sysgen: Loading done.
ca-manual_2_sysgen: Loading done.
ca-manual_3_sysgen: Loading done.
ca-manual_4_sysgen: Loading done.
ca-manual_5_sysgen: Loading done.
ca-manual_6_sysgen: Loading done.

...
ca-manual_195_sysgen: Loading done.
ca-manual_196_sysgen: Loading done.
ca-profile-group 'ca-manual' successfully loaded. Success[193] Skipped[3]
```


Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate enroll (Security)

Syntax request security pki ca-certificate enroll ca-profile *ca-profile-name*

Release Information Command introduced in Junos OS Release 7.5.

Description Request a digital certificate from a certificate authority (CA) online by using the Simple Certificate Enrollment Protocol (SCEP).

Options **ca-profile *ca-profile-name***—CA profile name.

Required Privilege Level maintenance

Related Documentation
- show security pki ca-certificate (View) on page 7080
- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki ca-certificate enroll on page 6993

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate enroll

```
user@host> request security pki ca-certificate enroll ca-profile entrust
Received following certificates:
  Certificate: C=us, O=example, CN=First Officer
    Fingerprint: 46:71:15:34:f0:a6:41:76:65:81:33:4f:68:47:c4:df:78:b8:e3:3f
  Certificate: C=us, O=example, CN=First Officer
    Fingerprint: bc:78:87:9b:a7:91:13:20:71:db:ac:b5:56:71:42:ad:1a:b6:46:17
  Certificate: C=us, O=example
    Fingerprint: 00:8e:6f:58:dd:68:bf:25:0a:e3:f9:17:70:d6:61:f3:53:a7:79:10
Do you want to load the above CA certificate ? [yes,no] (no) yes
```

Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate load (Security)

Syntax request security pki ca-certificate load ca-profile *ca-profile-name* filename *path/filename*

Release Information Command introduced in Junos OS Release 7.5.

Description Manually load a certificate authority (CA) digital certificate from a specified location.

Options **ca-profile *ca-profile-name***—Load the specified CA profile.

filename *path/filename*—Directory location and filename of the CA digital certificate.

Required Privilege Level maintenance

Related Documentation
- show security pki ca-certificate on page 793
- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki ca-certificate load on page 6994

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki ca-certificate load

```
user@host> request security pki ca-certificate load ca-profile 2Kkey filename /var/tmp/2Kkey.pem

Fingerprint:
  a0:08:bb:1f:75:96:76:cd:ee:db:36:10:b6:c6:d8:df:5e:02:05:05 (sha1)
  f5:58:6b:de:7c:d6:cd:90:5a:18:c3:0e:3d:95:da:25 (md5)
Do you want to load this CA certificate ? [yes,no] (no) yes

CA certificate for profile 2Kkey loaded successfully
```


Copyright © 2016, Juniper Networks, Inc.

request security pki ca-certificate verify (Security)

Syntax **request security pki ca-certificate verify ca-profile** ***ca-profile-name***

Release Information Command introduced in Junos OS Release 8.5.

Description Verify the digital certificate installed for the specified certificate authority (CA).

Options **ca-profile** ***ca-profile-name*** —Display the specified CA profile.

Required Privilege Level maintenance and security

Related Documentation

- ca-profile (Security PKI) on page 6847
- show security pki ca-certificate (View) on page 7080
- Understanding Certificates and PKI on page 6347

List of Sample Output

request security pki ca-certificate verify ca-profile ca1 (CRL downloaded) on page 6995
request security pki ca-certificate verify ca-profile ca1 (CRL not downloaded) on page 6995

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

This user has downloaded the certificate revocation list (CRL).

request security pki ca-certificate verify ca-profile ca1 (CRL downloaded)

```
user@host> request security pki ca-certificate verify ca-profile ca1
CA certificate ca1 verified successfully
```

Sample Output

This user has not downloaded the certificate revocation list (CRL).

request security pki ca-certificate verify ca-profile ca1 (CRL not downloaded)

```
user@host> request security pki ca-certificate verify ca-profile ca1
CA certificate ca1: CRL verification in progress. Please check the PKId debug
logs for completion status
```

Copyright © 2016, Juniper Networks, Inc.

request security pki crl load (Security)

Syntax request security pki crl load ca-profile *ca-profile-name* filename *path/filename*

Release Information Command introduced in Junos OS Release 8.1.

Description Manually install a certificate revocation list (CRL) on the device from a specified location.

Options **ca-profile *ca-profile-name*** —Load the specified certificate authority (CA) profile.

filename *path/filename* —Directory location and filename of the CRL.

Required Privilege Level maintenance

Related Documentation
- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki crl load on page 6996

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki crl load

```
user@host> request security pki crl load ca-profile ca-test filename example-inter-ca.crl
CRL for CA profile ca-test loaded successfully
```


Copyright © 2016, Juniper Networks, Inc.

request security pki generate-certificate-request (Security)

Syntax

request security pki generate-certificate-request certificate-id *certificate-id-name*
domain-name *domain-name* subject *subject-distinguished-name*
<add-ca-constraint>
<digest (sha1 | sha256)>
<email *email-address*>
<filename (*path* | terminal)>
<ip-address *ip-address*>

Release Information

Command introduced in Junos OS Release 7.5. Support for **digest** option added in Junos OS Release 12.1X45-D10.

Description

Manually generate a local digital certificate request in the Public-Key Cryptography Standards #10 (PKCS-10) format.

Options

certificate-id *certificate-id-name*—Name of the local digital certificate and the public/private key pair.

domain-name *domain-name*—Fully qualified domain name (FQDN) provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

subject *subject-distinguished-name*—Distinguished name format contains the following information:

- **DC**—Domain component
- **CN**—Common name
- **OU**—Organizational unit name
- **O**—Organization name
- **L**—Locality
- **ST**—State
- **C**—Country

digest—(Optional) Hash algorithm used to sign the certificate request.

- **sha1**—SHA-1 digests (default value for RSA or DSA only).
- **sha256**—SHA-256 digests for RSA or ECDSA only (default value for ECDSA).
- **sha-384**—SHA-384 digests for ECDSA only.

email *email-address*—(Optional) E-mail address of the certificate holder.

filename (*path* | terminal)—(Optional) Location where the local digital certificate request should be placed or the login terminal.

ip-address *ip-address*—(Optional) IP address of the router.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level maintenance

Related Documentation
- show security pki certificate-request (View) on page 7084

List of Sample Output request security pki generate-certificate-request on page 6998

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-certificate-request

```
user@host> request security pki generate-certificate-request certificate-id local-entrust2
domain-name router2.example.net filename entrust-req2 subject cn=router2.example.net

Generated certificate request
-----BEGIN CERTIFICATE REQUEST-----
MIIBoTCCAQoCAQAwGjEYMBYGA1UEAxMPdHAxLmp1bmlwZXIubmV0MIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQCiUFklQws1Ud+AqN5DDxRs2kVyKEhh9qoVFnz+
Hz4c9vsy3B8ElwTJlkmIt2cB3yifB6zePd+6WYpf57Crwre7YqPkiXM31F6z3YjX
H+1BPNbCxNWYvyrnSyVYDbFj8oOXyqog8ACDfVL2JBWrPNBYy7imq/K9soDBbAs6
5hZqqwIDAQABoEcwRQYJKoZIhvcNAQkOMTgwNjAOBgNVHQ8BAf8EBAMCB4AwJAYD
VR0RAQH/BBowGIIWdHAxLmVuZ2xhYi5qdW5pcGVyLm5ldDANBgkqhkiG9w0BAQQF
AAOBgQBc2rq1v5SOQXH7LCb/FdqAL8ZM6GoaN5d6cGwq4bB6a7UQFgtoH4O6gQ3G
3iHOZfz4xMIBpJYuGd1dkqgvcDoH3AgTsLkfn7Wi3x5H2qeQVs9bvL4P5nvEZLND
EIMUHwteolZCiZ7OfO9Fer9cXWHSQs1UtXtgPqQJy2xIeImLgw==
-----END CERTIFICATE REQUEST-----
Fingerprint:
0d:90:b8:d2:56:74:fc:84:59:62:b9:78:71:9c:e4:9c:54:ba:16:97 (sha1)
1b:08:d4:f7:90:f1:c4:39:08:c9:de:76:00:86:62:b8 (md5)
```


Copyright © 2016, Juniper Networks, Inc.

request security pki generate-key-pair (Security)

Syntax request security pki generate-key-pair certificate-id *certificate-id-name*
<size (256 | 384 | 512 | 1024 | 2048 | 4096)>
<type (dsa | ecdsa | rsa)>

Release Information Command introduced in Junos OS Release 11.1. Options to support Elliptic Curve Digital Signature Algorithm (ECDSA) added in Junos OS Release 12.1X45-D10.

Description Generate a public key infrastructure (PKI) public/private key pair for a local digital certificate.

Options

certificate-id *certificate-id-name*—Name of the local digital certificate and the public/private key pair.

size—Key pair size. The key pair size can be 256, 384, 512, 1024, 2048, or 4096 bits. Key pair sizes of 256 and 384 bits are compatible with ECDSA. If a key pair size is not specified, the default value, **1024** bits, is applied.

type—The algorithm to be used for encrypting the public/private key pair:

- **ecdsa**—ECDSA encryption
- **dsa**—Digital Signal Algorithm (DSA) encryption
- **rsa**—Rivest Shamir Adleman (RSA) encryption (default)

Required Privilege Level maintenance

Related Documentation

- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki generate-key-pair on page 6999

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki generate-key-pair

```
user@switch> request security pki generate-key-pair type rsa size 1024 certificate-id test
Generated key pair test, key size 1024 bits
```

Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate enroll cmpv2

Syntax

```
request security pki local-certificate enroll cmpv2
  ca-dn subject-dn
  ca-profile ca-profile name
  ca-reference reference
  ca-secret shared-secret
  certificate-id certificate-id-name
  domain-name domain-name
  email email-address
  ip-address ip-address
  ipv6-address ipv6-address
  subject subject-distinguished-name
```

Release Information Command introduced in Junos OS Release 15.1X49-D40.

Description Enroll and install a local digital certificate online by using CMPv2. This command loads both end-entity (EE) and CA certificates based on the CA server configuration. Certificate revocation list (CRL) or Online Certificate Status Protocol (OCSP) can be used to check the revocation status of a certificate.

Options

ca-dn *subject-dn*—The distinguished name (DN) of the CA enrolling the EE certificate must be specified during enrollment. This optional parameter is mandatory if the CA certificate is not already enrolled. If the CA certificate is already enrolled, the subject DN is extracted from the CA certificate.

ca-profile *ca-profile-name*—CA profile name.

ca-reference *reference*—Out-of-band reference value received from the CA server.

ca-secret *shared-secret*—Out-of-band secret value received from the CA server.

certificate-id *certificate-id-name*—Name of the local digital certificate and the public/private key pair.

domain-name *domain-name*—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

email *email-address*—E-mail address of the certificate holder.

ip-address *ip-address*—IP address of the router.

ipv6-address *ipv6-address*—IPv6 address of the router for the alternate subject.

subject *subject-distinguished-name*—Distinguished Name (DN) format that contains the domain component, common name, department, serial number, company name, state, and country in the following format: DC, CN, OU, O, SN, L, ST, C.

- **DC**—Domain component
- **CN**—Common name
- **OU**—Organizational unit name

Copyright © 2016, Juniper Networks, Inc.

- **O**—Organization name
- **SN**—Serial number of the device

NOTE: If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR).

- **ST**—State
- **C**—Country

Required Privilege Level maintenance and security

Related Documentation
- show security pki local-certificate (View) on page 797
- clear security pki local-certificate (Device) on page 6990

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> request security pki local-certificate enroll cmpv2 ca-profile root-552 ca-dn
DC=Example,CN=root-552 certificate-id tc552 email tc552-root@example.net domain-name
example.net ip-address 11.0.0.22 ca-secret example ca-reference 51892 subject
CN=Example,OU=Sales,O=552-22

Certificate enrollment has started. To view the status of your enrollment, check
 the public key infrastructure log (pkid) log file at /var/log/pkid.
```

Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate enroll scep

Syntax

```
request security pki local-certificate enroll scep
  ca-profile ca-profile name
  certificate-id certificate-id-name
  challenge-password challenge-password
  digest (sha-1 | sha-256)
  domain-name domain-name
  email email-address
  ip-address ip-address
  ipv6-address ipv6-address
  scep-digest-algorithm (md5 | sha-1)
  scep-encryption-algorithm (des | des3)
  subject subject-distinguished-name
```

Release Information Command introduced in Junos OS Release 9.1. Serial number (SN) option added to the subject string output field in Junos OS Release 12.1X45. **scep** keyword and **ipv6-address** option added in Junos OS Release 15.1X49-D40.

Description Enroll and install a local digital certificate online by using Simple Certificate Enrollment Protocol (SCEP).

Options

ca-profile *ca-profile-name*—CA profile name.

certificate-id *certificate-id-name*—Name of the local digital certificate and the public/private key pair.

challenge-password *password*—Password set by the administrator and normally obtained from the SCEP enrollment webpage of the CA. The password is 16 characters in length.

digest (sha-1 | sha-256)—Hash algorithm used for signing RSA certificates, either SHA-1 or SHA-256. SHA-1 is the default.

domain-name *domain-name*—Fully qualified domain name (FQDN). The FQDN provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

email *email-address*—E-mail address of the certificate holder.

ip-address *ip-address*—IP address of the router.

ipv6-address *ipv6-address*—IPv6 address of the router for the alternate subject.

scep-digest-algorithm (md5 | sha-1)—Hash algorithm digest, either MD5 or SHA-1; SHA-1 is the default.

scep-encryption-algorithm (des | des3)—Encryption algorithm, either DES or DES3; DES3 is the default.

subject *subject-distinguished-name*—Distinguished Name (DN) format that contains the domain component, common name, department, serial number, company name, state, and country in the following format: DC, CN, OU, O, SN, L, ST, C.

Copyright © 2016, Juniper Networks, Inc.

- **DC**—Domain component
- **CN**—Common name
- **OU**—Organizational unit name
- **O**—Organization name
- **SN**—Serial number of the device

NOTE: If you define SN in the subject field without the serial number, then the serial number is read directly from the device and added to the certificate signing request (CSR).

- **ST**—State
- **C**—Country

Required Privilege Level maintenance and security

Related Documentation
- show security pki local-certificate (View) on page 797
- clear security pki local-certificate (Device) on page 6990

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> request security pki local-certificate enroll scep certificate-id r3-entrust-scep
ca-profile entrust domain-name router3.example.net subject
"CN=router3,OU=Sales,O=example,C=US" challenge-password 123

Certificate enrollment has started. To view the status of your enrollment, check
 the public key infrastructure log (pkid) log file at /var/log/pkid. Please save
 the challenge-password for revoking this certificate in future.  Note that this
 password is not stored on the router.
```

Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate export

Syntax request security pki local-certificate export

Release Information Command introduced in Junos OS Release 12.1.

Description Export a generated self-signed certificate from the default location (var/db/certs/common/local) to a specific location within the device.

Options

certificate id *certificate-id-name*—Name of the local digital certificate.

filename *path/filename*—Target directory location and filename of the CA digital certificate.

type (der | pem)—Certificate format: DER (distinguished encoding rules) or PEM (privacy-enhanced mail).

Required Privilege Level maintenance

Related Documentation

- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki local-certificate export on page 7004

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate export

```
user@host> request security pki local-certificate export filename /var/tmp/my-cert.pem
certificate-id nss-cert type pem
certificate exported successfully
```


Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate generate-self-signed (Security)

Syntax

request security pki local-certificate generate-self-signed certificate-id *certificate-id-name*domain-name *domain-name* subject *subject-distinguished-name*
<add-ca-constraint>
<digest (sha1 | sha256)>
<email *email-address*>
<ip-address *ip-address*>

Release Information

Command introduced in Junos OS Release 9.1. Support for **digest** option added in Junos OS Release 12.1X45-D10.

Description

Manually generate a self-signed certificate for the given distinguished name.

Options

certificate-id *certificate-id-name*—Name of the certificate and the public/private key pair.

domain-name *domain-name*—Fully qualified domain name (FQDN) provides the identity of the certificate owner for Internet Key Exchange (IKE) negotiations and provides an alternative to the subject name.

subject *subject-distinguished-name*—Distinguished name format contains the following information:

- **DC**—Domain component
- **CN**—Common name
- **OU**—Organizational unit name
- **O**—Organization name
- **L**—Locality
- **ST**—State
- **C**—Country

add-ca-constraint—(Optional) Specifies that the certificate can be used to sign other certificates.

digest—(Optional) Hash algorithm used to sign the certificate.

- **sha1**—SHA-1 digest (default)
- **sha256**—SHA-256 digest

email *email-address*—(Optional) E-mail address of the certificate holder.

Required Privilege Level

maintenance and security

Related Documentation

- clear security pki local-certificate (Device) on page 6990
- show security pki local-certificate (View) on page 797

Copyright © 2016, Juniper Networks, Inc.

List of Sample Output request security pki local-certificate generate-self-signed certificate-id self-cert subject cn=abc domain-name example.net email mholmes@example.net on page 7006

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate generate-self-signed certificate-id self-cert subject cn=abc domain-name example.net email mholmes@example.net

```
user@host> request security pki local-certificate generate-self-signed certificate-id self-cert
subject cn=abc domain-name example.net email mholmes@example.net
Self-signed certificate generated and loaded successfully
```


Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate load

Syntax request security pki local-certificate load filename *ssl_proxy_ca.crt* key *ssl_proxy_ca.key* certificate-id *certificate id*

Release Information Command introduced in Junos OS Release 11.4.

Description Manually load a local digital certificate from a specified location.

Options **filename** — Filename that contains the certificate to load

key— File pathname that contains the private key/key-pair to loaded

certificate-id —Name of the certificate identifier

Required Privilege Level maintenance and security

Related Documentation

- show security pki local-certificate (View) on page 797
- clear security pki local-certificate (Device) on page 6990
- request security pki local-certificate verify (Security) on page 7008
- Understanding Certificates and PKI on page 6347

List of Sample Output request security pki local-certificate load on page 7007

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki local-certificate load

```
user@host> request security pki local-certificate load filename cert_name.crt key key_name.key
certificate-id test
Local certificate cert_name.crt loaded successfully
```

Copyright © 2016, Juniper Networks, Inc.

request security pki local-certificate verify (Security)

Syntax request security pki local-certificate verify certificate-id *certificate-id-name*

Release Information Command introduced in Junos OS Release 8.5.

Description Verify the validity of the local digital certificate identifier.

Options certificate-id *certificate-id-name* — Name of the local digital certificate identifier.

Required Privilege Level maintenance and security

Related Documentation

- request security pki local-certificate load on page 7007
- show security pki local-certificate (View) on page 797
- clear security pki local-certificate (Device) on page 6990
- Understanding Certificates and PKI on page 6347

List of Sample Output

request security pki local-certificate verify certificate-id bme1 (not downloaded) on page 7008
request security pki local-certificate verify certificate bme1 (downloaded) on page 7008

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

You receive the following response before the certificate revocation list (CRL) is downloaded:

request security pki local-certificate verify certificate-id bme1 (not downloaded)

```
user@host> request security pki local-certificate verify certificate-id bme1
Local certificate bme1: CRL verification in progress. Please check the PKId debug
 logs for completion status
```

Sample Output

You receive the following response after the certificate revocation list (CRL) is downloaded:

request security pki local-certificate verify certificate bme1 (downloaded)

```
user@host> request security pki local-certificate verify certificate-id bme1
Local certificate bme1 verification success
```


Copyright © 2016, Juniper Networks, Inc.

request security pki verify-integrity-status

Syntax request security pki verify-integrity-status

Release Information Command introduced in Junos OS Release 11.2.

Description Verify the integrity of public key infrastructure (PKI) files.

Required Privilege Level maintenance

Related Documentation
- [edit security pki] Hierarchy Level on page 646

List of Sample Output request security pki verify-integrity-status on page 7009

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security pki verify-integrity-status

```
user@host> request security pki verify-integrity-status
All PKI objects: verification success
```

show security group-vpn member ike security-associations

Syntax show security group-vpn member ike security-associations [brief | detail] [index *sa-index*] [*peer-ipaddress*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display IKE security associations (SAs) for group members.

Options

- none—Display summary information about all IKE SAs for the group members.
- **brief**—(Optional) Display summary output.
- **detail**—(Optional) Display detailed output.
- **index** *sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- *peer-ipaddress*—(Optional) Display information about the SA with the specified peer.

Required Privilege Level view

Related Documentation

- clear security group-vpn member ike security-associations on page 6976
- Group VPNv2 Overview on page 6649

List of Sample Output

show security group-vpn member ike security-associations on page 7012
show security group-vpn member ike security-associations detail on page 7012

Output Fields Table 572 lists the output fields for the **show security group-vpn member ike security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 572: show security group-vpn member ike security-associations Output Fields

Field Name	Field Description
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
State	State of the IKE security associations: • **DOWN**—SA has not been negotiated with the peer. • **UP**—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.

Copyright © 2016, Juniper Networks, Inc.

Table 572: show security group-vpn member ike security-associations Output Fields *(continued)*

Field Name	Field Description
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
Remote Address	IP address of the destination peer with which the local peer communicates.
IKE Peer	IP address of the destination peer with which the local peer communicates.
Exchange type	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
Authentication method	Method the server uses to authenticate the source of IKE messages: • **pre-shared-keys**—Preshared key for encryption and decryption that both participants must have before beginning tunnel negotiations.
Local	Address of the local peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Algorithms	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	
	• **Authentication**—Type of authentication algorithm used.
	  ◦ **sha-256**—Secure Hash Algorithm 256 authentication.
	  ◦ **sha-384**—Secure Hash Algorithm 384 authentication.
	• **Encryption**—Type of encryption algorithm used.
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption
	  ◦ **aes-128-cbc**—AES 128-bit encryption.

Copyright © 2016, Juniper Networks, Inc.

Table 572: show security group-vpn member ike security-associations Output Fields *(continued)*

Field Name	Field Description
Traffic statistics	• **Input bytes**—Number of bytes received. • **Output bytes**—Number of bytes transmitted. • **Input packets**—Number of packets received. • **Output packets**—Number of packets transmitted.

Sample Output

show security group-vpn member ike security-associations

```
user@host> show security group-vpn member ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

4736345 UP     70611c65603d53da  6e0888777ad10f8d  Main           106.1.1.53
```

Sample Output

show security group-vpn member ike security-associations detail

```
user@host> show security group-vpn member ike security-associations detail
IKE peer 106.1.2.53, Index 5824842, Gateway Name: group1_2
  Role: Initiator, State: UP
  Initiator cookie: fc866556b8afe4cd, Responder cookie: 1238de6b8a89de44
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 106.1.2.170:848, Remote: 106.1.2.53:848
  Lifetime: Expires in 2 seconds
  Peer ike-id: 106.1.2.53
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-2
  Traffic statistics:
   Input  bytes  :                2044
   Output bytes  :                 900
   Input  packets:                   7
   Output packets:                   7
  Flags: IKE SA is created
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn member ipsec inactive-tunnels

Syntax show security group-vpn member ipsec inactive-tunnels <brief> <detail> <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Show inactive Group VPNs.

Options **none**—Display information for all groups.

brief—(Optional) Display summary output.

detail—(Optional) Display detailed output.

group-id *group-id*—(Optional) Display information for the specified group identifier.

Required Privilege Level view

Related Documentation
- Group VPNv2 Overview on page 6649

List of Sample Output
show security group-vpn member ipsec inactive-tunnels on page 7014
show security group-vpn member ipsec inactive-tunnels detail on page 7014

Output Fields Table 573 lists the output fields for the **show security group-vpn member ipsec inactive-tunnels** command. Output fields are listed in the approximate order in which they appear.

Table 573: show security group-vpn member ipsec inactive-tunnels Output Fields

Field Name	Field Description
Server	Server on which group member is registered.
Port	UDP port number.
GId	Group identifier.
lsys	Logical system.
Reason	Reason that the tunnel is inactive: • The tunnel was cleared through the CLI. • The hard lifetime has expired. • There are too many TEKs. • There was a configuration change. • There was an SA installation error. • The TEK is stale. • The tunnel was deleted from the server.

Copyright © 2016, Juniper Networks, Inc.

Table 573: show security group-vpn member ipsec inactive-tunnels Output Fields *(continued)*

Field Name	Field Description
Virtual-system	Logical system name.
Group VPN Name	Name of the Group VPN.
Local Gateway	IP address of the local IKE gateway.
GDOI Server	IP address of the group server.
Group Id	Group identifier.
Recovery Probe	Status of the recovery probe, either enabled or disabled (default).
DF-bit	Fragmentation of IPsec traffic on the group member—clear (default), copy, or set.
Stats	Statistics for GDOI groupkey-pull and groupkey-push exchanges, server failovers, deletes received, number of times the maximum number of keys and policies were exceeded, and the number of unsupported algorithms received.
Down Reason	Reason that the tunnel is inactive: • The tunnel was cleared through the CLI. • The hard lifetime has expired. • There are too many TEKs. • There was a configuration change. • There was an SA installation error. • The TEK is stale. • The tunnel was deleted from the server. • The tunnel is not initiated.

Sample Output

show security group-vpn member ipsec inactive-tunnels

```
user@host> show security group-vpn member ipsec inactive-tunnels
  Total inactive tunnels: 1
  Server           Port  GId lsys  Reason
  200.1.1.50       848   1000 root uninitiated
```

show security group-vpn member ipsec inactive-tunnels detail

```
user@host> show security group-vpn member ipsec inactive-tunnels detail
  Virtual-system: root Group VPN Name: group1000
  Local Gateway: 200.1.1.101, GDOI Server: 200.1.1.50
  Group Id: 1000
  Recovery Probe: Disabled
  DF-bit: clear
  Stats:
      Pull Succeeded             :   0
```


Copyright © 2016, Juniper Networks, Inc.

```
    Pull Failed                :   8841
    Pull Timeout               :   7996
    Pull Aborted               :   0
    Push Succeeded             :   0
    Push Failed                :   0
    Server Failover            :   0
    Delete Received            :   0
    Exceed Maximum Keys(4)     :   0
    Exceed Maximum Policies(10):   0
    Unsupported Algo           :   0
Down Reason: uninitiated
```

show security group-vpn member ipsec security-associations

Syntax show security group-vpn member ipsec security-associations [brief | detail] [index *sa-index*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display group VPN security associations (SAs) for a group member.

Options
- none—Display information about all group VPN SAs for the group member.
- **brief**—(Optional) Display summary output.
- **detail**—(Optional) Display detailed output.
- **index** *sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

Required Privilege Level view

Related Documentation
- clear security group-vpn member ipsec security-associations on page 6977
- Group VPNv2 Overview on page 6649

List of Sample Output
show security group-vpn member ipsec security-associations on page 7018
show security group-vpn member ipsec security-associations detail on page 7018

Output Fields Table 574 lists the output fields for the **show security group-vpn member ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 574: show security group-vpn member ipsec security-associations

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.
ID	Index number of the SA. You can use this number to get additional information about the SA.
Server	IP address of the group server (remote gateway).
Port	If Network Address Translation-Traversal (NAT-T) is used, this value is 4500. Otherwise it is the standard IKE port, 500.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes • An authentication algorithm used to authenticate exchanges between the peers. Options are **sha-256** or **sha-384** • An encryption algorithm used to encrypt data traffic. Options are **aes-128**, **aes-192**, and **aes-256**.

Copyright © 2016, Juniper Networks, Inc.

Table 574: show security group-vpn member ipsec security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
GId	Group identifier.
vsys or Virtual-system	The root system.
Local Gateway	Gateway address of the local system.
GDOI Server	IP address of the group server.
Local Identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IPv4 address, fully qualified domain name, e-mail address, or distinguished name.
Remote Identity	IPv4 address of the destination peer gateway.
DF-bit	State of the don't fragment bit: set or cleared.
Policy name	Name of the applicable policy.
Direction	Direction of the security association; it can be inbound or outbound.
AUX-SPI	Value of the auxiliary security parameter index. • When the value is AH or ESP, AUX-SPI is always 0. • When the value is AH+ESP, AUX-SPI is always a positive integer.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • **Expires in seconds**—Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • **Expires in kilobytes**—Number of kilobytes left until the SA expires.
Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in seconds**—Number of seconds left until the SA expires.
Mode	Mode of the security association: • transport—Protects host-to-host connections. • tunnel—Protects connections between security gateways.

Copyright © 2016, Juniper Networks, Inc.

Table 574: show security group-vpn member ipsec security-associations *(continued)*

Field Name	Field Description
Protocol	Protocol supported. Transport mode supports Encapsulation Security Protocol (ESP).
Anti-replay service	State of the service that prevents packets from being replayed. It can be **Enabled** or **Disabled**.

Sample Output

show security group-vpn member ipsec security-associations

```
user@host> show security group-vpn member ipsec security-associations
  Total active tunnels: 2
  ID    Server           Port  Algorithm       SPI      Life:sec/kb  GId lsys
  <>49157 110.1.1.1.53     848   ESP:3des/sha1   c0792f86 114/  unlim   2000 root
  <>49156 110.1.1.1.53     848   ESP:aes-256/md5 7def169d 18/   unlim   2000 root
  <>49156 110.1.1.1.53     848   ESP:aes-256/md5 86c48448 146/  unlim   2000 root
```

Sample Output

show security group-vpn member ipsec security-associations detail

```
user@host> show security group-vpn member ipsec security-associations detail
  Virtual-system: root Group VPN Name: group2000
  Local Gateway: 110.1.1.70, GDOI Server: 110.1.1.53
  Group Id: 2000
  Routing Instance: vr1
  Recovery Probe: Enabled
  DF-bit: clear

  Stats:
      Pull Succeeded            :   3
      Pull Failed               :   0
      Pull Timeout              :   6
      Pull Aborted              :   0
      Push Succeeded            :   1773
      Push Failed               :   0
      Server Failover           :   0
      Delete Received           :   0
      Exceed Maximum Keys(4)    :   0
      Exceed Maximum Policies(10):   0
      Unsupported Algo          :   0
  Flags:
      Rekey Needed:   no

    List of policies received from server:
    Tunnel-id: 49157
      Source IP: ipv4_subnet(any:900,[0..7]=90.1.1.0/24)
      Destination IP: ipv4_subnet(any:901,[0..7]=90.2.1.0/24)

      Direction: bi-directional, SPI: c0792f86
      Protocol: ESP, Authentication: sha1, Encryption: 3des
      Hard lifetime: Expires in 81 seconds, Activated
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expired
      Mode: Tunnel, Type: Group VPN, State: installed
      Anti-replay service: D3P enabled, Window size: 3000 milliseconds
```


Copyright © 2016, Juniper Networks, Inc.

```
Direction: bi-directional, SPI: a645b381
Protocol: ESP, Authentication: sha1, Encryption: 3des
Hard lifetime: Expires in 207 seconds, Activated in 51 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 117 seconds
Mode: Tunnel, Type: Group VPN, State: installed
Anti-replay service: D3P enabled, Window size: 3000 milliseconds
```

Copyright © 2016, Juniper Networks, Inc.

show security group-vpn member ipsec statistics

Syntax show security group-vpn member ipsec statistics <index *index*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Show IPsec statistics.

Options **none**—Display information for all IPsec SAs.

index *index*—(Optional) Display detailed information about the specified SA, identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

Required Privilege Level view

Related Documentation
- Group VPNv2 Overview on page 6649

List of Sample Output show security group-vpn member ipsec statistics on page 7021

Output Fields Table 575 lists the output fields for the **show security group-vpn member ipsec statistics** command. Output fields are listed in the approximate order in which they appear.

Table 575: show security group-vpn member ipsec statistics Output Fields

Field Name	Field Description
ESP Statistics	Numbers of encrypted and decrypted bytes and encrypted and decrypted packets.
AH Statistics	Numbers of input and output bytes and input and output packets.
Errors	Numbers of AH failures, replay errors, ESP authentication failures, ESP decryption failures, bad headers, and bad trailers.
D3P Statistics	Numbers of old timestamp packets, new timestamp packets, no timestamp packets, unexpected D3P header packets, invalid type packets, invalid length packets, and invalid next header packets.
Exclude Statistics	Numbers of created and invalidated sessions.
Dynamic Policy Statistics	Numbers of created and invalidated sessions.
Fail-Open Statistics	Numbers of created and invalidated sessions.
Fail-Close Statistics	Number of dropped packets.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security group-vpn member ipsec statistics

```
user@host> show security group-vpn member ipsec statistics
ESP Statistics:
  Encrypted bytes:            54712
  Decrypted bytes:            16800
  Encrypted packets:            381
  Decrypted packets:            200
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
D3P Statistics:
  Old timestamp packets:                  0
  New timestamp packets:                  0
  No timestamp packets:                   0
  Unexpected D3P header packets:          0
  Invalid type packets:                   0
  Invalid length packets:                 0
  Invalid next header packets:            0
Exclude Statistics:
  Created sessions:                       0
  Invalidated sessions:                   0
Dynamic Policy Statistics:
  Created sessions:                     381
  Invalidated sessions:                   0
Fail-Open Statistics:
  Created sessions:                       0
  Invalidated sessions:                   0
Fail-Close Statistics:
  Dropped packets:                        0
```

Copyright © 2016, Juniper Networks, Inc.

show security group-vpn member kek security-associations

Syntax show security group-vpn member kek security-associations [brief | detail] [index *sa-index*] [*peer-ipaddress*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display group VPN security associations (SAs) for a group member.

Options

- none—Display information about all group VPN SAs for the group member.
- **brief**—(Optional) Display summary output.
- **detail**—(Optional) Display detailed output.
- **index** *sa-index*—(Optional) Display detailed information about the specified SA identified by index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- *peer-ipaddress*—(Optional) Display information about the SA with the specified peer.

Required Privilege Level view

Related Documentation

- Group VPNv2 Overview on page 6649

List of Sample Output

show security group-vpn member kek security-associations on page 7023
show security group-vpn member kek security-associations detail on page 7024

Output Fields Table 576 lists the output fields for the **show security group-vpn member kek security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 576: show security group-vpn member kek security-associations

Field Name	Field Description
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Remote Address	IP address of the destination peer with which the local peer communicates.
State	State of the KEK security associations: • **DOWN**—SA is not active. • **UP**—SA is active.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received.

Copyright © 2016, Juniper Networks, Inc.

Table 576: show security group-vpn member kek security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
GroupID	Group identifier.
KEK Peer	IP address of the destination peer with which the local peer communicates.
Role	For the member, it is always responder.
State	State of the KEK security associations, which is always up.
Authentication method	RSA is the supported authentication method.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Algorithms	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• **Sig-hash**—Type of authentication algorithm used.
	  ◦ **sha-256**–Secure Hash Algorithm 256 (sha-256) authentication.
	  ◦ **sha-384**–Secure Hash Algorithm 394 (sha-384) authentication.
	• **Encryption**—Type of encryption algorithm used.
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption
	  ◦ **aes-128-cbc**—AES 128-bit encryption.

Traffic statistics	• **Input bytes**—Number of bytes received. • **Output bytes**—Number of bytes transmitted. • **Input packets**—Number of packets received. • **Output packets**—Number of packets transmitted.
Server Info Version	Identify the latest set of information maintained in the server.
Server Replay Window	Antireplay time window value in milliseconds. 0 means antireplay is disabled.
Group Key Push sequence number	Sequence number of the KEK SA groupkey-push message. This number is incremented with every groupkey-push message.

Sample Output

show security group-vpn member kek security-associations

```
user@host> show security group-vpn member kek security-associations
```

Copyright © 2016, Juniper Networks, Inc.

```
Index   Server Address  Life:sec  Initiator cookie  Responder cookie  GroupId
5824843 106.1.2.53      166       46871e26227f08f3  f0a463a4d5c3737b  1
```

Sample Output

show security group-vpn member kek security-associations detail

```
user@host> show security group-vpn member kek security-associations detail
  Index 5824843, Group Id: 1
  Group VPN Name: group1_2
  Local Gateway: 106.1.2.170, GDOI Server: 106.1.2.53
  Initiator cookie: 46871e26227f08f3, Responder cookie: f0a463a4d5c3737b
  Lifetime: Expires in 155 seconds
  Group Key Push Sequence number: 0

  Algorithms:
   Sig-hash             : hmac-md5-96
   Encryption           : 3des-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Stats:
      Push received           :   0
      Delete received         :   0
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn member policy

Syntax show security group-vpn member policy <vpn *vpn-name*> <group-id *group-id*>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Show Group VPN policies.

Options **none**—Display information for all groups.

vpn *vpn-name*—(Optional) Display policy information for the specified group name.

group-id *group-id*—(Optional) Display policy information for the specified group identifier.

Required Privilege Level view

Related Documentation
- Group VPNv2 Overview on page 6649

List of Sample Output show security group-vpn member policy on page 7025

Output Fields Table 577 lists the output fields for the **show security group-vpn member policy** command. Output fields are listed in the approximate order in which they appear.

Table 577: show security group-vpn member policy Output Fields

Field Name	Field Description
Group VPN Name	Group name.
Group Id	Group identifier.
From-zone	From zone configured for the policy.
To-zone	To zone configured for the policy.
Tunnel-id	Tunnel identifier.
Policy type	Secure, fail-open, fail-close, or exclude.
Source	IP address, port, and protocol of the source traffic.
Destination	IP address, port, and protocol of the destination traffic.

Sample Output

show security group-vpn member policy

```
user@host> show security group-vpn member policy
Group VPN Name: group1000, Group Id: 1000
From-zone: trust_1, To-zone: untrust
```

Copyright © 2016, Juniper Networks, Inc.

```
  Tunnel-id: 63490, Policy type: Exclude
    Source      : IP <55.16.0.0 - 55.16.255.255>, Port <0 - 65535>, Protocol <17>

    Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <17>


  Tunnel-id: 49153, Policy type: Secure
    Source      : IP <188.0.0.0 - 188.255.255.255>, Port <0 - 65535>, Protocol
<0>
    Destination : IP <188.0.0.0 - 188.255.255.255>, Port <0 - 65535>, Protocol
<0>

  Tunnel-id: 49152, Policy type: Secure
    Source      : IP <55.0.0.0 - 55.255.255.255>, Port <0 - 65535>, Protocol <1>

    Destination : IP <55.0.0.0 - 55.255.255.255>, Port <0 - 65535>, Protocol <1>


  Tunnel-id: 63491, Policy type: Fail-open (Inactivated)
    Source      : IP <55.5.0.0 - 55.5.255.255>, Port <0 - 65535>, Protocol <17>
    Destination : IP <55.5.0.0 - 55.5.255.255>, Port <0 - 65535>, Protocol <17>

  Tunnel-id: 63489, Policy type: Fail-close
    Source      : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
Destination : IP <0.0.0.0 - 255.255.255.255>, Port <0 - 65535>, Protocol <0>
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server ike security-associations

Syntax show security group-vpn server ike security-associations [brief | detail] [group *group-name* | group-id *group-id*] [index *sa-index*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display IKE security associations (SAs).

Options

- none—Display all IKE SAs for all groups.
- brief—(Optional) Display summary output.
- **detail**—(Optional) Display detailed level of output.
- **group**—(Optional) Display IKE SAs for the specified group.
- **group-id**—(Optional) Display IKE SAs for the specified group.

NOTE: An IKE SA can be used by a group member to register to multiple groups. When you specify the group or group-id options to list the IKE SAs for a specified group, all existing IKE SAs that could be used to register to the group are displayed.

- **index**—(Optional) Display information for a particular SA based on the index number of the SA. To obtain the index number for a particular SA, display the list of existing SAs by using the command with no options.

Required Privilege Level view

Related Documentation

- show security group-vpn member ike security-associations on page 7010
- Group VPNv2 Overview on page 6649

List of Sample Output

show security group-vpn server ike security-associations on page 7029
show security group-vpn server ike security-associations detail on page 7030

Output Fields Table 578 lists the output fields for the **show security group-vpn server ike security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 578: show security group-vpn server ike security-associations Output Fields

Field Name	Field Description
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Remote Address	IP address of the destination peer with which the local peer communicates.

Copyright © 2016, Juniper Networks, Inc.

Table 578: show security group-vpn server ike security-associations Output Fields *(continued)*

Field Name	Field Description
State	State of the IKE security associations: • **DOWN**—SA has not been negotiated with the peer. • **UP**—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
IKE Peer	IP address of the destination peer with which the local peer communicates.
Exchange type	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected.
Authentication method	Method the server uses to authenticate the source of IKE messages: • **pre-shared-keys**—Preshared key for encryption and decryption that both participants must have before beginning tunnel negotiations. **rsa-signatures**—Digital signature, a certificate that confirms the identity of the certificate holder.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Copyright © 2016, Juniper Networks, Inc.

Table 578: show security group-vpn server ike security-associations Output Fields *(continued)*

Field Name	Field Description

Algorithms	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	• **Authentication**—Type of authentication algorithm used.
	  ◦ **sha-256**—Secure Hash Algorithm 256 authentication.
	  ◦ **sha-384**—Secure Hash Algorithm 384 authentication..
	• **Encryption**—Type of encryption algorithm used.
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption
	  ◦ **aes-128-cbc**—AES 128-bit encryption.

Traffic statistics	• **Input bytes—**Number of bytes received. • **Output bytes—**Number of bytes transmitted. • **Input packets—**Number of packets received. • **Output packets—**Number of packets transmitted.
IPSec security associations	• ***number* created:** The number of SAs created. • ***number* deleted**: The number of SAs deleted.

Phase 2 negotiations in progress	Number of Phase 2 IKE negotiations in progress and status information:
	• **Negotiation type**—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• **Message ID**—Unique identifier for a Phase 2 negotiation.
	• **Local identity**—Identity of the local Phase 2 negotiation. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)
	• **Remote identity**—Identity of the remote Phase 2 negotiation. The format is id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)
	• **Flags**—Notification to the key management process of the status of the IKE negotiation:
	  ◦ **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  ◦ **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  ◦ **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  ◦ **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.

Sample Output

show security group-vpn server ike security-associations

```
user@host> show security group-vpn server ike security-associations
  Index   State  Initiator cookie  Responder cookie  Mode         Remote Address

  738879  UP     0fa7c5fdcb74669f  8c21f5d1b533010c  Aggressive     110.1.1.120
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security group-vpn server ike security-associations detail

```
user@host> show security group-vpn server ike security-associations detail
IKE peer 110.1.1.120, Index 738879, Gateway Name: Gvpn
  Role: Responder, State: UP
  Initiator cookie: 0fa7c5fdcb74669f, Responder cookie: 8c21f5d1b533010c
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 110.1.1.50:848, Remote: 110.1.1.120:848
  Lifetime: Expires in 3541 seconds
  Peer ike-id: test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha-256-128
   Encryption            : aes-256-cbc
   Pseudo random function: hmac-sha-256
   Diffie-Hellman group  : DH-group-14
  Traffic statistics:
   Input  bytes  :                  600
   Output bytes  :                  932
   Input  packets:                    4
   Output packets:                    3
  Flags: IKE SA is created
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 0

Flags: IKE SA is created
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server ipsec security-associations

Syntax show security group-vpn server ipsec security-associations [brief | detail] [group *group-name* | group-id *group-id*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display IPsec security associations (SAs).

Options
- none—Display all IPsec SAs for all groups.
- **brief**—(Optional) Display summary output.
- **detail**—(Optional) Display detailed level of output.
- **group**—(Optional) Display IPsec SAs for the specified group.
- **group-id**—(Optional) Display IPsec SAs for the specified group.

Required Privilege Level view

Related Documentation
- show security group-vpn member ipsec security-associations on page 7016
- Group VPNv2 Overview on page 6649

List of Sample Output
show security group-vpn server ipsec security-associations on page 7032
show security group-vpn server ipsec security-associations detail on page 7032

Output Fields Table 579 lists the output fields for the **show security group-vpn server ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 579: show security group-vpn server ipsec security-associations

Field Name	Field Description
Group	Group name.
Group ID	Group identifier.
Total IPsec SAs	The total number of IPsec SAs for each group is shown.
IPsec SA	Name of the SA.
Protocol	Protocol supported. Transport mode supports Encapsulation Security Protocol (ESP).

Copyright © 2016, Juniper Networks, Inc.

Table 579: show security group-vpn server ipsec security-associations *(continued)*

Field Name	Field Description
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes • An authentication algorithm used to authenticate exchanges between the peers. Options are **sha-256** and **sha-384**. • An encryption algorithm used to encrypt data traffic. Options are **aes-128-cbc**, **aes-192-cbc**, or **aes-256-cbc**.
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI.
Lifetime	The lifetime of the SA, after which it expires, expressed in seconds.
Policy Name	Group policy associated with the IPsec SA. The source address, destination address, source port, destination port, and protocol defined for the policy are displayed.

Sample Output

show security group-vpn server ipsec security-associations

```
user@host> show security group-vpn server ipsec security-associations
  Group: group200, Group Id: 200
    Total IPsec SAs: 1
    IPsec SA          Algorithm             SPI              Lifetime
    sa1               ESP:aes-256/sha-256    55837dfe         17
    sa1               ESP:aes-256/sha1-256   760088d          137
```

Sample Output

show security group-vpn server ipsec security-associations detail

```
user@host> show security group-vpn server ipsec security-associations detail
Group: group1, Group Id: 1
Total IPsec SAs: 10
  IPsec SA: sa1
    Protocol: ESP, Authentication: sha-256, Encryption: aes-256
    Anti-replay: D3P enabled, window size 10 milliseconds
    SPI: e68c9525
    Lifetime: Expires in 66 seconds, Activated
    Policy Name: pol1
      Source: 99.0.1.0/24
      Destination: 99.0.1.0/24
      Source Port: 0
      Destination Port: 0
      Protocol: 0
  IPsec SA: sa1
    Protocol: ESP, Authentication: sha-256, Encryption: aes-256
    Anti-replay: D3P enabled, window size 10 milliseconds
    SPI: 7ee14902
    Lifetime: Expires in 276 seconds, Activated in 36 seconds
    Rekey in 186 seconds
    Policy Name: pol1
      Source: 99.0.1.0/24
      Destination: 99.0.1.0/24
      Source Port: 0
```


Copyright © 2016, Juniper Networks, Inc.

```
Destination Port: 0
Protocol: 0
```

Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server kek security-associations

Syntax show security group-vpn server kek security-associations [brief | detail] [group *group-name* | group-id *group-id* | index *sa-index*]

Release Information Command introduced in Junos OS Release 10.2.

Description Display configured server-member communications.

Options

- none—Display server-member communications configured for all groups.
- **brief**—(Optional) Display summary output.
- **detail**—(Optional) Display detailed output.
- **group**—(Optional) Display server-member communications configured for the specified group.
- **group-id**—(Optional) Display server-member communications configured for the specified group.
- **index**—(Optional) Display information for a particular SA based on the index number of the SA. To obtain the index number for a particular SA, display the list of existing SAs by using the command with no options.

Required Privilege Level view

Related Documentation

- show security group-vpn member kek security-associations on page 7022
- Group VPNv2 Overview on page 6649

List of Sample Output

show security group-vpn server kek security-associations on page 7036
show security group-vpn server kek security-associations detail on page 7036

Output Fields Table 580 lists the output fields for the **show security group-vpn server kek security-assocations** command. Output fields are listed in the approximate order in which they appear.

Table 580: show security group-vpn server kek security-associations Output Fields

Field Name	Field Description
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Remote Address	Identifier of the remote/peer. Because there could be multiple members, the remote address always contains the IP address 0.0.0.0.
State	State of the KEK security associations: • **DOWN**—SA is not active. • **UP**—SA is active.

Copyright © 2016, Juniper Networks, Inc.

Table 580: show security group-vpn server kek security-associations Output Fields *(continued)*

Field Name	Field Description
Initiator cookie	Random number generated by the server. This is used when the server needs to push data to a member, or a member needs to reply to the server.
Responder cookie	Random number generated by the server. This is used when the server needs to push data to a member, or a member needs to reply to the server.
GroupId	Group identifier.
KEK Peer	IP address of the destination peer with which the local peer communicates. For KEK SAs, it always contains 0.0.0.0 which means any IP address.
Role	For the server, it is always initiator.
Authentication method	RSA is the supported authentication method.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Algorithms	Internet Key Exchange (IKE) algorithms used to encrypt and secure exchanges between the peers during the Phase 2 process:
	• **Sig-hash**—Type of authentication algorithm used.
	  ◦ **sha-256**—Secure Hash Algorithm 256 authentication.
	  ◦ **sha-384**—Secure Hash Algorithm 384 authentication.
	• **Encryption**—Type of encryption algorithm used.
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption
	  ◦ **aes-128-cbc**—AES 128-bit encryption.

Traffic statistics	• **Input bytes—**Number of bytes received. • **Output bytes—**Number of bytes transmitted. • **Input packets—**Number of packets received. • **Output packets—**Number of packets transmitted.
Server Info Version	Identify the latest set of information maintained in the server.
The following fields are the configured **server-member-communication** options:	
Server Replay Window	Antireplay time in milliseconds. This is 0 if antireplay is disabled.
Retransmission Period	Number of seconds between a rekey transmission and the first retransmission when there is no reply from the member.

Copyright © 2016, Juniper Networks, Inc.

Table 580: show security group-vpn server kek security-associations Output Fields *(continued)*

Field Name	Field Description
Number of Retransmissions	For unicast communications, the number of times the server retransmits rekey messages to a member when there is no reply.
Lifetime Seconds	Configured lifetime, in seconds, for the KEK.
Group Key Push sequence number	Sequence number of the KEK SA groupkey-push message. This number is incremented with every groupkey-push message.

Sample Output

show security group-vpn server kek security-associations

```
user@host> show security group-vpn server kek security-associations
Index   Life:sec  Initiator cookie  Responder cookie  GroupId
  739031  18995     7e17278bf0a65975  0616de443d1beb77  200
```

Sample Output

show security group-vpn server kek security-associations detail

```
user@host> show security group-vpn server kek security-associations detail
Index 738879, Group Name: GROUP_ID-0001, Group Id: 1
Initiator cookie: 114e4a214891e42f, Responder cookie: 4b2848d14372e5bd
Authentication method: RSA
Lifetime: Expires in 4186 seconds, Activated
Rekey in 3614 seconds
  Algorithms:
   Sig-hash             : sha256
   Encryption           : aes256-cbc
  Traffic statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
  Server Member Communication: Unicast
  Retransmission Period: 10, Number of Retransmissions: 2
  Group Key Push sequence number: 0

PUSH negotiations in progress: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server registered-members

Syntax show security group-vpn server registered-members <group *group-name*> <group-id *group-id*> <detail>

Release Information Command introduced in Junos OS Release 10.2.

Description Display currently registered group members.

Options
- none—Display all group members for all groups.
- brief—(Optional) Display summary output.
- detail—(Optional) Display detailed output.
- **group**—(Optional) Display group members for the specified group.
- **group-id**—(Optional) Display group members for the specified group.

Required Privilege Level view

Related Documentation
- clear security group-vpn server on page 6980
- Group VPNv2 Overview on page 6649

List of Sample Output
show security group-vpn server registered-members on page 7038
show security group-vpn server registered-members detail on page 7038

Output Fields Table 581 lists the output fields for the **show security group-vpn server registered-members**command. Output fields are listed in the approximate order in which they appear.

Table 581: show security group—vpn server registered-members Output Fields

Field Name	Field Description
Group	Group name.
Group Id	Group identifier.
Member Gateway	IP address of the gateway for the group member.
Member IP	IP address of the group member.
Last Update	The last time that members registered or sent acknowledgements to the server.
Vsys	The root system.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security group-vpn server registered-members

```
user@host> show security group-vpn server registered-members
  Group: group200, Group Id: 200
    Total number of registered members: 1
    Member Gateway                  Member IP       Last Update
Vsys
    gvpn_simpleman                  110.1.1.100     Fri Dec 20 2013 07:27:33
root
```

Sample Output

show security group-vpn server registered-members detail

```
user@host> show security group-vpn server registered-members detail
Group: group1, Group Id: 1
      Total number of registered members: 1

  Member gateway: gateway_group1_1, Member IP: 100.1.1.2, Vsys: root
  Last Update: Fri May 16 2014 03:37:17
  Stats:
      Pull Succeeded                : 321
      Pull Failed                   : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server server-cluster

Syntax **show security group-vpn server server-cluster <brief> <detail> <group *group-name*> <group-id *group-id*> <peer-gateway *gateway-name*>**

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Show information about servers in the Group VPNv2 server cluster.

Options

none—Display Group VPNv2 server cluster information for all groups.

brief—(Optional) Display summary output.

detail—(Optional) Display detailed output, including information about exchanges with peer servers in the cluster.

group *group-name*—(Optional) Display Group VPNv2 server cluster information for the specified group name.

group-id *group-id*—(Optional) Display Group VPNv2 server cluster information for the specified group identifier.

peer-gateway *gateway-name*—(Optional) Display Group VPNv2 server cluster information for the specified peer.

Required Privilege Level view

Related Documentation

- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699

List of Sample Output

show security group-vpn server server-cluster on page 7040
show security group-vpn server server-cluster detail on page 7040

Output Fields Table 582 lists the output fields for the **show security group-vpn server server-cluster** command. Output fields are listed in the approximate order in which they appear.

Table 582: show security group-vpn server server-cluster Output Fields

Field Name	Field Description
Group	Group name.
Group Id	Group identifier.
Role	Role of this server in the Group VPNv2 server cluster.

Copyright © 2016, Juniper Networks, Inc.

Table 582: show security group-vpn server server-cluster Output Fields *(continued)*

Field Name	Field Description
Version Number	32-bit version number included in **cluster-update** exchanges and DPD probes to support anti-replay. The first **cluster-update** message sent from the root-server has version number 1. Subsequent **cluster-update** messages increment the version number by one. (Retransmit messages do not increment the version number.) Upon receipt of a **cluster-update** message, the sub-server validates the received version number. The received version number must be greater than the version number in the last received message, otherwise the message is discarded. The sub-server responds to a **cluster-update** message with an ACK message that contains the same version number as the received message. Upon receipt of the ACK message, the root-server checks that the version number is the same as in the message it sent. If the version number is valid, the exchange is considered successful. If the version number is not valid, the original message is retransmitted or the exchange is considered failed.
Peer Gateway	Name of the peer server in the Group VPNv2 server cluster.
Peer IP	IP address of the remote peer server in the Group VPNv2 server cluster.
Role	Role of the peer server in the Group VPNv2 server cluster.
Status	Status of the peer server in the Group VPNv2 server cluster.

Sample Output

show security group-vpn server server-cluster

```
user@host> show security group-vpn server server-cluster
Group: group200, Group Id: 200
Role: Root-server, Version Number: 1,
  Peer Gateway                    Peer IP         Role                 Status

  sub_server1                     110.1.1.112     Sub-server           Active

  sub_server2                     110.1.1.113     Sub-server           Active
```

show security group-vpn server server-cluster detail

```
user@host> show security group-vpn server server-cluster detail
GGroup: group200, Group Id: 200
Role: Root-server, Version Number: 1,

Peer gateway: sub_server1,
  Peer IP: 110.1.1.112, Local IP: 110.1.1.111, VR: vr1,
  Role: Sub-server, Status: Active,
  CLUSTER-INIT send:                0
  CLUSTER-INIT recv:                1
  CLUSTER-INIT success:             1
  CLUSTER-INIT fail:                0
  CLUSTER-INIT dup:                 0
  CLUSTER-INIT abort:               0
  CLUSTER-INIT timeout:             0
  CLUSTER-UPDATE send:              1
  CLUSTER-UPDATE recv:              0
```


Copyright © 2016, Juniper Networks, Inc.

```
  CLUSTER-UPDATE success:            1
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          5
  DPD send fail:                     0
  DPD ACK recv:                      5
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0

Peer gateway: sub_server2,
  Peer IP: 110.1.1.113, Local IP: 110.1.1.111, VR: default,
  Role: Sub-server, Status: Active,
  CLUSTER-INIT send:                 0
  CLUSTER-INIT recv:                 1
  CLUSTER-INIT success:              1
  CLUSTER-INIT fail:                 0
  CLUSTER-INIT dup:                  0
  CLUSTER-INIT abort:                0
  CLUSTER-INIT timeout:              0
  CLUSTER-UPDATE send:               1
  CLUSTER-UPDATE recv:               0
  CLUSTER-UPDATE success:            1
  CLUSTER-UPDATE fail:               0
  CLUSTER-UPDATE abort:              0
  CLUSTER-UPDATE timeout:            0
  CLUSTER-UPDATE pending:            0
  CLUSTER-UPDATE max retry reached:  0
  DPD send:                          6
  DPD send fail:                     0
  DPD ACK recv:                      6
  DPD ACK invalid seqno:             0
  IPsec SA policy mismatch:          0
  IPsec SA proposal mismatch:        0
  KEK SA proposal mismatch:          0
```

Copyright © 2016, Juniper Networks, Inc.

show security group-vpn server statistics

Syntax **show security group-vpn server statistics <group *group-name*> <group-id *group-id*>**

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Show Group VPNv2 server statistics.

Options **none**—Display Group VPNv2 server statistics for all groups.

group *group-name*—(Optional) Display Group VPNv2 server statistics for the specified group name.

group-id *group-id*—(Optional) Display Group VPNv2 server statistics for the specified group identifier.

Required Privilege Level view

Related Documentation
- Group VPNv2 Overview on page 6649
- Understanding Group VPNv2 Server Clusters on page 6699

List of Sample Output show security group-vpn server statistics on page 7042

Output Fields Table 583 lists the output fields for the **show security group-vpn server statistics** command. Output fields are listed in the approximate order in which they appear.

Table 583: show security group-vpn server statistics Output Fields

Field Name	Field Description
Group	Group name.
Group Id	Group identifier.
Stats	Server events and number of occurrences.

Sample Output

show security group-vpn server statistics

```
user@host> show security group-vpn server statistics
Group: group1, Group Id: 1
  Stats:
      Pull Succeeded                : 321
      Pull Failed                   : 0
      Pull Exceed Member Threshold  : 0
      Push Sent                     : 0
      Push Acknowledged             : 0
      Push Unacknowledged           : 0
```


Copyright © 2016, Juniper Networks, Inc.

show security ike active-peer

Syntax show security ike active-peer

Release Information Command introduced in Junos OS Release 10.4. Support to display dead peer detection (DPD) statistics added in Junos OS Release 12.3X48-D10.

Description Display the list of connected active users with details about the peer addresses and ports they are using.

Required Privilege Level view

Related Documentation
- show security ike security-associations on page 3673
- show security ipsec security-associations on page 3681

List of Sample Output
show security ike active-peer on page 7043
show security ike active-peer detail on page 7043

Sample Output

show security ike active-peer

```
user@host> show security ike active-peer

Remote Address   Port     Peer IKE-ID        XAUTH username   Assigned IP
172.27.6.136     8034     tleungjtac@650a     tleung          10.123.80.225
```

show security ike active-peer detail

```
user@host> show security ike active-peer detail
Peer address: 31.0.0.6, Port: 500,
Peer IKE-ID: C=US, ST=California, L=Sunnyvale, O=Example, OU=sales, CN=SPOKE9061
XAUTH username: not available
Assigned network attributes:
IP Address: 0.0.0.0 ,   netmask         : 0.0.0.0
DNS Address     : 0.0.0.0 ,   DNS2 Address    : 0.0.0.0
WINS Address    : 0.0.0ß.0 ,   WINS2 Address   : 0.0.0.0

Previous Peer address   : 0.0.0.0, Port               : 0
Active IKE SA indexes   : 75203629
IKE SA negotiated       : 1
IPSec tunnels active    : 1, IPSec Tunnel IDs   : 68157442

DPD Config Info     :  Mode: always-send    Interval: 60   Threshold: 5
p1sa_index:75203629
DPD Statistics      : DPD-flags: REMOTE_ACCESS
DPD Statistics      : DPD TTL                   :     0       DPD seq-no
      :       0
DPD Statistics      : DPD Req Sent              :     0       DPD Resp Rcvd
      :       0
```

Copyright © 2016, Juniper Networks, Inc.

show security ike pre-shared-key

Syntax **show security ike pre-shared key**
<master-key ***master-key*** **>**
<user-id ***user-id*** **>**

Release Information Command introduced in Junos OS Release 8.5.

Description Display the Internet Key Exchange (IKE) preshared key used by the Virtual Private network (VPN) gateway to authenticate the remote access user.

Options
- **master-key** ***master-key*** —(Optional) Master preshared key.
- **user-id** ***user-id*** —(Optional) IKE user ID value.

Required Privilege Level view

Related Documentation
- pre-shared-key (Security IKE Policy) on page 6927

List of Sample Output show security ike pre-shared-key on page 7044

Sample Output

show security ike pre-shared-key

```
user@host> show security ike pre-shared-key user-id a@example.net master-key example
Preshared Key:3b33ec3631a561ec5a710f5d02f208033b108bb4
```


Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations

Syntax

show security ike security-associations
peer-address
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number*
sa-type shortcut <detail>

Release Information

Command introduced in Junos OS Release 8.5 . Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for Auto Discovery VPN added in Junos OS Release 12.3X48-D10.

Description

Display information about Internet Key Exchange security associations (IKE SAs).

Options

- none—Display standard information about existing IKE SAs, including index numbers.
- ***peer-address***—(Optional) Display details about a particular SA based on the IPv4 or IPv6 address of the destination peer. This option and **index** provide the same level of output.
- **brief**—(Optional) Display standard information about all existing IKE SAs. (Default)
- **detail**—(Optional) Display detailed information about all existing IKE SAs.
- **family**—(Optional) Display IKE SAs by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc *slot-number***—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index *SA-index-number***—(Optional) Display information for a particular SA based on the index number of the SA. For a particular SA, display the list of existing SAs by using the command with no options. This option and ***peer-address*** provide the same level of output.
- **kmd-instance** —(Optional) Display information about existing IKE SAs in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic *slot-number*** —(Optional) Display information about existing IKE SAs in this PIC slot. This option is used to filter the output.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

List of Sample Output
show security ike security-associations (IPv4) on page 7048
show security ike security-associations (IPv6) on page 7048
show security ike security-associations detail (Branch SRX Series Devices) on page 7049
show security ike security-associations detail (High-End SRX Series Devices) on page 7049
show security ike security-associations family inet6 on page 7049
show security ike security-associations index 8 detail on page 7050
show security ike security-associations 1.1.1.2 on page 7050
show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 7051
show security ike security-associations detail (ADVPN Suggester, Static Tunnel) on page 7051
show security ike security-associations detail (ADVPN Partner, Static Tunnel) on page 7051
show security ike security-associations detail (ADVPN Partner, Shortcut) on page 7051
show security ike security-associations sa-type shortcut (ADVPN) on page 7052
show security ike security-associations sa-type shortcut detail (ADVPN) on page 7052

Output Fields Table 362 lists the output fields for the **show security ike security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 584: show security ike security-associations Output Fields

Field Name	Field Description
IKE Peer or Remote Address	IP address of the destination peer with which the local peer communicates.
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Gateway Name	Name of the IKE gateway.
Location	• **FPC**—Flexible PIC Concentrator (FPC) slot number. • **PIC**—PIC slot number. • **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances are running on each SPU, and any particular IKE negotiation is carried out by a single KMD instance.
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.
State	State of the IKE SAs: • **DOWN**—SA has not been negotiated with the peer. • **UP**—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.

Copyright © 2016, Juniper Networks, Inc.

Table 584: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode or Exchange type	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between one another. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are: • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected. NOTE: IKEv2 protocol does not use the mode configuration for negotiation. Therefore, mode displays the version number of the security association.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.

Algorithms	IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	
	• **Authentication**—Type of authentication algorithm used:
	  ◦ **sha1**—Secure Hash Algorithm 1 authentication.
	  ◦ **md5**—MD5 authentication.
	• **Encryption**—Type of encryption algorithm used:
	  ◦ **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  ◦ **aes-192-cbc**— AES192-bit encryption.
	  ◦ **aes-128-cbc**—AES 128-bit encryption.
	  ◦ **3des-cbc**—3 Data Encryption Standard (DES) encryption.
	  ◦ **des-cbc**—DES encryption.

Diffie-Hellman group	Specifies the IKE Diffie-Hellman group.
Traffic statistics	• **Input bytes**—Number of bytes received. • **Output bytes**—Number of bytes transmitted. • **Input packets**—Number of packets received. • **Output packets**—Number of packets transmitted.

Copyright © 2016, Juniper Networks, Inc.

Table 584: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
Flags	Notification to the key management process of the status of the IKE negotiation: • **caller notification sent**—Caller program notified about the completion of the IKE negotiation. • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.
IPSec security associations	• ***number* created:** The number of SAs created. • ***number* deleted**: The number of SAs deleted.

Phase 2 negotiations in progress	Number of Phase 2 IKE negotiations in progress and status information:
	• **Negotiation type**—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• **Message ID**—Unique identifier for a Phase 2 negotiation.
	• **Local identity**—Identity of the local Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• **Remote identity**—Identity of the remote Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• **Flags**—Notification to the key management process of the status of the IKE negotiation:
	  ◦ **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  ◦ **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  ◦ **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  ◦ **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.

Sample Output

show security ike security-associations (IPv4)

```
user@host> show security ike security-associations
Index  Remote Address    State  Initiator cookie     Responder cookie     Mode
8  1.1.1.2    UP  3a895f8a9f620198     9040753e66d700bb     Main
Index  Remote Address    State  Initiator cookie     Responder cookie     Mode
9  1.2.1.3     UP  5ba96hfa9f65067   1  70890755b65b80b   d  Main
```

show security ike security-associations (IPv6)

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode           Remote Address

5       UP     e48efd6a444853cf  0d09c59aafb720be  Aggressive     1212::1112
```


Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations detail (Branch SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 25.191.134.245, Index 2577565, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: b869b3424513340a, Responder cookie: 4cb3488cb19397c3
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 25.191.134.241:500, Remote: 25.191.134.245:500
  Lifetime: Expires in 169 seconds
  Peer ike-id: 25.191.134.245
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1012
   Output bytes  :                 1196
   Input  packets:                    4
   Output packets:                    5
  Flags: IKE SA is created
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 25.191.134.241:500, Remote: 25.191.134.245:500
    Local identity: 25.191.134.241
    Remote identity: 25.191.134.245
    Flags: IKE SA is created
```

show security ike security-associations detail (High-End SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 1.1.1.2, Index 914039858, Gateway Name: tropic
  Location: FPC 3, PIC 1, KMD-Instance 3
  Role: Initiator, State: UP
  Initiator cookie: 219a697652bdde37, Responder cookie: b49c30b229d36bcd
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 1.1.1.1:500, Remote: 1.1.1.2:500
  Lifetime: Expires in 26297 seconds
  Peer ike-id: 1.1.1.2
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  IPSec security associations: 0 created, 0 deleted
  Phase 2 negotiations in progress: 1
```

show security ike security-associations family inet6

```
user@host> show security ike security-associations family inet6
```

```
IKE peer 1212::1112, Index 5, Gateway Name: tropic
Role: Initiator, State: UP
Initiator cookie: e48efd6a444853cf, Responder cookie: 0d09c59aafb720be
Exchange type: Aggressive, Authentication method: Pre-shared-keys
Local: 1212::1111:500, Remote: 1212::1112:500
Lifetime: Expires in 19518 seconds
Peer ike-id: not valid
Xauth assigned IP: 0.0.0.0
Algorithms:
 Authentication        : sha1
 Encryption            : 3des-cbc
 Pseudo random function: hmac-sha1
 Diffie-Hellman group  : DH-group-5
Traffic statistics:
 Input  bytes  :                 1568
 Output bytes  :                 2748
 Input  packets:                    6
 Output packets:                   23
Flags: Caller notification sent
IPSec security associations: 5 created, 0 deleted
Phase 2 negotiations in progress: 1

  Negotiation type: Quick mode, Role: Initiator, Message ID: 2900338624
  Local: 1212::1111:500, Remote: 1212::1112:500
  Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Flags: Caller notification sent, Waiting for done
```

show security ike security-associations index 8 detail

```
user@host> show security ike security-associations index 8 detail
IKE peer 1.1.1.2, Index 8, Gateway Name: tropic
 Role: Responder, State:UP
 Initiator cookie: 3a895f8a9f620198, Responder cookie: 9040753e66d700bb
 Exchange type; main, Authentication method: Pre-shared-keys
 Local: 1.1.1.1:500, Remote: 1.1.1.2:500
 Lifetime: Expired in 381 seconds
 Algorithms:
  Authentication:       md5
  Encryption:      3des-cbc
  Pseudo random function      hmac-md5
     Diffie-Hellman group  : DH-group-5
    Traffic statistics:
  Input bytes:       11268
  Output bytes:        6940
  Input packets:        57
  Output packets:        57
 Flags: Caller notification sent
 IPsec security associations: 0 created, 0 deleted
 Phase 2 negotiations in progress: 1

  Negotiation type: Quick mode, Role: Responder, Message ID: 1765792815
  Local: 1.1.1.1:500, Remote: 1.1.1.2:500
  Local identity: No Id
  Remote identity: No Id
  Flags: Caller notification sent, Waiting for remove
```

show security ike security-associations 1.1.1.2

```
user@host> show security ike security-associations 1.1.1.2
```


Copyright © 2016, Juniper Networks, Inc.

```
Index     State  Initiator cookie  Responder cookie  Mode Remote Address
   8        UP     3a895f8a9f620198  9040753e66d700bb  Main 1.1.1.2
```

show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ike security-associations fpc 6 pic 1 kmd-instance all
Index   Remote Address  State  Initiator cookie  Responder cookie  Mode

1728053250 1.1.1.2      UP     fc959afd1070d10b  bdeb7e8c1ea99483  Main
```

show security ike security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.105, Index 13563297, Gateway Name: zth_hub_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent            :  12
     Suggestion response accepted:  12
     Suggestion response declined:   0
  Role: Responder, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.154:500, Remote: 23.0.0.105:500
  Lifetime: Expires in 26429 seconds
  Peer ike-id: DC=example, CN=test, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.154, Index 4980720, Gateway Name: zth_spoke_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:  12
     Suggestions accepted:  12
     Suggestions declined:   0
  Role: Initiator, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.105:500, Remote: 23.0.0.154:500
  Lifetime: Expires in 26252 seconds
  Peer ike-id: DC=example, CN=vsrxvpn01, OU=SBU, O=Example, L=Sunnyvale, ST=CA,
C=US
```

show security ike security-associations detail (ADVPN Partner, Shortcut)

```
user@host> show security ike security-associations detail
IKE peer 23.0.0.106, Index 4980737, Gateway Name:
GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173323
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: e1ed0c655929debc, Responder cookie: 437de6ed784ba63e
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 23.0.0.105:500, Remote: 23.0.0.106:500
  Lifetime: Expires in 28796 seconds
  Peer ike-id: DC=example, CN=paulyd, L=Sunnyvale, ST=CA, C=US
```

Copyright © 2016, Juniper Networks, Inc.

show security ike security-associations sa-type shortcut (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut
Initiator cookie  Responder cookie  Mode           Remote Address
4980742Index   State  UP     b56fbe694eaee5b6  064dbccbfa3b2aab  IKEv2
23.0.0.106
```

show security ike security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ike security-associations sa-type shortcut detail
IKE peer 23.0.0.106, Index 4980742, Gateway Name:
GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173327
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Role: Partner, Peer Role: Partner
  Role: Responder, State: UP
```


Copyright © 2016, Juniper Networks, Inc.

show security ike tunnel-map

Syntax **show security ike tunnel-map <brief> <fpc *slot-number*> <kmd-instance (all | *kmd-instance-name*> <pic *slot-number*> <summary>**

Release Information Command introduced in Junos OS Release 12.1X44-D10.

Description Display the tunnel mapping on different Services Processing Units (SPUs) for site-to-site and manual VPNs. You can insert an SPC on a device in a chassis cluster without disrupting traffic on the existing VPN tunnels. After inserting the SPC, you can view the tunnel mapping using this command.

Options **brief**—Display standard information about all existing IKE SAs. This is the default.

fpc *slot-number*—Display information about existing IKE SAs in the specified Flexible PIC Concentrator (FPC) slot.

kmd-instance (all | *kmd-instance-name*)—Display information about existing IKE SAs in the KMD key management process. You can specify one of the following options:

- all—All KMD instances running on the SPU.
- *kmd-instance-name*—Name of the KMD instance running on the SPU.

pic *slot-number*—Display information about existing IKE SAs in the specified PIC slot.

summary—Display the tunnel-mapping load on each SPU. The load is the number of times an SPU has been chosen as an anchor SPU. For site-to-site VPNs, the load should be equal to the number of gateways mapped to an SPU.

Required Privilege Level view

Related Documentation

- Understanding VPN Support for Inserting Services Processing Cards on page 6068

List of Sample Output

show security ike tunnel-map on page 7054
show security ike tunnel-map brief on page 7054
show security ike tunnel-map fpc 1 pic 0 on page 7054
show security ike tunnel-map kmd-instance kmd1 on page 7054
show security ike tunnel-map kmd-instance all on page 7054
show security ike tunnel-map summary on page 7055

Output Fields Table 585 lists the output fields for the **show security ike tunnel-map** command. Output fields are listed in the approximate order in which they appear.

Table 585: show security ike tunnel-map Output Fields

Field Name	Field Descripton
Gateway ID	Gateway identifier

Copyright © 2016, Juniper Networks, Inc.

Table 585: show security ike tunnel-map Output Fields *(continued)*

Field Name	Field Descripton
Gateway Name	Name of the IKE gateway
FPC	FPC slot number
PIC	PIC slot number
IKED Instance	IKE process instance identifier
SPU Load	Number of times an SPU has been chosen as an anchor SPU

Sample Output

show security ike tunnel-map

```
user@host> show security ike tunnel-map
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        ike_gw1      4    0       1
    3        ike_gw2      7    0       1
    4        ike_gw3      7    0       2
    5        ike_gw4      4    0       2
```

show security ike tunnel-map brief

```
user@host> show security ike tunnel-map brief
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01        1    0       1
    3        LAN_1        1    0       2
    4        LAN_2        1    0       1
    5        LAN_3        1    0       2
    6        LAN_4        1    0       1
```

show security ike tunnel-map fpc 1 pic 0

```
user@host> run show security ike tunnel-map fpc 1 pic 0
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01        1    0       1
    3        LAN_1        1    0       2
    4        LAN_2        1    0       1
    5        LAN_3        1    0       2
    6        LAN_4        1    0       1
```

show security ike tunnel-map kmd-instance kmd1

```
user@host> show security ike tunnel-map kmd-instance kmd1
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01        1    0       1
    4        LAN_2        1    0       1
    6        LAN_4        1    0       1
```

show security ike tunnel-map kmd-instance all

```
user@host> show security ike tunnel-map kmd-instance all
Gateway ID   Gateway Name FPC  PIC  IKED Instance
    2        gw-01        1    0       1
```


Copyright © 2016, Juniper Networks, Inc.

```
3          LAN_1          1     0        2
4          LAN_2          1     0        1
5          LAN_3          1     0        2
6          LAN_4          1     0        1
```

show security ike tunnel-map summary

```
user@host> show security ike tunnel-map summary
FPC  PIC  SPU Load
1    0      5
```

show security ipsec control-plane-security-associations

Syntax show security ipsec control-plane-security-associations
<brief | detail>
<sa-name *sa-name*>

Release Information Command introduced in Junos OS Release 12.1X46-D20.

Description Display information about manual IPsec security associations (SAs) applied to OSPF or OSPFv3 interfaces or virtual links.

Options
- **brief** | **detail**—(Optional) Display the specified level of output.
- **sa-name** ***sa-name***—Name of the manual SA.

Required Privilege Level view

Related Documentation
- *Understanding OSPF and OSPFv3 Authentication on SRX Series Devices*

List of Sample Output
show security ipsec control-plane-security-associations on page 7056
show security ipsec control-plane-security-associations sa-name on page 7057
show security ipsec control-plane-security-associations detail on page 7057

Output Fields Table 586 lists the output fields for the **show security ipsec control-plane-security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 586: show security ipsec control-plane-security-associations Output Fields

Field Name	Field Description
Name	Name of the SA.
Algorithm	IPsec protocol followed by encryption algorithm and authentication algorithm.
SPI	SPI value.
Total active security-associations	Total number of active manual SAs for application to OSPF or OSPFv3 interfaces or virtual links.

Sample Output

show security ipsec control-plane-security-associations

```
user@host> show security ipsec control-plane-security-associations
Name      Algorithm       SPI
test_sa   ESP:3des/md5   3e8
test_sa   ESP:3des/md5   3e8
test_sa2  ESP:3des/sha1  7d1
test_sa2  ESP:3des/sha1  7d1
Total active security-associations: 2
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec control-plane-security-associations sa-name

```
user@host> show security ipsec control-plane-security-associations sa-name test_sa
Name      Algorithm       SPI
test_sa   ESP:3des/md5    3e8
test_sa   ESP:3des/md5    3e8
Total active security-associations: 1
```

show security ipsec control-plane-security-associations detail

```
user@host> show security ipsec control-plane-security-associations detail
Direction: inbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 1,

Direction: outbound, SA Name: test_sa,
Protocol: ESP:, Authentication: md5
SPI: 3e8, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 2,

Direction: inbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 3,

Direction: outbound, SA Name: test_sa2,
Protocol: ESP:, Authentication: sha1
SPI: 7d1, AUX-SPI: 0,
Mode: transport, Type: manual,
ID: 4,
```

Copyright © 2016, Juniper Networks, Inc.

show security ipsec inactive-tunnels

Syntax **show security ipsec inactive-tunnels**
brief | detail
family (inet | inet6)
fpc *slot-number*
index *index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number*
sa-type shortcut
vpn-name *vpn-name*

Release Information Command introduced in Junos OS Release 11.4R3. Support for Auto Discovery VPN added in Junos OS Release 12.3X48-D10.

Description Display security information about the inactive tunnel.

Options

- none—Display information about all inactive tunnels.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **family**—(Optional) Display the inactive tunnel by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc** ***slot-number***—(Optional) Display information about inactive tunnels in the Flexible PIC Concentrator (FPC) slot.
- **index** ***index-number***—(Optional) Display detailed information about the specified inactive tunnel identified by this index number. For a list of all inactive tunnels with their index numbers, use the command with no options.
- **kmd-instance** —(Optional) Display information about inactive tunnels in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number***—Display information about inactive tunnels in the PIC slot.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.
- **vpn-name** ***vpn-name***—(Optional) Name of the VPN.

NOTE: The **fpc** ***slot-number***, **kmd-instance (all |** ***kmd-instance-name*****)**, and **pic** ***slot-number*** parameters apply to SRX5600 and SRX5800 devices only.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- show security ipsec security-associations on page 3681

List of Sample Output
show security ipsec inactive-tunnels on page 7060
show security ipsec inactive-tunnels index 131073 on page 7060
show security ipsec inactive-tunnels sa-type shortcut on page 7060

Output Fields Table 587 lists the output fields for the **show security ipsec inactive-tunnels** command. Output fields are listed in the approximate order in which they appear.

Table 587: show security ipsec inactive-tunnels Output Fields

Field Name	Field Description
Total inactive tunnels	Total number of inactive IPsec tunnels.
Total inactive tunnels which establish immediately	Total number of inactive IPsec tunnels that can establish a session immediately.
ID	Identification number of the inactive tunnel. You can use this number to get more information about the inactive tunnel.
Gateway	IP address of the remote gateway.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Def-Del#	Number of deferred deletions of a dial-up IPsec VPN.
Virtual system	Virtual system to which the VPN belongs.
VPN name	Name of the IPsec VPN.
Local gateway	Gateway address of the local system.
Remote gateway	Gateway address of the remote system.
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote identity	IP address of the destination peer gateway.
Version	Version of IKE.
DF-bit	State of the don't fragment bit: **set** or **clear**.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Policy-name	Name of the applicable policy.

Copyright © 2016, Juniper Networks, Inc.

Table 587: show security ipsec inactive-tunnels Output Fields *(continued)*

Field Name	Field Description
Tunnel Down Reason	Reason for which the tunnel is inactive.
Tunnel events	Tunnel event and the number of times the event has occurred. See "Tunnel Events" on page 6797 for descriptions of tunnel events and the action you can take.

Sample Output

show security ipsec inactive-tunnels

```
user@host> show security ipsec inactive-tunnels
Total inactive tunnels: 1
  Total inactive tunnels with establish immediately: 0
  ID     Gateway    Port  Tunnel down reason
  131073 100.1.1.2  500   Phase1 proposal mismatch detected
```

show security ipsec inactive-tunnels index 131073

```
user@host> show security ipsec inactive-tunnels index 131073
ID: 131073 Virtual-system: root, VPN Name: vpn1
  Local Gateway: 100.1.1.100, Remote Gateway: 100.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.0
  Port: 500, Nego#: 2, Fail#: 0, Def-Del#: 0 Flag: 600a29
  Tunnel events:
    Wed Jul 16 2014 06:18:02 +0800: User cleared IPSec SA from CLI (1 times)
    Wed Jul 16 2014 06:17:58 +0800: IPSec SA negotiation successfully completed
(1 times)
    Wed Jul 16 2014 06:17:54 +0800: User cleared IPSec SA from CLI (1 times)
    Wed Jul 16 2014 06:16:58 +0800: IPSec SA negotiation successfully completed
(1 times)
   Wed Jul 16 2014 06:16:58 +0800: Bind interface's address received. Information
 updated (1 times)
    Wed Jul 16 2014 06:16:58 +0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Wed Jul 16 2014 06:16:58 +0800: External interface's address received.
Information updated (1 times)
    Wed Jul 16 2014 06:16:58 +0800: Bind interface's zone received. Information
updated (1 times)
    Wed Jul 16 2014 06:16:58 +0800: IKE SA negotiation successfully completed (1
 times)
```

show security ipsec inactive-tunnels sa-type shortcut

```
user@host> show security ipsec inactive-tunnels sa-type shortcut
  Total inactive tunnels: 1
  Total inactive tunnels with establish immediately: 0
  ID     Port  Nego#  Fail#  Flag      Gateway          Tunnel Down Reason
  268173322 500 0     0      40608aa9  23.0.0.105       Cleared via CLI
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec next-hop-tunnels

Syntax **show security ipsec next-hop-tunnels**
< interface-name ***interface-name*** **>**

Release Information Command introduced in Junos OS Release 8.5.

Description Display security information about the secure tunnel interface.

Options
- none—Display information about all secure tunnel interface.
- **interface-name** ***interface-name*** —(Optional) Name of the secure tunnel logical interface.

Required Privilege Level view

Related Documentation
- show security ipsec security-associations on page 3681

List of Sample Output show security ipsec next-hop-tunnels on page 7061

Output Fields Table 588 lists the output fields for the **show security ipsec next-hop-tunnels** command. Output fields are listed in the approximate order in which they appear.

Table 588: show security ipsec next-hop-tunnels Output Fields

Field Name	Field Description
Next-hop gateway	IP address of the next gateway.
Interface	Name of the secure tunnel logical interface.
IPsec VPN name	Name of the IPsec VPN tunnel.
Flag	• **Static**—IP address manually configured. • **Auto**—IP address obtained from the remote peer automatically.

Sample Output

show security ipsec next-hop-tunnels

```
user@host> show security ipsec next-hop-tunnels
Next-hop gateway  interface   IPsec VPN name                 Flag
11.1.1.2          st0.0       autokey                        Static
11.1.1.3          st0.0       pbd-4-6                        Auto
```

Copyright © 2016, Juniper Networks, Inc.

show security ipsec security-associations

Syntax **show security ipsec security-associations**
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number>*
sa-type shortcut
vpn-name *vpn-name* <traffic-selector *traffic-selector-name*>

Release Information Command introduced in Junos OS Release 8.5. Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for the **vpn-name** option added in Junos OS Release 11.4R3. Support for the **traffic-selector** option and traffic selector field added in Junos OS Release 12.1X46-D10. Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10.

Description Display information about the IPsec security associations (SAs).

Options

- none—Display information about all SAs.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **family**—(Optional) Display SAs by family. This option is used to filter the output.
  - **inet**—IPv4 address family.
  - **inet6**—IPv6 address family.
- **fpc** ***slot-number***—(Optional) Display information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index** ***SA-index-number***—(Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.
- **kmd-instance**—(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number***—(Optional) Display information about existing IPsec SAs in this PIC slot. This option is used to filter the output.
- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.
- **vpn-name** ***vpn-name***—Name of the VPN. If configured, **traffic-selector** ***traffic-selector-name*** can optionally be specified.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- clear security ipsec security-associations on page 6986
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 3335

List of Sample Output
show security ipsec security-associations (IPv4) on page 7066
show security ipsec security-associations (IPv6) on page 7066
show security ipsec security-associations index 131073 on page 7066
show security ipsec security-associations brief on page 7067
show security ipsec security-associations detail on page 7067
show security ipsec security-associations family inet6 on page 7068
show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 7068
show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel) on page 7068
show security ike sa index 222075191 detail on page 7069
show security ipsec security-associations detail (ADVPN Partner, Static Tunnel) on page 7070
show security ike sa index 788674 detail on page 7071
show security ipsec security-associations sa-type shortcut (ADVPN) on page 7071
show security ipsec security-associations sa-type shortcut detail (ADVPN) on page 7071
show security ipsec security-associations family inet detail on page 7072

Output Fields Table 363 lists the output fields for the **show security ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 589: show security ipsec security-associations

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.
ID	Index number of the SA. You can use this number to get additional information about the SA.
VPN name	IPsec name for VPN.
Gateway	IP address of the remote gateway.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are **hmac-md5-95**, **hmac-sha1-96,** or **ESP**. • An encryption algorithm used to encrypt data traffic. Options are **3des-cbc, aes-128-cbc, aes-192-cbc**, **aes-256-cbc**, or **des-cbc.**

Copyright © 2016, Juniper Networks, Inc.

Table 589: show security ipsec security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
Sta	State has two options, **Installed** and **Not Installed**. • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.
Mon	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays **U** (up) or **D** (down). A hyphen (-) means VPN monitoring is not enabled for this SA.
vsys or Virtual-system	The root system.
Tunnel index	Numeric identifier of the specific IPsec tunnel for the SA.
Local gateway	Gateway address of the local system.
Remote gateway	Gateway address of the remote system.
Traffic selector	Name of the traffic selector.
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote identity	IP address of the destination peer gateway.
DF-bit	State of the don't fragment bit: **set** or **cleared**.
Policy-name	Name of the applicable policy.
Location	**FPC**—Flexible PIC Concentrator (FPC) slot number. **PIC**—PIC slot number. **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance.
Tunnel events	Tunnel event and the number of times the event has occurred. See "Tunnel Events" on page 6797 for descriptions of tunnel events and the action you can take.
Direction	Direction of the SA; it can be inbound or outbound.

Copyright © 2016, Juniper Networks, Inc.

Table 589: show security ipsec security-associations *(continued)*

Field Name	Field Description
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is **AH** or **ESP**, **AUX-SPI** is always 0. • When the value is **AH+ESP**, **AUX-SPI** is always a positive integer.
Mode	Mode of the SA: • **transport**—Protects host-to-host connections. • **tunnel—**Protects connections between security gateways.
Type	Type of the SA: • **manual**—Security parameters require no negotiation. They are static and are configured by the user. • **dynamic**—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.
State	State of the SA: • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.

Protocol	Protocol supported.
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  ◦ **Authentication**—Type of authentication used.
	  ◦ **Encryption**—Type of encryption used.

Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in seconds**—Number of seconds left until the SA expires.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • **Expires in seconds**—Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • **Expires in kilobytes**—Number of kilobytes left until the SA expires.
Anti-replay service	State of the service that prevents packets from being replayed. It can be **Enabled** or **Disabled**.

Copyright © 2016, Juniper Networks, Inc.

Table 589: show security ipsec security-associations *(continued)*

Field Name	Field Description
Replay window size	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Copy-Outer-DSCP	Indicates if copying outer IP header DSCP and ECN to inner IP header is enabled or disabled.

Sample Output

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys
  131075 11.0.28.241    500   ESP:3des/sha1   86758ff0 6918/ unlim   -   0
  131075 11.0.28.241    500   ESP:3des/sha1   3183ff26 6918/ unlim   -   0
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon vsys Port  Gateway
  131074 ESP:3des/sha1 14caf1d9 3597/ unlim   -   root 500   1212::1112
  131074 ESP:3des/sha1 9a4db486 3597/ unlim   -   root 500   1212::1112
```

show security ipsec security-associations index 131073

```
user@host> show security ipsec security-associations index 131073
  ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
```


Copyright © 2016, Juniper Networks, Inc.

```
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1774 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1151 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ipsec security-associations brief

```
user@host> show security ipsec security-associations brief
Total active tunnels: 2
ID Gateway Port Algorithm SPI Life:sec/kb Mon vsys
<16384 1.1.1.1 500 ESP:3des/sha1 af88baa 28795/unlim D 0
>16384 1.1.1.1 500 ESP:3des/sha1 f4e3e5f4 28795/unlim D 0
```

show security ipsec security-associations detail

```
user@host> show security ipsec security-associations detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Version: IKEv1
DF-bit: clear
, Copy-Outer-DSCP Enabled
Bind-interface: st0.99

Port: 500, Nego#: 8, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
Tunnel events:
Mon Oct 26 2015 22:27:50 -0700: IPSec SA rekey successfully completed (7 times)
Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:41:07 -0700: IKE SA negotiation successfully completed (1
times)
Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information
 updated (1 times)
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: inbound, SPI: 81ed9998, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled
```

Copyright © 2016, Juniper Networks, Inc.

```
, Replay window size: 64
Location: FPC 0, PIC 1, KMD-Instance 1
Direction: outbound, SPI: 80565248, AUX-SPI: 0
Hard lifetime: Expires in 2296 seconds
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1688 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```

show security ipsec security-associations family inet6

```
user@host> show security ipsec security-associations family inet6
  Virtual-system: root
  Local Gateway: 1212::1111, Remote Gateway: 1212::1112
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
    DF-bit: clear
    Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64

    Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                              , VPN Monitoring: -
    Hard lifetime: Expires in 3440 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 2813 seconds
    Mode: tunnel, Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
  Total active tunnels: 1

ID    Gateway          Port  Algorithm       SPI      Life:sec/kb  Mon vsys

<2    1.1.1.2          500   ESP:3des/sha1   67a7d25d 28280/unlim   -   0

>2    1.1.1.2          500   ESP:3des/sha1   a23cbcdc 28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
  Local Gateway: 11.1.1.1, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear
  Bind-interface: st0.1

  Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
```


Copyright © 2016, Juniper Networks, Inc.

```
  Tunnel events:
  Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
  Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding
IKE/IPSec SAs are deleted (1 times)
  Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 3, KMD-Instance 2
  Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 1335 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 996 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
```

show security ike sa index 222075191 detail

```
user@host> show security ike sa index 222075191 detail
node0:
--------------------------------------------------------------------------
IKE peer 31.1.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent    :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 11.1.1.1:500, Remote: 31.1.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
```

```
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                20474
   Output bytes  :                21091
   Input  packets:                  237
   Output packets:                  237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 11.1.1.1:500, Remote: 31.1.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 31.1.1.2, Remote Gateway: 11.1.1.1
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local
certificate. Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate
 not found. Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```


Copyright © 2016, Juniper Networks, Inc.

show security ike sa index 788674 detail

```
user@host> show security ike sa index 788674 detail
IKE peer 11.1.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 31.1.1.2:500, Remote: 11.1.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                22535
   Output bytes  :                21918
   Input  packets:                  256
   Output packets:                  256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 31.1.1.2:500, Remote: 11.1.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=Example,
OU=engineering, CN=test
    Flags: IKE SA is created
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID    Algorithm       SPI      Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes-cbc-256/sha1 6f164ee0 3580/ unlim - root 500 23.0.0.111
>268173318 ESP:aes-cbc-256/sha1 e6f29cb0 3580/ unlim - root 500 23.0.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 21.1.1.2, Remote Gateway: 31.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
```

Copyright © 2016, Juniper Networks, Inc.

```
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
  Tunnel events:
    Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed
(1 times)
    Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
    Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 62.1.1.1, Remote Gateway: 62.1.1.2
  Local Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Remote Identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
```


Copyright © 2016, Juniper Networks, Inc.

```
Lifesize Remaining:  Unlimited
Soft lifetime: Expires in 1090 seconds
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```

Copyright © 2016, Juniper Networks, Inc.

show security ipsec statistics

Syntax

show security ipsec statistics
<**fpc** *slot-number* >
<**index** *SA-index-number* >
<**kmd-instance** *kmd-instance-name* >
pic *slot-number*

Release Information

Command introduced in Junos OS Release 8.5. **fpc** and **pic** options added in Junos OS Release 9.3. **kmd-instance** option added in Junos OS Release 10.4.

Description

Display standard IPsec statistics.

Options

- none—Display statistics about all IPsec security associations (SAs).
- **fpc** ***slot-number*** —Specific to SRX Series devices. Display statistics about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.
- **index** ***SA-index-number*** —(Optional) Display statistics for the SA with this index number.
- **kmd-instance** ***kmd-instance-name*** —Specific to SRX Series devices. Display information about existing IKE SAs in the key management process (the daemon, which in this case is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.
  - **all**—All KMD instances running on the Services Processing Unit (SPU).
  - ***kmd-instance-name***—Name of the KMD instance running on the SPU.
- **pic** ***slot-number*** —Specific to SRX Series devices. Display statistics about existing IPsec SAs in this PIC slot. This option is used to filter the output.

Required Privilege Level

view

Related Documentation

- clear security ipsec statistics on page 6987

List of Sample Output

show security ipsec statistics on page 7075
show security ipsec statistics index 5 on page 7076
show security ipsec statistics fpc 6 pic 1 (SRX Series devices) on page 7076

Output Fields

Table 379 lists the output fields for the **show security ipsec statistics** command. Output fields are listed in the approximate order in which they appear.

Table 590: show security ipsec statistics Output Fields

Field Name	Field Description
Virtual-system	The root system.

Copyright © 2016, Juniper Networks, Inc.

Table 590: show security ipsec statistics Output Fields *(continued)*

Field Name	Field Description
ESP Statistics	• **Encrypted bytes**—Total number of bytes encrypted by the local system across the IPsec tunnel. • **Decrypted bytes**—Total number of bytes decrypted by the local system across the IPsec tunnel. • **Encrypted packets**—Total number of packets encrypted by the local system across the IPsec tunnel. • **Decrypted packets**—Total number of packets decrypted by the local system across the IPsec tunnel.
AH Statistics	• **Input bytes**—Total number of bytes received by the local system across the IPsec tunnel. • **Output bytes**—Total number of bytes transmitted by the local system across the IPsec tunnel. • **Input packets**—Total number of packets received by the local system across the IPsec tunnel. • **Output packets**—Total number of packets transmitted by the local system across the IPsec tunnel.
Errors	• **AH authentication failures**—Total number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • **Replay errors**—Total number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • **ESP authentication failures**—Total number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • **ESP decryption failures**—total number of ESP decryption errors. • **Bad headers**—Total number of invalid headers detected. • **Bad trailers**—Total number of invalid trailers detected.

Sample Output

show security ipsec statistics

```
user@host> show security ipsec statistics
Virtual-system: Root
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security ipsec statistics index 5

```
user@host> show security ipsec statistics index 5
Virtual-system: Root
SA index: 5
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:              0
  Decrypted packets:              0
AH Statistics:
  Input bytes:                    0
  Output bytes:                   0
  Input packets:                  0
  Output packets:                 0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, ESP decryption failures: 0
  Bad headers: 0, Bad trailers: 0
```

Sample Output

show security ipsec statistics fpc 6 pic 1 (SRX Series devices)

```
user@host> show security ipsec statistics fpc 6 pic 1
ESP Statistics:
Encrypted bytes:           536408
Decrypted bytes:           696696
Encrypted packets:           1246
Decrypted packets:            888
AH Statistics:
Input bytes:                    0
Output bytes:                   0
Input packets:                  0
Output packets:                 0
Errors:
AH authentication failures: 0, Replay errors: 0
ESP authentication failures: 0, ESP decryption failures: 0
Bad headers: 0, Bad trailers: 0
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec traffic-selector

Syntax

show security ipsec traffic-selector interface-name *interface-name*
<**brief | detail**>
<**destination-address** *address*>
<**fpc** *slot-number*>
<**kmd-instance (all |** *kmd-instance-name*)>
<**pic** *slot-number*>
<**source-address** *address*>

Release Information

Command introduced in Junos OS Release 12.3X48-D10.

Description

Display information about the traffic selectors that have been negotiated between the initiator and responder.

Options

interface-name *interface-name*—Name of the secure tunnel logical interface.

brief | detail —(Optional) Display the specified level of output.

destination-address *address*—(Optional) Destination IP address.

fpc *slot-number*—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.

kmd-instance—(Optional) Display information about existing traffic selectors in the key management process (in this case, it is KMD) identified by FPC slot-number and PIC slot-number. This option is used to filter the output.

- **all**—All KMD instances running on the Services Processing Unit (SPU).
- *kmd-instance-name*—Name of the KMD instance running on the SPU.

pic *slot-number*—(Optional) Display information about existing traffic selectors in this PIC slot. This option is used to filter the output.

source-address *address*—(Optional) Source IP address.

Required Privilege Level

view

Related Documentation

- IPsec VPN Overview on page 6041

List of Sample Output

show security ipsec traffic-selector interface-name st0.1 on page 7078
show security ipsec traffic-selector interface-name st0.1 detail on page 7078

Output Fields

Table 591 lists the output fields for the **show security ipsec traffic-selector** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2016, Juniper Networks, Inc.

Table 591: show security ipsec traffic-selector Output Fields

Field Name	Field Description
Source IP	Source IP address for the negotiated traffic selector.
Destination IP	Destination IP address for the negotiated traffic selector.
Interface	Secure tunnel (st0) interface for the traffic selector.
Tunnel-id	Tunnel ID.
IKE-ID	Peer IKE ID for the negotiated traffic selector.

Sample Output

show security ipsec traffic-selector interface-name st0.1

```
user@host> show security ipsec traffic-selector interface-name st0.1
 Source IP                          Destination IP                     Interface
Tunnel-id    IKE-ID
 80.1.1.0-80.1.1.255                30.1.1.0-30.1.1.255                st0.1
69206018     DC=Common_component, CN=enodeA, OU=Dept, O=Company, L=City, ST=CA,
C=US
 80.1.1.0-80.1.1.255                50.1.1.0-50.1.1.255                st0.1
77594626     DC=Common_component, CN=enodeB, OU=Det, O=Company, L=City, ST=CA,
C=US
```

show security ipsec traffic-selector interface-name st0.1 detail

```
user@host> show security ipsec traffic-selector interface-name st0.1 detail
 Source IP                                        Destination IP
 Interface   Tunnel-id    IKE-ID
 80.0.0.0-80.0.0.255                30.0.0.0-30.0.0.255                st0.1
208666625    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia1
 80.0.1.0-80.0.1.255                30.0.1.0-30.0.1.255                st0.1
213909505    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia2
 80.0.2.0-80.0.2.255                30.0.2.0-30.0.2.255                st0.1
214958081    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia3
 80.0.3.0-80.0.3.255                30.0.3.0-30.0.3.255                st0.1
216006657    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia4
 80.0.4.0-80.0.4.255                30.0.4.0-30.0.4.255                st0.1
217055233    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia5
 80.0.5.0-80.0.5.255                30.0.5.0-30.0.5.255                st0.1
218103809    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia6
 80.0.6.0-80.0.6.255                30.0.6.0-30.0.6.255                st0.1
219152385    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia7
 80.0.7.0-80.0.7.255                30.0.7.0-30.0.7.255                st0.1
220200961    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia8
 80.0.8.0-80.0.8.255                30.0.8.0-30.0.8.255                st0.1
221249537    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia9
 80.0.9.0-80.0.9.255                30.0.9.0-30.0.9.255                st0.1
222298113    C=US, ST=CA, L=City, O=Ixia, OU=IxLoad, CN=ixia10
```


Copyright © 2016, Juniper Networks, Inc.

show security ipsec tunnel-events-statistics

Syntax **show security ipsec tunnel-events-statistics**

Release Information Command introduced in Junos OS Release 12.3X48-D10.

Description Show tunnel event statistics.

Required Privilege Level view

Related Documentation
- *clear security ipsec tunnel-events-statistics*

List of Sample Output show security ipsec tunnel-events statistics on page 7079

Sample Output

show security ipsec tunnel-events statistics

```
user@host> show security ipsec tunnel-events statistics
IPSec SA delete payload received from peer                                : 153
Configuration change triggered clearing of IPSec SA                       : 1
Peer's remote IKE-ID validation failed during negotiation                 : 2
Phase1 proposal mismatch detected                                         : 2
Phase2 proposal mismatch detected                                         : 2
Peer proposed traffic-selectors are not in configured range              : 8576
Negotiation failed as peer did not respond                                : 4
IKE SA negotiation successfully completed                                 : 19
IPSec SA negotiation successfully completed                              : 154
Tunnel is ready. Waiting for trigger event or peer to trigger negotiation  : 1
```

Copyright © 2016, Juniper Networks, Inc.

show security pki ca-certificate (View)

Syntax

show security pki ca-certificate
<**brief** | **detail**>
<**ca-profile** *ca-profile-name* >

Release Information

Command modified in Junos OS Release 8.5. Subject string output field added in Junos OS Release 12.1X44-D10. Policy identifier output field added in Junos OS Release 12.3X48-D10.

Description

Display information about the certificate authority (CA) public key infrastructure (PKI) digital certificates configured on the device.

NOTE: The FIPS image does not permit the use of MD5 fingerprints. Therefore, MD5 fingerprints are not included when a certificate is displayed using this command. The SHA-1 fingerprint that is currently displayed is retained in the FIPS image. The Simple Certificate Enrollment Protocol (SCEP) is disabled in the FIPS image.

Options

- none—Display basic information about all configured CA certificates.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **ca-profile** *ca-profile-name*- (Optional) Display information about only the specified CA certificate.

Required Privilege Level

view

Related Documentation

- ca-profile (Security PKI) on page 6847
- *request security pki ca-certificate verify (Security)*

List of Sample Output

show security pki ca-certificate ca-profile RootCA brief on page 7081
show security pki ca-certificate ca-profile RootCA detail on page 7082
show security pki ca-certificate ca-profile ca-tmp detail on page 7082

Output Fields

Table 592 lists the output fields for the **show security pki ca-certificate** command. Output fields are listed in the approximate order in which they appear.

Table 592: show security pki ca-certificate Output Fields

Field Name	Field Description
Certificate identifier	Name of the digital certificate.
Certificate version	Revision number of the digital certificate.
Serial number	Unique serial number of the digital certificate.

Copyright © 2016, Juniper Networks, Inc.

Table 592: show security pki ca-certificate Output Fields *(continued)*

Field Name	Field Description
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority.
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority. If the certificate contains multiple subfield entries, all entries are displayed.
Subject string	Subject field as it appears in the certificate.
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption**(**1024 bits**).
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.
Certificate Policy	**Policy Identifier**—One or more policy object identifiers (OIDs).
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Data encipherment**.
Fingerprint	Secure Hash Algorithm (**SHA1**) and Message Digest 5 (**MD5**) hashes used to identify the digital certificate.
Distribution CRL	Distinguished name information and the URL for the certificate revocation list (CRL) server.

Sample Output

show security pki ca-certificate ca-profile RootCA brief

```
user@host> show security pki ca-certificate ca-profile RootCA brief
Certificate identifier: RootCA
  Issued to: RootCA, Issued by: C = US, O = Example, CN = RootCA
```

Copyright © 2016, Juniper Networks, Inc.

```
  Validity:
    Not before: 05- 3-2012 07:15
    Not after: 05- 2-2017 07:15
  Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki ca-certificate ca-profile RootCA detail

```
user@host> show security pki ca-certificate ca-profile RootCA detail
Certificate identifier: RootCA
  Certificate version: 3
  Serial number: 0712dc31
  Issuer:
    Organization: Example, Country: US, Common name: RootCA
  Subject:
    Organization: Example, Country: US, Common name: RootCA
  Subject string:
    C=US, O=Example, CN=RootCA
  Validity:
    Not before: 05- 3-2012 07:15
    Not after: 05- 2-2017 07:15
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:ac:b0:c0:11:ac:0c:34:37:04:97:65:c2:b1
    ae:7e:68:e0:fa:37:23:a1:f0:eb:4d:eb:03:89:c9:d9:0d:34:f3:66
    91:97:8c:e9:9c:d4:b5:55:8d:c1:e2:8b:95:08:9d:29:f8:ab:ac:ff
    ae:af:f7:bc:4b:33:f2:eb:b9:e6:13:6d:18:d7:64:a7:85:78:99:41
    4e:b4:fa:bc:3e:1b:5c:26:25:89:03:af:e9:c6:e9:9e:7b:74:1a:1a
    5b:b4:2a:48:78:57:68:e2:5c:0b:71:71:78:ac:a2:23:5f:ca:d2:4a
    38:4c:35:5a:20:cc:44:39:96:26:20:43:bd:75:fd:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Use for key: CRL signing, Certificate signing, Key encipherment,
  Digital signature
  Fingerprint:
    eb:2a:2a:eb:d3:c7:cb:62:65:2e:6a:76:56:b8:af:88:51:8a:30:c9 (sha1)
    cd:43:ae:a4:b2:11:9e:cf:1a:47:fd:7f:0c:ce:d9:fd (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Sample Output

show security pki ca-certificate ca-profile ca-tmp detail

```
user@host> show security pki ca-certificate ca-profile ca-tmp detail
Certificate identifier: ca-tmp
  Certificate version: 3
  Serial number: 00000047
  Issuer:
    Organization: U.S. Government,
    Organizational unit: DoD, Organizational unit: Testing, Country: US,
    Common name: Trust Anchor
  Subject:
    Organization: U.S. Government,
    Organizational unit: Dod, Organizational unit: Testing, Country: US,
    Common name: test
  Subject string:
    C=US, O=U.S. Government, OU=Dod, OU=Testing, CN=test
  Validity:
    Not before: 01- 1-1998 12:01 UTC
    Not after: 01- 1-2048 12:01 UTC
```


Copyright © 2016, Juniper Networks, Inc.

```
Public key algorithm: rsaEncryption(1024 bits)
  30:81:89:02:81:81:00:cb:fd:78:0c:be:87:ac:cd:c0:33:66:a3:18
  9e:fd:40:b7:9b:bc:dc:66:ff:08:45:f7:7e:fe:8e:d6:32:f8:5b:75
  db:76:f0:4d:21:9a:6e:4f:04:21:4c:7e:08:a1:f9:3d:ac:8b:90:76
  44:7b:c4:e9:9b:93:80:2a:64:83:6e:6a:cd:d8:d4:23:dd:ce:cb:3b
  b5:ea:da:2b:40:8d:ad:a9:4d:97:58:cf:60:af:82:94:30:47:b7:7d
  88:c3:76:c0:97:b4:6a:59:7e:f7:86:5d:d8:1f:af:fb:72:f1:b8:5c
  2a:35:1e:a7:9e:14:51:d4:19:ae:c7:5c:65:ea:f5:02:03:01:00:01
Signature algorithm: sha1WithRSAEncryption
Certificate Policy:
  Policy Identifier = 2.16.840.1.101.3.1.48.2
Use for key: CRL signing, Certificate signing
Fingerprint:
  e0:b3:2f:2e:a1:c5:ee:ad:af:dd:96:85:f6:78:24:c5:89:ed:39:40 (sha1)
  f3:47:6e:55:bc:9d:80:39:5a:40:70:8b:10:0e:93:c5 (md5)
```

Copyright © 2016, Juniper Networks, Inc.

show security pki certificate-request (View)

Syntax
show security pki certificate-request
<**brief** | **detail**>
<certificate-id *certificate-id-name* >

Release Information
Command modified in Junos OS Release 8.5.

Description
Display information about manually generated local digital certificate requests that are stored on the device.

Options
- none—Display basic information about all local digital certificate requests.
- **brief | detail**—(Optional) Display the specified level of output.
- **certificate-id** ***certificate-id-name*** —(Optional) Display information about only the specified local digital certificate requests.

Required Privilege Level
view

Related Documentation
- clear security pki key-pair (Local Certificate) on page 6989

List of Sample Output
show security pki certificate-request certificate-id test brief on page 7085
show security pki certificate-request certificate-id test detail on page 7085

Output Fields
Table 593 lists the output fields for the **show security pki certificate-request** command. Output fields are listed in the approximate order in which they appear.

Table 593: show security pki certificate-request Output Fields

Field Name	Field Description
Certificate identifier	Name of the digital certificate.
Certificate version	Revision number of the digital certificate.
Issued to	Device that was issued the digital certificate.
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority.
Alternate subject	Domain name or IP address of the device related to the digital certificate.
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption**(**1024 bits**).

Copyright © 2016, Juniper Networks, Inc.

Table 593: show security pki certificate-request Output Fields *(continued)*

Field Name	Field Description
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.
Fingerprint	Secure Hash Algorithm (**SHA1**) and Message Digest 5 (**MD5**) hashes used to identify the digital certificate.
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Data encipherment**.

Sample Output

show security pki certificate-request certificate-id test brief

```
user@host> show security pki certificate-request certificate-id test brief
Certificate identifier: test
      Issued to: test@example.net
      Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki certificate-request certificate-id test detail

```
user@host> show security pki certificate-request certificate-id test detail
    Certificate identifier: test
      Certificate version: 3
      Subject:
        Organization: example, Organizational unit: example, Country: IN,
                      Common name: test
      Alternate subject: 102.168.72.124
      Public key algorithm: rsaEncryption(1024 bits)
      Public key verification status: Passed
        c7:a4:fb:e7:8c:4f:31:e7:eb:01:d8:32:65:21:f2:eb:6f:7d:49:1a:c3:9b
        63:47:e2:4f:f6:db:f6:c8:75:dd:e6:ec:0b:35:0a:62:32:45:6b:35:1f:65
        c9:66:b7:40:b2:f9:2a:ab:5b:60:f7:c7:73:36:da:68:25:fc:40:4b:12:3c
        d5:c8:c6:66:f6:10:1e:86:67:a8:95:9b:7f:1c:ae:a7:55:b0:28:95:a7:9a
        a2:24:28:e4:5a:b2:a9:06:7a:69:37:20:15:e1:b6:66:eb:22:b5:b6:77:f6
        65:88:b0:94:2b:91:4b:99:78:4a:e3:56:cc:14:45:d7:97:fd
      Fingerprint:
        8f:22:1a:f2:9f:27:b0:21:6c:da:46:64:31:34:1f:68:42:5a:39:e0 (sha1)
        09:15:11:aa:ea:f9:5a:b5:70:d7:0b:8e:be:a6:d3:cb (md5)
      Use for key: Digital signature
```

Copyright © 2016, Juniper Networks, Inc.

show security pki crl (View)

Syntax
show security pki crl
< **brief** | **detail**>
<ca-profile *ca-profile-name* >

Release Information
Command modified in Junos OS Release 8.5.

Description
Display information about the certificate revocation lists (CRLs) configured on the device.

Options
- none—Display basic information about all CRLs.
- **brief** | **detail**—(Optional) Display the specified level of output.
- **ca-profile** *ca-profile-name*- (Optional) Display information about only the specified CA profile.

Required Privilege Level
view

Related Documentation
- crl (Security) on page 6852

List of Sample Output
show security pki crl ca-profile ca2 on page 7087
show security pki crl ca-profile ca2 brief on page 7087
show security pki crl ca-profile ca2 detail on page 7087

Output Fields
Table 594 lists the output fields for the **show security pki crl** command. Output fields are listed in the approximate order in which they appear.

Table 594: show security pki crl Output Fields

Field Name	Field Description
CA profile	Name of the configured CA profile.
CRL version	Revision number of the certificate revocation list.
CRL issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **emailAddress**—Mail address of the issuing authority. • **C**—Country of origin. • **ST**—State of origin. • **L**—Locality of origin. • **O**—Organization of origin. • **OU**—Department within an organization. • **CN**—Name of the authority.
Effective date	Date and time the certificate revocation list becomes valid.

Copyright © 2016, Juniper Networks, Inc.

Table 594: show security pki crl Output Fields *(continued)*

Field Name	Field Description
Next update	Date and time the routing platform will download the latest version of the certificate revocation list.
Revocation List	List of digital certificates that have been revoked before their expiration date. Values are: • **Serial number**—Unique serial number of the digital certificate. • **Revocation date**—Date and time that the digital certificate was revoked.

Sample Output

show security pki crl ca-profile ca2

```
user@host> show security pki crl ca-profile ca2
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = test@example.net, C = US, ST = ca, L = sunnyvale, O
 =  Example, OU = SPG QA, CN = test-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
```

Sample Output

show security pki crl ca-profile ca2 brief

```
user@host> show security pki crl ca-profile ca2 brief
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = test@example.net, C = US, ST = ca, L = sunnyvale, O
 = Example, OU = SPG QA, CN = test-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
```

Sample Output

show security pki crl ca-profile ca2 detail

```
user@host> show security pki crl ca-profile ca2 detail
CA profile: ca2
  CRL version: V00000001
  CRL issuer: emailAddress = test@example.net, C = US, ST = ca, L = sunnyvale, O
 = Example, OU = SPG QA, CN = test-example-net
  Effective date: 04-26-2007 18:47
  Next update: 05- 4-2007 07:07
  Revocation List:
    Serial number              Revocation date
    174e6399000000000506       03-16-2007 23:09
    174ef3f3000000000507       03-16-2007 23:09
    17529cd6000000000508       03-16-2007 23:09
    1763ac26000000000509       03-16-2007 23:09
    21904e5700000000050a       03-16-2007 23:09
    2191cf7900000000050b       03-16-2007 23:09
    21f10eb600000000050c       03-16-2007 23:09
    2253ca2a00000000050f       03-16-2007 23:09
    2478939b000000000515       03-16-2007 23:09
```

Copyright © 2016, Juniper Networks, Inc.

```
24f35004000000000516        03-16-2007 23:09
277ddfa8000000000517        03-16-2007 23:09
277e97bd000000000518        03-16-2007 23:09
27846a76000000000519        03-16-2007 23:09
2785176f00000000051a        03-16-2007 23:09
```


Copyright © 2016, Juniper Networks, Inc.

show security pki local-certificate (View)

Syntax

show security pki local-certificate
< **brief** | **detail** >
< certificate-id ***certificate-id-name*** >
<system-generated>

Release Information

Command modified in Junos OS Release 9.1. Subject string output field added in Junos OS Release 12.1X44-D10.

Description

Display information about the local digital certificates, corresponding public keys, and the automatically generated self-signed certificate configured on the device.

Options

- none—Display basic information about all configured local digital certificates, corresponding public keys, and the automatically generated self-signed certificate.
- **brief** | **detail**—(Optional) Display the specified level of output.
- certificate-id ***certificate-id-name*** —(Optional) Display information about only the specified local digital certificates and corresponding public keys.
- **system-generated**—Display information about the automatically generated self-signed certificate.

Required Privilege Level

view

Related Documentation

- clear security pki local-certificate (Device) on page 6990
- request security pki local-certificate generate-self-signed (Security) on page 7005

List of Sample Output

show security pki local-certificate certificate-id hello on page 7091
show security pki local-certificate certificate-id hello detail on page 7091
show security pki local-certificate system-generated on page 7092
show security pki local-certificate system-generated detail on page 7092
show security pki local-certificate certificate-id mycert - (local certificate enrolled online using SCEP) on page 7092
show security pki local-certificate certificate-id mycert detail - (local certificate enrolled online using SCEP) on page 7093

Output Fields

Table 50 lists the output fields for the **show security pki local-certificate** command. Output fields are listed in the approximate order in which they appear.

Table 595: show security pki local-certificate Output Fields

Field Name	Field Description
Certificate identifier	Name of the digital certificate.
Certificate version	Revision number of the digital certificate.
Serial number	Unique serial number of the digital certificate.

Copyright © 2016, Juniper Networks, Inc.

Table 595: show security pki local-certificate Output Fields *(continued)*

Field Name	Field Description
Issued to	Device that was issued the digital certificate.
Issued by	Authority that issued the digital certificate.
Issuer	Authority that issued the digital certificate, including details of the authority organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority.
Subject	Details of the digital certificate holder organized using the distinguished name format. Possible subfields are: • **Organization**—Organization of origin. • **Organizational unit**—Department within an organization. • **Country**—Country of origin. • **Locality**—Locality of origin. • **Common name**—Name of the authority. • **Serial number**—Serial number of the device. If the certificate contains multiple subfield entries, all entries are displayed.
Subject string	Subject field as it appears in the certificate.
Alternate subject	Domain name or IP address of the device related to the digital certificate.
Validity	Time period when the digital certificate is valid. Values are: • **Not before**—Start time when the digital certificate becomes valid. • **Not after**—End time when the digital certificate becomes invalid.
Public key algorithm	Encryption algorithm used with the private key, such as **rsaEncryption**(**1024 bits**).
Public key verification status	Public key verification status: **Failed** or **Passed**. The **detail** output also provides the verification hash.
Signature algorithm	Encryption algorithm that the CA used to sign the digital certificate, such as **sha1WithRSAEncryption**.
Fingerprint	Secure Hash Algorithm (**SHA1**) and Message Digest 5 (**MD5**) hashes used to identify the digital certificate.
Distribution CRL	Distinguished name information and URL for the certificate revocation list (**CRL**) server.
Use for key	Use of the public key, such as **Certificate signing**, **CRL signing**, **Digital signature**, or **Data encipherment**.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate certificate-id hello

```
user@host> show security pki local-certificate certificate-id hello
Certificate identifier: hello
  Issued to: cn1, Issued by: DC = local, DC = demo, CN = example-CN1
  Validity:
    Not before: 08- 8-2012 17:02
    Not after: 08- 8-2014 17:02
  Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki local-certificate certificate-id hello detail

```
user@host> show security pki local-certificate certificate-id hello detail
Certificate identifier: hello
  Certificate version: 3
  Serial number: 61ba9da000000000d72e
  Issuer:
    Common name: example-cn1,
    Domain component: local, Domain component: demo
  Subject:
    Organization: o1, Organization: o2,
    Organizational unit: ou1, Organizational unit: ou2, Country: US, State: CA,
    Locality: Sunnyvale, Common name: cn1, Common name: cn2,
    Domain component: dc1, Domain component: dc2
  Subject string:
    C=US, DC=dc1, DC=dc2, ST=CA, L=Sunnyvale, O=o1, O=o2, OU=ou1, OU=ou2, CN=cn1,
 CN=cn2
  Alternate subject: "user@example.net", user.example.net, 10.1.2.3
  Validity:
    Not before: 08- 8-2012 17:02
    Not after: 08- 8-2014 17:02
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:b4:14:01:d5:4f:79:87:d5:bb:e6:5e:c1:14
    97:da:b4:40:ad:1a:77:3e:ec:2e:68:8e:e4:93:a3:fe:7c:0b:58:af
    e1:20:27:82:ca:8d:6f:f0:97:d1:ad:fe:df:6c:cb:3c:b0:4f:cc:dd
    ac:d8:69:3f:3c:59:b5:2a:c6:83:e8:b3:94:5e:0a:2d:cd:e2:b0:15
    3e:97:a7:8a:4e:fb:59:f7:20:4c:ba:a8:80:3e:ba:be:69:ef:2b:32
    e4:1a:1c:24:53:1b:d5:c3:aa:d4:25:73:96:76:ea:49:d4:da:7e:3e
    0c:c6:6b:22:43:cb:04:84:0d:25:33:07:6b:49:41:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    ldap:///CN=example-cn1,CN=example,CN=CDP,CN=Public%20Key
%20Services,CN=Services,CN=Configuration,DC=demo,DC=local?certificateRevocationList?base?
objectClass=cRLDistributionPoint
    http://user.device.example.net/CertEnroll/device-user.crl
  Use for key: Key encipherment, Digital signature, 1.3.6.1.5.5.8.2.2,
  1.3.6.1.5.5.8.2.2
  Fingerprint:
    76:a8:5f:65:b4:bf:bd:10:d8:56:82:65:ff:0d:04:3a:a5:e9:41:dd (sha1)
    8f:99:a4:15:98:10:4b:b6:1a:3d:81:13:93:2a:ac:e7 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate system-generated

```
user@host> show security pki local-certificate system-generated
Certificate identifier: system-generated
  Issued to: JN1XXXXXX, Issued by: CN = JN10BXXXXX, CN = system generated, CN =
self-signed
  Validity:
    Not before: 10-30-2009 23:02
    Not after: 10-29-2014 23:02
  Public key algorithm: rsaEncryption(1024 bits)
```

Sample Output

show security pki local-certificate system-generated detail

```
user@host> show security pki local-certificate system-generated detail
Certificate identifier: system-generated
  Certificate version: 3
  Serial number: a1bc122bcaaaabbc1
  Issuer:
    Common name: JN1XXXXXXX, Common name: system generated, Common name:
self-signed
  Subject:
    Common name: JN1XXXXXXX, Common name: system generated, Common name:
self-signed
  Subject string:
    CN=JN10B9390AGB, CN=system generated, CN=self-signed
  Validity:
    Not before: 10-30-2009 23:02
    Not after: 10-29-2014 23:02
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:cb:c8:3f:e6:d3:e5:ca:9d:dc:2d:e9:ca:c7
    5f:b1:f5:3a:f0:1c:a7:55:43:0f:ef:fd:1c:fe:29:09:d5:37:d0:fa
    d6:ee:bc:b8:3f:58:d4:31:fb:96:4f:4f:cc:a9:1a:8f:2e:1b:50:6f
    2b:88:34:74:b2:6d:ad:94:b5:dd:3d:80:87:56:d0:42:50:4d:ac:d7
    8c:21:06:2d:07:1e:f4:d0:c7:85:2e:25:60:ad:1b:b5:b2:d2:1d:c8
    79:67:8c:56:06:04:75:6e:be:4e:99:b8:07:e6:9a:11:fe:b5:ec:c0
    1e:68:da:47:99:1b:b2:c8:07:ab:cd:6e:fe:c1:fd:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Fingerprint:
    be:1f:21:13:71:cd:9d:de:7a:41:d7:4c:52:8d:3e:d6:ba:db:75:96 (sha1)
    ba:fc:90:4b:5f:a8:66:a3:b9:64:89:9f:e2:45:b5:84 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Sample Output

show security pki local-certificate certificate-id mycert - (local certificate enrolled online using SCEP)

```
user@host> show security pki local-certificate certificate-id mycert
Certificate identifier: mycert
  Issued to: user1, Issued by: DC = local, DC = demo, CN = example-cn1
  Validity:
    Not before: 11-15-2012 18:58
    Not after: 11-15-2014 18:58
  Public key algorithm: rsaEncryption(1024 bits)
```


Copyright © 2016, Juniper Networks, Inc.

Sample Output

show security pki local-certificate certificate-id mycert detail - (local certificate enrolled online using SCEP)

```
user@host> show security pki local-certificate certificate-id mycert detail
Certificate identifier: mycert
  Certificate version: 3
  Serial number: 1f00b50a0000000XXXX
  Issuer:
    Common name: example-cn1,
    Domain component: local, Domain component: demo
  Subject:
    Organization: example, Organizational unit: SSD, Country: US,
    Common name: user1, Serial number: SRX240-11152012
  Subject string:
    serialNumber=SRX240-11152012, C=US, O=example, OU=SSD, CN=user
  Alternate subject: "user@example.net", user.example.net, 10.150.1.2
  Validity:
    Not before: 11-15-2012 18:58
    Not after: 11-15-2014 18:58
  Public key algorithm: rsaEncryption(1024 bits)
    30:81:89:02:81:81:00:e3:e5:ae:c0:82:af:db:94:01:2f:56:46:50
    7d:3d:0b:0c:f0:1f:1d:7d:c3:aa:d4:4c:a0:cd:23:8b:3f:47:05:ee
    7b:65:42:a0:dc:c4:ac:a7:b6:a6:9f:5c:ea:d8:22:b0:bf:03:75:09
    be:fa:77:cb:d6:67:19:e6:80:fa:a5:7c:93:af:96:66:9f:cc:45:d5
    eb:ab:c1:f0:32:a6:d9:27:1b:80:bb:57:ec:31:a2:e0:2b:e1:42:c0
    92:8a:9b:ed:a6:d2:ec:7c:84:5a:8a:d9:96:a7:7e:40:c3:80:0e:f4
    d6:a2:5d:78:93:3b:7d:d5:8a:f5:de:fb:bc:0d:6d:02:03:01:00:01
  Signature algorithm: sha1WithRSAEncryption
  Distribution CRL:
    ldap:///CN=example-cn1,CN=example,CN=CDP,CN=Public%20Key%20Services,
CN=Services,CN=Configuration,DC=demo,DC=local?certificateRevocationList?
base?objectClass=cRLDistributionPoint
    http://user.device.example.net/CertEnroll/device-user.crl
  Use for key: Key encipherment, Digital signature, 1.3.6.1.5.5.8.2.2,
  1.3.6.1.5.5.8.2.2
  Fingerprint:
    1f:2f:a9:22:a8:d5:a9:36:cc:c4:bd:81:59:9d:9c:58:bb:40:15:72 (sha1)
    51:27:e4:d5:29:90:f7:85:9e:67:84:a1:75:d1:5b:16 (md5)
  Auto-re-enrollment:
    Status: Disabled
    Next trigger time: Timer not started
```

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.